

DSO 435 – Enterprise Data Architecture

Time: Monday & Wednesdays 2:00 - 3:50 PM
Place: HOH 421
Professor: Douglas Shook, Ph.D.
Office hours: JHH 102
Wed. 4:00-5:00 p.m., Thu. 8:30-9:30 a.m., after class and by appointment
Phone: office: 740-4623
FAX: office: 821-0264
Email: shook@esd.usc.edu
Blackboard: <https://blackboard.usc.edu>

Overview of the Course:

Database architecture approaches have become the foundation of most business information systems. Efforts to support Data Analytics (DA) and Business Intelligence (BI), Customer Relationship Management (CRM), increased reliance on ERP packages (e.g., SAP, PeopleSoft, Baan, etc.), exposing business data to customers through web-based customer interfaces, compliance with Sarbanes-Oxley/HIPAA/FERPA, recognition of the importance of “Big Data,” synchronizing/integrating legacy data with Cloud-based applications, etc., continue to accelerate the need to move toward more comprehensive data integration and enterprise-wide management of data architecture.

This course is motivated by a recognition of the importance of enterprise data architectures for understanding, designing and building effective information systems for today's business organizations.

The objective of this course is to provide the prospective business or information systems professional with fundamental concepts and skills in data modeling (conceptual, logical and physical), as well as in designing, building and managing the data layer to support business applications. In addition to data modeling, considerable proficiency with Structured Query Language (SQL) will be obtained. Although object-oriented, and other database approaches will be discussed, the course focuses primarily on the use of state-of-the-art relational and object-relational databases. Oracle will be used as the primary pedagogical software in class and for required assignments.

Prerequisites: None

Required Readings/Materials:

- Hoffer, J., Topi, H & Ramesh, V. (2014). *Essentials of Database Management*. Pearson (ISBN-13: 978-0-13-340568-2)
- Additional articles/handouts to be distributed in class.
- Data processing template (e.g., Staedtler 977/112 available at bookstore), or similar template.

Optional, Supplemental Readings:

- Loney, K. (2009). *Oracle 11g - The Complete Reference*, Osborne Oracle Press (ISBN: 0071598758).

On Reserve in Crocker Library (three hour checkout):

- Hoffer, J., Ramesh, V. & Topi, H. (2011). *Modern Database Management*, 10th Edition. Pearson-Prentice Hall: New Jersey
- Hoffer, J., Prescott, M. & Topi, H. (2009). *Modern Database Management*, 9th Edition. Pearson-Prentice Hall: New Jersey (2 copies)
- Loney, K. (2004). *Oracle 10g - The Complete Reference*, Osborne Oracle Press (ISBN: 0072253517).
- Earp, R. and Bahui, S. (2003). *Learning SQL: a step-by-step guide using Oracle*. Addison Wesley: New York.

Course Details

Grading:

Class Participation, In-Class Problem Solutions	10%
Oracle/SQL Programming Assignments	20%
Term Project	30%
In-Class Workshops (First @10%, Second and Third @15% each)	40%

(Extraordinary work in any of the above components will be rewarded accordingly with "grade overflow" into the other components. As per the 2014-2015 grading guidelines established by the Marshall School of Business, the average grade in this course will be 3.3.)

Class Participation:

Class participation be based upon class attendance, involvement in class discussions and responses to questions asked in class. Additionally, problems from the Hoffer text labeled "Problems and Exercises" will be assigned. Class participants will be chosen randomly to present their solutions on the board for class analysis and discussion. (individual work)

Oracle/SQL Problems:

Two hands-on programming assignments will be done using Oracle. You will turn in the computer generated output from your work. (individual work)

Workshops:

There will be three in-class workshops where you will perform data modeling (conceptual, logical, physical), as well as some SQL (DDL, DML, etc.). There will be no exams in this class. (individual work)

Project:

The term project will be the "cap stone" of the semester and requires the design of a database application and the implementation of this design using a database management system. The project will apply most of the issues/concepts covered during the semester and will enable you to obtain first-hand experience in designing and implementing a basic DBMS application from start to finish at an actual organizational site. It will be your responsibility to find a suitable project.

In the past, students have done projects in the organization where they worked, for some department in the university, local retailers, businesses they frequent, churches, social organizations, governmental agencies, or for their friends' or relatives' businesses. I would encourage you to develop an application for a not-for-profit, or similar organization, which may be able to afford to undertake such a project for a fee. The term project will result in a coded application, a high quality written report and an oral presentation. The project will be undertaken in the small teams (generally between two and four members) with a minimum of three to five entities per member. Additional details regarding project requirements will be provided later. (group work)

Class Administrivia:

Unfortunately, a certain amount of administrivia, or housekeeping, is required for any course. Although I dislike including this section in the syllabus, it is only fair to make the course expectations clearly known at the outset of the class (hopefully there will be no surprises down the road). Following are the "rules" and expectations of the class:

- It is expected you will attend all class meetings. If you must miss a class, please inform me beforehand. There will be no "make-ups" for any work missed due to un-excused absences (this includes workshops).
- You will be expected to have prepared for each class by reading the assigned chapters and handouts. If you do not keep up with the readings, you will neither enjoy, nor benefit from the class.
- Typically I will use a randomized class list to call upon students to present their solutions to the assigned "Problems and Exercises" at the end of each chapter.

General Comments

The content of this course is by nature somewhat narrow focused and somewhat technical. It is designed to provide the participant with a reasonably thorough understanding of database design and implementation in business organizations. It also provides the participant with first-hand exposure to the entire database design and implementation process in an actual organization.

Mastery of business data modeling, and data layer-related concepts, increasingly are becoming essential skills for any individual pursuing a career in Information Systems. This class will provide a foundation for professional development. It is assumed that the participants in this course wish to achieve some level of mastery over the database process.

The course assumes some previous knowledge of, and exposure to, information systems. This course is not intended to be a general survey course on information technologies (see DSO 431), nor is it an explicit course in systems analysis and design (see DSO 433).

Academic Integrity

The use of unauthorized material, communication with fellow students during an examination, attempting to benefit from the work of another student, and similar behavior that defeats the intent of an examination or other class work is unacceptable to the University. It is often difficult to distinguish between a culpable act and inadvertent behavior resulting from the nervous tensions accompanying examinations. Where a clear violation has occurred, however, the instructor may disqualify the student's work as unacceptable and assign a failing mark on the paper.

Violations of academic integrity will not be tolerated in this course. Violations of academic integrity include, but are not limited to:

1. **Plagiarism** - the submission of any material authored by another person, whether or not the original author gave permission to use the material.
2. **For individual assignments** - any inappropriate collaboration or copying of solutions. Students are encouraged to discuss assignments, to work together, and to share knowledge, as long as the exchange of information is related to understanding the material, and not for the purpose of sharing actual solutions to the assignment.
3. **For the workshops** - any communication of information, or reference to any sources of information, other than with the professor or the TA.
4. Any effort to uncover, and/or use, any materials from previous semesters, or from other sections, or from other students of this course (this includes data modeling cases and programming assignments).
5. Any violation of the academic integrity standards set forth in the student conduct code. Students are expected to be familiar with these standards, as well as being familiar with the ramifications of violating these standards.

In an effort to create a high integrity learning opportunity for all, and to ensure fairness in grading during the course of the semester, the professor and the teaching assistants will endeavor assiduously to uncover any violations of academic integrity. Depending upon the severity of the infraction, the professor will seek recourse up to, and including, expulsion from the university.

Students with Disabilities

Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me as early in the semester as possible. DSP is located in STU 301 and is open 8:30 a.m. - 5:00 p.m., Monday through Friday. The phone number for DSP is (213) 740-0776.

Classroom Recording

Various portions of this class may be captured/recorded for subsequent use by the students. The links to the recordings will be placed in Blackboard and only the students in this class should have access to them recorded. Please note that your voice and or image may be recorded in the process.

Course Schedule – Version 1.0

Date	Topic	Readings/Questions
Module 1 - Introduction to the Database Environment		
Aug. 25	Administrivia, Course Overview and Announcements	none / none
Aug. 27	The Database Environment I	Chap. 1 pp. 1-24 / none
Sep. 1	Labor Day (no class)	
Sep. 3	The Database Environment II & Team Composition	(same as previous session)
Sep. 8	Database Development Process	Chap. 1 pp. 24-43, handout/none
Module 2 - Conceptual Database Design		
Sep. 10	Conceptual Data Modeling/Entity Relationship Diagrams I	Chap. 2 & Appendix A / 5, 8, 15a-d, 18
Sep. 15	Conceptual Data Modeling/Entity Relationship Diagrams II	(same as previous session)
Sep. 17	Conceptual Data Modeling/Entity Relationship Diagrams III	(same as previous session)
Sep. 22	Conceptual Data Modeling/Entity Relationship Diagrams IV	(same as previous session)
Sep. 24	Advanced Conceptual Modeling I	Chap. 3 / 2, 7, 8, 12
Sep. 29	Advanced Conceptual Modeling II	(same as previous session)
Oct. 1	First Hands-on Workshop	(same as previous session)
Oct. 6	First Workshop Solutions, Discussion & Analysis	
Module 3 - Structured Query Language (SQL)		
Oct. 8	Introduction to Structured Query Language I	Chap. 6 & handout / 2-4, 7-10
Oct. 13	Introduction to Structured Query Language II	(same as previous session)
Oct. 15	Advanced Structured Query Language I	Chap. 7 & handout / 1-5
Oct. 20	Advanced Structured Query Language II	(same as previous session)
Module 4 - Logical Database Design		
Oct. 22	Logical Design and the Relational Model I	Chap. 4, Appendix B / 1A-2D, 3
Oct. 27	Logical Design and the Relational Model II	(same as previous session)
Oct. 29	Second Hands-on Workshop	none / none
Oct. 31	Friday - Conference with professor on project status	Bring EDM, Function vs. Entity, hard copy of user views, ERDs etc.
Nov. 1	Saturday - Conference with professor on project status	
Nov. 3	Second Workshop Solutions, Discussion, Analysis	none / none

Course Schedule

Date	Topic	Readings/Questions
------	-------	--------------------

Module 5 - Physical Database Design

Nov. 5	Physical Design I	Chap. 5, pp., Appendix C / 1-3, 5, 7, 8, 16
Nov. 10	Physical Design II	(same as previous session)
Nov. 12	Physical Database Optimization I	(same as previous session)
Nov. 17	Physical Database Optimization II	(same as previous session)
Nov. 19	Physical Design & PL/SQL	PL/SQL handout
Nov. 24	<i>Third Hands-on Workshop</i>	none / none
Nov. 26	<i>Thanksgiving Break (no class)</i>	

Module 6 - The "Wrap"

Dec. 1	Third Workshop Solutions, Discussion, Analysis	none / none
Dec. 3	Project "Spot" Presentations	none / none
Dec. 12	<i>Semester Projects Due @ JHH 102 from 2:00-5:00</i>	none / none

Components of the Database Design Project Report

A. Introduction

1. Cover page
2. Table of Contents
3. Executive Summary

B. Planning for the Database – Please see handout “*Practical Database Development*” for detailed instructions on Step B here (the textbook is not thorough on this step).

4. Business Function-to-Data Entity Matrix (page 24)
5. Preliminary Enterprise Data Model (EDM) (page 51, but we also will show minimum cardinality)

C. Conceptual Design

6. Entity Relationship Diagrams (ERDs) for each user view. You also must attach the original "hard copy" of the user view for comparison - please make it clear which ERD is for which user view. Some annotation for each user view is expected.
7. Create a conceptual data model (CDM) by integrating all the individual user views into one ERD
8. Business Rules/Constraints
 - a. Domain Definitions (create domain classes and refer individual attributes to them - page 162)
 - b. Other business rules, restrictions, constraints and requirements for the database

D. Logical Design with the Relational Model

9. Transform the conceptual data model into a set of Third Normal Form (3NF) relations indicating all primary and foreign keys, making note of any changes from the CDM due to normalization.

E. Physical Design and Implementation with the Relational Model

10. a. Functionally decompose/expand your Business Function-to-Data Entity Matrix into a Process versus Entity Matrix (see lecture handout)
 - b. Two or three Transaction Analysis Forms (TAFs) of key processes (see lecture handout).
11. Composite Usage Map - combine TAFs with remaining usage and volume estimates (lecture & p.210)
12. List of all DDL statements necessary to create the physical model, making note of any changes from the logical model due to physical design/performance considerations.
13. List contents of all populated tables/relations to show sample data (if output wraps on the page, Please format it so it is readable).
14. List six sample queries/views and the results of their executions (show the query statements/code and the output they generated). This is a good opportunity to showcase your SQL skills.
15. Summary discussion of your physical design (e.g., index creations, clustering, responses to table volumes and access patterns, denormalizations, data distribution, table partitioning, triggers, stored procedures, security issues, physical implementation of the constraints, etc.). This portion of the report "implements" your project, please be rigorous/thorough here.

F. Summary and Conclusion

Discuss difficulties encountered, how you solved problems, what you like most/least about your database and what you learned (a thorough description will be appreciated here).

Please note: You must submit files (i.e., a disk or USB drive) containing all SQL code and output for your database tables, views, queries, etc. The report must be page numbered; the entries in the table of contents should appear as headings in the appropriate report sections. It should be well written in a "normal form" for business reports. As always, the report must have sufficient annotation (text) to explain why each diagram or page is in the report and what it is telling the reader. If you have questions regarding report format, please see me as I have samples from previous semesters.

Please start this project early as possible; it will take longer to complete than you think.