Accepted Manuscript

Self-adaptive cloud monitoring with online anomaly detection

Tao Wang, Jiwei Xu, Wenbo Zhang, Zeyu Gu, Hua Zhong

PII: S0167-739X(17)30376-X
DOI: https://doi.org/10.1016/j.future.2017.09.067
Reference: FUTURE 3723

To appear in: *Future Generation Computer Systems*

Received date : 9 March 2017
Revised date : 29 August 2017
Accepted date : 25 September 2017

Please cite this article as: T. Wang, J. Xu, W. Zhang, Z. Gu, H. Zhong, Self-adaptive cloud monitoring with online anomaly detection, *Future Generation Computer Systems* (2017), https://doi.org/10.1016/j.future.2017.09.067

Self-Adaptive Cloud Monitoring with Online Anomaly Detection

Tao Wang[1], Jiwei Xu[1], Wenbo Zhang[1*], Zeyu Gu[2], and Hua Zhong[1]

1. Institute of Software, Chinese Academy of Sciences, Beijing 100190, China
2. SRE(Site Reliability Engineering) Group, Baidu Inc., Beijing 100085, China

*Corresponding author: zhangwenbo@otcaix.iscas.ac.cn

Abstract—Monitoring is the key to guarantee the reliability of cloud computing systems. By analyzing monitoring data, administrators can understand systems' statuses to detect, diagnose and solve problems. However, due to the enormous scale and complex structure of cloud computing, a monitoring system should collect, transfer, store and process a large amount of monitoring data, which brings a significant performance overhead and increases the difficulty of analyzing useful information. To address the above issue, this paper proposes a self-adaptive monitoring approach for cloud computing systems. First, we conduct correlation analysis between different metrics, and monitor selected important ones which represent the others and reflect the running status of a system. Second, we characterize the running status with Principal Component Analysis (PCA), estimate the anomaly degree, and predict the possibility of faults. Finally, we dynamically adjust the monitoring period based on the estimated anomaly degree and a reliability model. To evaluate our proposal, we have applied the approach in our open-source TPC-W benchmark Bench4Q deployed in our real cloud computing platform OnceCloud. The experimental results demonstrate that our approach can adapt to dynamic workloads, accurately estimate the anomaly degree, and automatically adjust monitoring periods. Thus, the approach can effectively improve the accuracy and timeliness of anomaly detection in an abnormal status, and efficiently lower the monitoring overhead in a normal status.

Key words—Cloud computing; Anomaly detection; Adaptive monitoring; Correlation analysis

1 INTRODUCTION

In recent years, we are witnessing a rapid development of cloud computing. Being applied in various fields, cloud computing has become a hot spot in IT industry. Many large IT companies all over the world have released their cloud platforms as their core businesses (e.g., Amazon EC2 and Microsoft Azure), while other open-source cloud platforms (e.g., Eucalyptus and OpenStack) also contribute to cloud computing. Currently, online services in Internet (e.g., e-commerce, online finance, social network) have become important parts of our lives, and more and more services have been deployed on cloud platforms (e.g., Saleforce CRM and Netflix). However, due to the diversity of deployed services and dynamism of deployment environment, cloud computing systems often come with unexpected faults which result in significant disruption in user experiences and enormous economic losses for companies [1]. For example, the Amazon website was down for about 45 minutes in August 2013, which brought about 5.3 million losses as customers could not make purchases during that period. Furthermore, the delivery models (e.g., IaaS, PaaS, SaaS) are vulnerable to security attacks [2]. Efficient monitoring is the prerequisite to timely detect anomalies and accurately locate root causes [3].

Cloud computing systems are always in large scales and have complex architectures. To track the running status of these systems, a monitoring system always acquires kinds of monitoring data from different layers (e.g., network, hardware, virtual machine, operating system, middleware, application) in lots of distributed nodes. However, collecting, storing and processing a large amount of monitoring data should cause a huge resource overhead, which affects the timeliness of anomaly detection, the accuracy of fault location, and even overall performance. Furthermore, it is difficult to mining valuable information from large-scale data. Thus, many commercial monitoring tools (e.g., Amazon CloudWatch) only support a fixed and relatively long monitoring period (e.g., collecting a data instance every minute) and manually pre-defined collected metrics. Meanwhile, the providers of cloud monitoring services charge their customers according to monitored metrics and monitoring frequency. Monitoring cost is usually proportional to the number of monitored metrics and monitoring frequency, which on average takes up 18% of total running cost for cloud computing systems [4]. Above all, on one hand, administrators and customers want to decrease the number of monitored metrics and decrease the monitoring frequency to reduce their costs; on the other hand, less monitored metrics and lower monitoring frequency result in less monitoring data, making it hard to locate faults efficiently. For example, the lower the monitoring frequency is, the higher the chance that faults may occur just in between two monitoring data instances. Hence, how to select monitored