

Warehouse-Scale Computing

Docente

Ozalp Babaoglu

Durata

12 ore

Periodo

Marzo 2023

Contenuto

As computations continue to move into the cloud, the platforms that they run on no longer resemble pizza boxes or refrigerators, but warehouses full of computers. These new large datacenters are quite different from traditional hosting facilities of earlier times and cannot be viewed simply as a collection of co-located servers. Large portions of the hardware and software resources in these facilities must continue to work correctly and in concert to efficiently deliver services, even in the presence of adverse effects such as errors and failures. This goal can only be achieved through a holistic approach to their design and deployment where the datacenter itself is treated as one massive warehouse-scale computer (WSC).

In this course we will examine the technical challenges that need to be confronted in the design and operation of WSC and try to give answers to questions such as: How are WSC built and programmed? Why have they become so ubiquitous? Who are the major players that operate them? What services are best provided through WSC? What makes them cost-effective and attractive for businesses? How do they store and process huge collections of information? How can they be made to be secure? How can they be made to be fault tolerant? How can their energy consumption be contained to practical levels? Which software technologies are most appropriate for their design and efficient operation?

Bibliografia

Barroso, Luiz André, Urs Hölzle, and Parthasarathy Ranganathan. "The Datacenter as a Computer: Designing Warehouse-Scale Machines", Third Edition, Morgan & Claypool Publishers series *Synthesis Lectures on Computer Architecture* Lecture #24 (2018): i-189.