

Course on “Cloud Computing & Big Data”

Prof. [Tommaso Cucinotta](#)

Goal

This course provides an overview of the challenges to face, and the technical solutions to embrace, when building large-scale, fault-tolerant, distributed and replicated real-time cloud services. These systems need to be capable of serving millions/billions of requests per second with industrial-grade reliability, availability and performance, and are composed of thousands of components spanning across millions of machines, worldwide. The course focuses on design, development, and operations of scalable software systems, including big-data processing and analytics, as used increasingly often for nowadays intensive computations needed to train large machine-learning and artificial intelligence models, where the huge volumes of data to handle mandates the use of heavily distributed algorithms. The course covers also basic concepts on networking architectures for data-centre and cloud computing infrastructures.

Program at a glance

- Cloud Computing
 - Basic concepts
 - Scalability and elasticity in cloud systems
 - Data center networks for Cloud & Big-Data
 - Fault-tolerance and replication
 - Real-time cloud services
 - Operations and devops engineering
 - Overview of public cloud services (EC2, GCP...)
 - Open Cloud orchestrators (OpenStack, Kubernetes)
- Big Data and Analytics
 - Basic concepts
 - Real-time data streaming and analytics
 - Distributed file-systems
 - SQL vs NoSQL data stores
 - Big-Data and the Internet of Things
 - Big-Data Platforms: Map Reduce, Spark

Requirements

Students need a basic understanding of software, computer architectures, distributed systems and communication protocols.

Why to attend

Students will acquire a unique insight into the world of cloud computing and big-data related technologies, and will be able to master key concepts behind them. This is a

fundamental brick in the background of a software engineer / computer scientist who will deal with modern distributed software systems in industry or academia, spanning across high-performance, cloud and even (increasingly connected) embedded systems.

About the course

Duration and format: front lessons, 30 hours (3 CFU).

Target programs: [PhD in AI](#), [PhD in Data Science](#), [PhD in Emerging Digital Technologies](#).

Schedule: flexible, to be agreed with students, with planned beginning on mid November.

Exam: oral.

More details: available on the [course web-page](#).

About the lecturer

[Prof. Tommaso Cucinotta](#) has a MSc in Computer Engineering from University of Pisa and a PhD from Scuola Superiore Sant'Anna. He spent more than 10 years at the [Real-Time Systems Laboratory \(ReTiS\)](#) of Scuola Superiore Sant'Anna carrying out research in security and smart-card based authentication, adaptive deadline-based scheduling in the Linux kernel for embedded, soft real-time and multimedia applications, temporal isolation in virtualized cloud services and novel OS designs for massively parallel and distributed systems. He has been a Member of Technical Staff (MTS) at [Bell Labs](#) in Dublin, carrying out industrial research on security and confidentiality, and real-time performance of cloud systems, with a focus on Telco applications. He has also been a Software Development Engineer in AWS DataBase Services in Dublin, Ireland, working on scalability and performance enhancements to the [AWS DynamoDB](#) NoSQL real-time data-base.

Since 2016, he is back at the ReTiS of Scuola Superiore Sant'Anna as associate professor. He is a member of the PhD board of the [Data Science PhD](#) program jointly offered by Scuola Sant'Anna, University of Pisa, Scuola Normale Superiore, IMT Lucca and CNR since its first edition in the a.y. 2017/2018. He is also a member of the PhD board of the [National PhD in AI](#) – AI for Society – program jointly offered by Scuola Sant'Anna, University of Pisa and other academic institutions. Since year 2019, he is coordinator of the real-time and embedded systems research area at the RETIS.