

Chapter 5

Quality of Service Monitoring Strategies in Service Oriented Architecture Environments using Processor Hardware Performance Metrics

Ernest Sithole

University of Ulster at Coleraine, UK

Sally McClean

University of Ulster at Coleraine, UK

Bryan Scotney

University of Ulster at Coleraine, UK

Gerard Parr

University of Ulster at Coleraine, UK

Adrian Moore

University of Ulster at Coleraine, UK

Dave Bustard

University of Ulster at Coleraine, UK

Stephen Dawson

SAP Research Belfast, UK

ABSTRACT

The sharp growth in data-intensive applications such as social, professional networking and online commerce services, multimedia applications, as well as the convergence of mobile, wireless, and internet technologies, is greatly influencing the shape and makeup of on-demand enterprise computing environments.

In response to the global needs for on-demand computing services, a number of trends have emerged, one of which is the growth of computing infrastructures in terms of the number of computing node entities and the widening in geophysical distributions of deployed node elements. Another development has been the increased complexity in the technical composition of the business computing space due to the diversity of technologies that are employed in IT implementations. Given the huge scales in infrastructure sizes

DOI: 10.4018/978-1-61350-432-1.ch005

and data handling requirements, as well as the dispersion of compute nodes and technology disparities that are associated with emerging computing infrastructures, the task of quantifying performance for capacity planning, Service Level Agreement (SLA) enforcements, and Quality of Service (QoS) guarantees becomes very challenging to fulfil. In order to come up with a viable strategy for evaluating operational performance on computing nodes, we propose the use of on-chip registers called Performance Monitoring Counters (PMCs), which form part of the processor hardware. The use of PMC measurements is largely non-intrusive and highlights performance issues associated with runtime execution on the CPU hardware architecture. Our proposed strategy of employing PMC data thus overcomes major shortcomings of existing benchmarking approaches such as overheads in the software functionality and the inability to offer detailed insight into the various stages of CPU and memory hardware operation.

MAJOR DEVELOPMENTS AND CHALLENGES IN ON-DEMAND COMPUTING

Given the current developments in eEnterprise implementations, the infrastructure planning tasks of accurately determining performance that can be delivered by on-demand computing resources and in turn, obtaining accurate estimates of the appropriate infrastructure hardware performance capabilities required for business computing solutions are becoming an increasingly challenging exercise to undertake. One major cause that has led to the difficulty in quantifying performance in business computing systems has been the huge amounts of user-generated data as a result of the exponential adoption of on-demand hosted computing services and applications such as social networking, e-commerce, and multi-media content sharing services. Yet another development that has caused a huge increase in consumer-generated data is the convergence of mobile, wireless and web technologies into a ubiquitous computing platform. In response to the challenges arising from these major trends, there has been a phenomenal growth in the size of deployed computing infrastructures, with the magnitude of the infrastructure expansion being characterised by three main dimensions of growth: (a) the increases in the number of computing machines brought together to form server domains, (b) the

wide geographic locations, over which participant server nodes are physically deployed and (c) the different types of technologies that are used to produce computing solutions.

Challenges for Performance Evaluation

Arising from the physical distribution of compute nodes due to the dispersion of resources in the infrastructures, are the performance-related challenges pertaining to the need to quantify network delays. The delays emanate from the communications of status and coordination messages as well as the actual data transfers between host machines. The calculation of overall performance metrics in distributed systems, which are dependent on network delays, is not straightforward to perform given that application routines running inside server nodes generate data in quantities that can vary dynamically. As a result of the changing loads, traffic levels introduced on network links usually follow irregular patterns leading to congestion and bandwidth delays that cannot be easily established. It is important to emphasize that while network delays do not feature in the proposed performance evaluation strategy considered in the latter sections of this chapter, they nevertheless make a key aspect which is captured in the Service Taxonomy presented in Figure 4.

27 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/quality-service-monitoring-strategies-service/60883

Related Content

Service-Oriented Enterprise Engineering: A Modeling Discipline Based on the Viable Systems Approach (vSa) for Strategic Sourcing Decision-Making

Laleh Rafatiand Geert Poels (2018). *International Journal of Information Systems in the Service Sector* (pp. 20-40).

www.irma-international.org/article/service-oriented-enterprise-engineering/206895

An Analysis of the Relationship between Quality Factors and Customer Loyalty for Success of Firms in the Service Industry: A Case Study

Cecilia Silvestri (2017). *Strategic Marketing Management and Tactics in the Service Industry* (pp. 74-106).

www.irma-international.org/chapter/an-analysis-of-the-relationship-between-quality-factors-and-customer-loyalty-for-success-of-firms-in-the-service-industry/179987

Towards High Maturity in SaaS Applications Based on Virtualization: Methods and Case Study

Shijun Liu, Yong Zhangand Xiangxu Meng (2013). *Implementation and Integration of Information Systems in the Service Sector* (pp. 223-236).

www.irma-international.org/chapter/towards-high-maturity-saas-applications/72552

Exploring the Factors Behind the Resistance to Mobile Banking in Portugal

Pedro Cruz, Tommi Laukkanenand Pablo Muñoz (2011). *Developing Technologies in E-Services, Self-Services, and Mobile Communication: New Concepts* (pp. 200-218).

www.irma-international.org/chapter/exploring-factors-behind-resistance-mobile/54964

Empirical Performance Assessment of Public Clouds Using System Level Benchmarks

Sanjay P. Ahuja, Thomas F. Furman, Kerwin E. Roslieand Jared T. Wheeler (2013). *International Journal of Cloud Applications and Computing* (pp. 81-91).

www.irma-international.org/article/empirical-performance-assessment-of-public-clouds-using-system-level-benchmarks/105511