

Experimental Evaluation of N-Tier Systems: Finding Bottlenecks beneath Low Average Resource Utilization

Simon Malkowski and Calton Pu
Center of Experimental Research in Computer Systems
Georgia Institute of Technology
{zmon, calton}@cc.gatech.edu

Abstract

It is important to undertake measurement-based performance characterization and analysis of large distributed systems to automate their design and operations management. Concretely, systems can be better prepared to meet expected deployment requirements through the analysis of resource accounting logs and an explicit bottleneck detection process. In this work we investigate bottleneck phenomena experimentally and find part-time as well as multiple-simultaneous bottlenecks to be commonly observable. We show that classical analysis through evaluation of average resource utilization values masks actual bottleneck behavior and that the observed phenomena necessitate the relaxation of classical assumptions such as dominant single system bottlenecks. Methodologically, we present an in-depth bottleneck analysis for a cluster running RUBiS and RUBBoS N-Tier benchmarks in various scenarios. The data include workloads between 300 and 13,000 concurrent users, up to 8 application and 9 database servers, low and high-end nodes, artificially throttled network bandwidth, and transaction-mixes between 0 and 70 percent read/write interactions.

Availability

Interested readers should refer to the authors for more information and a complete version of this technical report.