Cost Effective, Reliable and Secure Workflow Deployment over Federated Clouds

The significant growth in cloud computing has led to increasing number of cloud providers, each offering their service under different conditions – one might be more secure whilst another might be less expensive or more reliable. At the same time user applications have become more and more complex. Often, they consist of a diverse collection of software components, and need to handle variable workloads, which poses different requirements on the infrastructure. Therefore, many organisations are considering using a combination of different clouds to satisfy these needs. It raises, however, a non-trivial issue of how to select the best combination of clouds to meet the application requirements. This paper presents a novel algorithm to deploy workflow applications on federatedclouds. Firstly, we introduce an entropy-based method to quantify the most reliableworkflowdeployments. Secondly, we apply an extension of the Bell-LaPadula Multi-Level security model to address application security requirements. Finally, we optimise deployment in terms of its entropy and also its monetary cost, taking into account the cost of computing power, data storage and inter-cloud communication. We implemented our new approach and compared it against two existing scheduling algorithms: Extended Dynamic Constraint Algorithm (EDCA) and Extended Biobjective dynamic level scheduling (EBDLS). We show that our algorithm can find deployments that are of equivalent reliability but are less expensive and meet security requirements. We have validated our solution through a set of realistic scientific workflows, using well-known cloud simulation tools (WorkflowSim and DynamicCloudSim) and a realistic cloud based data analysis system (e-Science Central).