

Towards CloudComputing@home: Ensuring Collective Availability in Volatile Resource Pools via Forecasting

Derrick Kondo

INRIA, France

Increasingly services are being deployed over large-scale computational and storage infrastructures. To meet ever-increasing computational demands and to reduce both hardware and system administration costs, these infrastructures have begun to include Internet resources distributed over enterprise and residential broadband networks. As these infrastructures increase in scale to hundreds of thousands to millions of resources, issues of resource availability and service reliability inevitably emerge. Our goal in this study is to determine and evaluate predictive methods that ensure the availability of a collection of resources. We gather real-world availability data from over 48,000 Internet hosts participating in the SETI@home project. With this trace data, we show how to reliably and efficiently predict that a collection of N hosts will be available continuously for T time. The results indicate that by using replication it is feasible to deploy enterprise services or applications even on such volatile resource pools.

