

## Course on “Cloud Computing & Big Data Lab”

Prof. Tommaso Cucinotta

### Goal

This is a hands-on and applied course following up to the Cloud Computing & Big-Data course. Here, students will put in practice the theoretical/abstract concepts acquired in the general course on Cloud Computing & Big-Data. During the practical sessions, we'll have a deep dive on such concepts as: machine virtualization and OS-level virtualization on Linux; virtual networking on Linux; programming abstractions for cloud and distributed computing; elasticity in practice; big-data programming frameworks; command-line interface for major public cloud services; popular open-source cloud platforms.

### Program at a glance

- Virtualization Fundamentals
  - KVM Command-Line Interface
  - libvirt and virtual-manager
- Virtual Switching on Linux
  - brctl and OpenVSwitch
- Containers
  - LXC and netns
- Public Cloud Services
  - AWS EC2, S3, DynamoDB, CloudWatch
- Open-source cloud and big-data platforms
  - OpenStack Nova, Glance, Neutron
  - OpenStack Heat/Senlin, Ceilometer/Monasca
  - Apache Spark

### Requirements

Students need to have attended the Cloud Computing & Big-Data general course, and they must be fluent in programming and shell scripting on Linux.

### 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 tools behind their use. 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:* lab 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'April.

*Exam:* coding project and oral examination.

*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 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.