

The XenoServer Platform for Global-Scale Service Deployment

Evangelos Kotsovinos, Jon Crowcroft
Computer Laboratory, University of Cambridge

This talk proposes a new distributed computing paradigm, termed global public computing, which allows any user to run any code anywhere. Public computing platforms price computing resources, and ultimately charge users for resources consumed.

To facilitate global public computing, we have addressed several research challenges. We have devised mechanisms (reusable in other distributed computing settings, such as grids) for representing, advertising, and supporting the discovery of resources. To allow flexible and federated control of resource allocation by all stakeholders involved, we proposed a novel role-based resource management framework for expressing and combining distributed management policies. Furthermore, we implemented effective service deployment models for launching distributed services on large numbers of machines around the world easily, quickly, and efficiently. To keep track of resource consumption and pass charges on to consumers, we developed an accounting and charging infrastructure.

This talk will also outline the design and implementation of the XenoServer Open Platform. The efficiency and scalability of the developed mechanisms are demonstrated by experimental evaluation; the prototype platform allows the global-scale deployment of complex services in less than 45 seconds, and could scale to millions of concurrent sessions without presenting performance bottlenecks.