Building a sustainable, intelligent and power-efficient cloud

September 29 2016

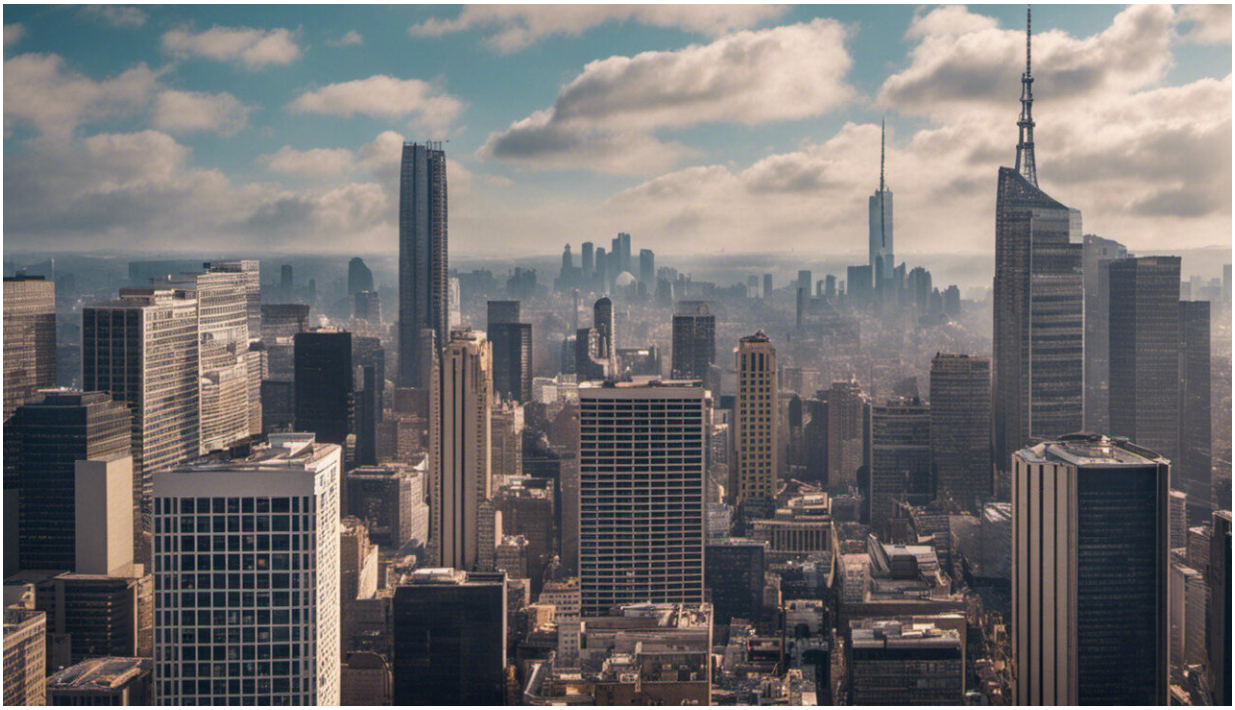

Credit: AI-generated image (disclaimer)

EU researchers are combining high-performance computational power with hardware and servers in order to build a more efficient, heterogeneous cloud.

Whether it's sharing files via Dropbox, saving images on Google Drive

or placing an order with Amazon, we all use what is commonly referred to as 'the cloud' for our basic computing needs. But did you know the typical cloud server only operates at approximately 20 % of its computing capacity? In fact, the cloud is now responsible for nearly 10 % of the world's total electricity consumption – meaning the current model of cloud computing is simply not sustainable.

The EU CLOUDLIGHTNING project intends to change this by creating a more efficient, sustainable and user-centric cloud. The project is working to develop an intelligent, power-efficient cloud computing infrastructure that will provide energy savings to cloud service providers and simplify access for cloud users. By using heterogeneous processing resources, the project believes it can increase the cloud's computing capacity from the dismal 20 % of today to a more sustainable 80 %.

Behind the cloud

It might be hard to understand how cloud computing can be so inefficient. After all, the whole idea of the cloud is that we no longer need to have physical files, data or hardware. However, behind the cloud sit huge, homogeneous data centres comprised of a large number of machines, components and hardware – a model that CLOUDLIGHTNING researchers believe restricts computational processing power and limits what certain cloud computing users are able to do.

As a case in point, take the high performance computing (HPC) used by the technical and scientific community – a community whose needs are not met by the commoditised nature of today's homogeneous cloud. For example, their use of cloud resources is not always predictable and tends to scale up and down. As today's cloud providers do not offer the ability to configure available resources in a way that can meet these unique needs, cloud computing is simply not being used for HPC.

According to a recently published report by the CLOUDLIGHTNING project, the project team has found that such inflexibility has led to HPC users to develop a general sense of distrust in the cloud. In particular, they have concerns about cloud computing's data management capabilities, including the lack of cloud infrastructure capable of meeting the communications and I/O requirements of high-end technical computing.

Another problem with the current system is that it encourages cloud providers to over-provision their computing resources as a way of accommodating users' unpredicted peaks in demand. The downside of this practice, however, is that it results in tens of thousands of cloud servers located in data centres around the world running – and consuming energy – even though they're not actually being used.

A heterogeneous idea

CLOUDLIGHTING aims to solve these barriers to cloud computing by developing a power-efficient cloud infrastructure that will simplify access to cloud recourses. Its proposed solution involves the creation of a heterogeneous cloud system that combines high performance computational power with the power-efficient use of different types of hardware and servers capable of working together. More technically, the project proposes a novel cloud management and delivery architecture based on the principles of self-organisation and self-management, which shifts the deployment and optimisation effort from the consumer to the software stack running on the cloud infrastructure.

The project's ultimate goal is to address the inefficient use of resources and, consequently, to deliver savings to the cloud provider and the cloud consumer in terms of reduced power consumption and improved service delivery. Although the project is still a work in progress, it has already published use cases for three application domains – genomics, oil and

gas exploration and ray tracing – which will be used to validate the CLOUDLIGHTNING management and delivery models.

The project's next step after its first promising results is to launch a testbed that runs the CLOUDLIGHTNING software stack. Researchers note that this testbed will gather the instrumentation data that will subsequently serve as the basis for large-scale simulations of self-organised and self-managed hyperscale heterogeneous clouds.

More information: Project website: cloudlightning.eu/

Provided by CORDIS