

MISQ Archivist

Shared or Dedicated Infrastructures? On the Impact of Re provisioning Ability

Roch Guerin, Kartik Hosanagar, Xinxin Li, and Soumya Sen

Abstract

New technologies such as virtualization, are transforming the way in which software and services are deployed and delivered to their users. They are behind the emergence of IT offerings such as cloud computing and converged networks, and manifest themselves through two important trends: (1) lower cost of sharing a common infrastructure across multiple services with disparate resource requirements, and (2) dynamic provisioning of capacity in response to demand. Conventional wisdom is that both of these capabilities are synergistic, with greater provisioning flexibility improving the benefits derived from sharing computing or network resources. Consequently, a service operator should now always favor the use of a shared infrastructure over dedicated solutions when hosting multiple services. In this paper, we ask whether this is indeed the case, and investigate the dual impact of lower costs of sharing and provisioning flexibility on shared and dedicated infrastructures. The investigation reveals that while lower costs are always expected to favor infrastructure sharing, dynamic provisioning plays an ambiguous role. Re provisioning improves both shared and dedicated solutions, but can do so differently and can sometimes favor a dedicated infrastructure. Our findings help illustrate that the technology trends, such as virtualization, behind cloud computing need not always favor the deployment of services on a shared infrastructure.

Keywords: Network infrastructure, infrastructure design, service deployment, virtualization, resource provisioning