Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 37 (2014) 244 – 251

Procedia Computer Science

The 5th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN-2014)

Classification of Post-Deployment Performance Diagnostic Techniques for Large-Scale Software Systems

Haroon Malik[a]*, Elhadi M. Shakshuki[b]

[a]*David R. Cheriton School of Computing, University of Waterloo, Waterloo , ON, N2L3G1, Canada*
[b]*Jodery School of Computing, Acadia University, Wolfvilee, NS, B4P2R6, Canada*

Abstract

Today's large-scale software systems (LSSs) such as Facebook, Google, Amazon and many other contemporary datacenters comprise hundreds or thousands of machines running complex applications that require high availability and responsiveness. These LSSs must be carefully monitored for performance bottlenecks before a serious harm is done. Performance analysts have to deal with the tedious task of monitoring the performance of these LSSs to avoid any service level agreements (SLA) violations and to ensure their failure free operations. There do exist several post-deployment performance diagnostic (PPD) techniques for to help analysts diagnose performance problems in the field, i.e., after the software is deployed. However, there is no classification of the proposed PPD techniques to understand their objectives and characteristics. In this paper, we classify the existing PPD techniques along multiple categories. The classification of PPD techniques will provide a guideline for performance analysts and practitioners of LSS to choose techniques suitable for their need. Moreover, the classification will also help researcher understand and fill gaps, i.e., dedicate their research efforts to categories that have received little attention in the past.

Peer-review under responsibility of the Program Chairs of EUSPN-2014 and ICTH 2014.

Keywords: Large-Scale Software System; Performance Monitoring; Performance Diagnostics; Anomaly Detection;

1. Introduction

Today's large-scale systems (LSS) consisting of datacenters and server farms have experienced an extraordinary explosion in size and complexity. For example, Google the fifth largest data center alone maintains a pool of more than one million servers [1]. Facebook has doubled the size of its data center within a year from 60,000 to 120,000

* Corresponding author. Tel.: +1-519-957-2262;
E-mail address: malikh@uwaterloo.ca

Peer-review under responsibility of the Program Chairs of EUSPN-2014 and ICTH 2014.
doi:10.1016/j.procs.2014.08.036

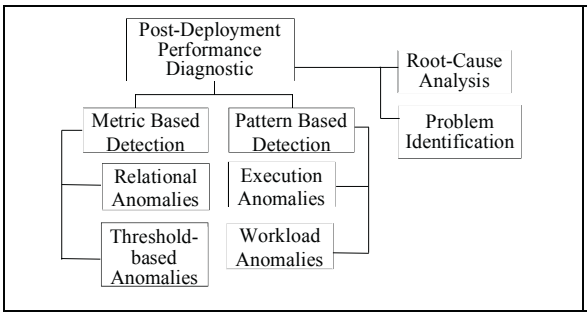

Fig.1. Taxonomy for PPD techniques for large-scale software systems

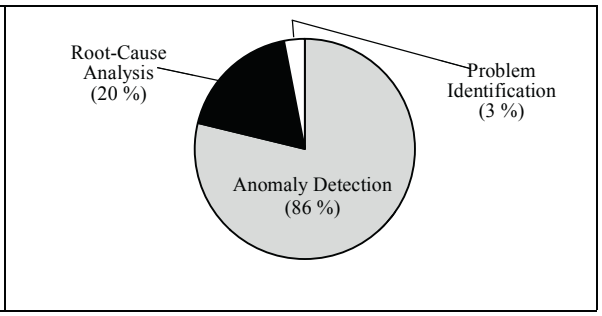

Fig. 2. Composition of PPD work

servers and continues to grow [2, 3]. Similarly, the collective server count of Microsoft, eBay, Yahoo and Amazon is over 550,000 servers [4-6].

LSSs are large capital investments for the service providers. Any discrepancy in their performance can result in large monetary losses. For example, an hour-long PayPal outage may have prevented up to $7.2 million in customer transactions [7]. Therefore, operators of LSS closely monitor their LSS to anticipate and identify performance bugs before a violation of service level agreements (SLAs) and unplanned system downtime that can cost as much as $550,000 per hour in lost revenues of LSS [8].

A performance problem includes an application or a system not responding fast enough, taking too much of an important resource, crashing or hanging under heavy load, or not meeting the desired service-level agreements (SLAs) [9]. Performance problems with LSS typically exhibit symptoms such as high response time, increased latency and low throughput under load.

The occurrence of performance problems in ever-growing LSSs has become a norm, rather than an exception [10]. Performance problems after deployment/expansion (i.e., post-deployment) of LSS are seldom due to feature errors, rather due to systems not scaling to field workloads. LSSs like AT&T and BlackBerry also report their concerns about performance degradation and resource saturation as the fundamental post-deployment problems [11]. For example, a robust performance monitoring and diagnostics could have alerted the operators of Skype for the system overload that resulted in system availability disruption for 48 hours, leaving millions of users without service [12].

Although many techniques have been introduced in the literature to diagnose post-deployment performance problems (e.g., [3,4,5,6,8,9,10,11,12,13,14,15]), there is no classification to understand the objectives and characteristics of these techniques. Moreover, non-existence of a classification provides a) little or no direction to practitioners in choosing the appropriate PPD techniques to satisfy their respective performance diagnostic objective and b) for researcher, identify gaps, i.e., dedicate their efforts to the classification category, where little work is done.

In this paper, we survey the state-of-the-art PPD techniques aimed to diagnose performance problems of LSS in the field, i.e., post-deployment performance problem. We first classify the PPD techniques into three major categories, i.e., anomaly detection, root/cause analysis and problem identification. We further categories anomaly detection category into two major sub-categories based on detection techniques, i.e., a) metric based detection techniques that capture both relational and threshold-based anomalies and b) pattern detection techniques that are geared to diagnose both execution and workload anomalies as shown in Fig.1. The survey will help researchers to differentiate existing PPD techniques. Moreover, it will provide a guideline to practitioners to choose an appropriate PPD technique suitable for their needs.

2. Performance Diagnostic

LSS contains multiple subsystems that interact across multiple nodes in sometimes unforeseen and complicated ways. As a result, detecting, isolating and identifying the root-cause of a performance problem is frustrating and can