

Module Description, available in: EN***Advanced Data Management – Foundations of data-intensive systems*****General Information****Number of ECTS Credits**

3

Module code

TSM_AdvDataMgmt

Valid for academic year

2024-25

Last modification

2023-10-24

Coordinator of the module

Roberto Mastropietro (SUPSI, roberto.mastropietro@supsi.ch)

Explanations regarding the language definitions for each location:

- Instruction is given in the language defined below for each location/each time the module is held.
- Documentation is available in the languages defined below. Where documents are in several languages, the percentage distribution is shown (100% = all the documentation).
- The examination is available 100% in the languages shown for each location/each time it is held.

	Lausanne			Lugano	Zurich		
Instruction				X E 100%			
Documentation				X E 100%			
Examination				X E 100%			

Module Category

TSM Technical scientific module

Lessons

2 lecture periods and 1 tutorial period per week

Entry level competences**Prerequisites, previous knowledge**

- Database design
- Relational Model
- Normalization
- SQL
- Object relational database systems
- Object-oriented programming languages

Brief course description of module objectives and content

The course covers the following core topics:

- Distributed and parallel database systems architectures, internals and services such as transaction processing, concurrency control and query processing
- No-SQL Systems
- Data processing architectures

Aims, content, methods

Learning objectives and acquired competencies

Students understand how to use database technologies to process and manage large data collections.

- They know databases alternatives beyond Relational and Object Relational systems and are able to decide which database system is appropriate depending on the context, and depending on the kind of data available
- They can design and implement Systems based on different architectures
- They understand the functioning of internal components of a database system
- They can reuse the material acquired in this course in their own working environment and apply them to solve their specific problems
- They know the current research directions of these domains.

Contents of module with emphasis on teaching content

The Internet, new types of data and applications and new business requirements have driven the development of data management systems having data models and architectures beyond the traditional relational and object-relational systems and centralised architectures.

The module is organised around the following core subject areas:

- Parallel query processing (20%)
- Replication and partitioning (40%)
- Consistency (25%)
- DLT and Blockchain (15%)

Contents:

- Introduction, data intensive architectures, relational vs document-based systems
- Data processing architectures, parallel queries on a single node
- Serializability
- Storage and retrieval, row-oriented vs column-oriented storage
- Partitioning
- Consistent hashing
- Replication
- Distributed consensus (2PC, 3PC, Paxos, Raft)
- Zookeeper (CAP, ZAB, Leader election)
- Batch processing: Distributed storage: block vs object storage
- Stream processing
- Blockchain

Teaching and learning methods

- Lectures with integrated exercises
- Self study of literature
- case studies

Literature

Lecture slides, references to internet resources and books

Assessment

Certification requirements

Module uses certification requirements

Certification requirements for final examinations (conditions for attestation)

The successful delivery of solved exercises is condition for entering the examination, but will not contribute to final mark.

Basic principle for exams

As a rule, all the standard final exams for modules and also all resit exams are to be in written form

Standard final exam for a module and written resit exam

Kind of exam

written

Duration of exam

120 minutes

Permissible aids

No aids permitted

Special case: Resit exam as oral exam

Kind of exam

oral

Duration of exam

30 minutes

Permissible aids

No aids permitted