

Accepted Manuscript

A self-organized resource provisioning for cloud block storage

Babak Ravandi, Ioannis Papapanagiotou

PII: S0167-739X(17)33012-1
DOI: <https://doi.org/10.1016/j.future.2018.06.045>
Reference: FUTURE 4307

To appear in: *Future Generation Computer Systems*

Received date : 31 December 2017
Revised date : 21 April 2018
Accepted date : 24 June 2018

Please cite this article as: B. Ravandi, I. Papapanagiotou, A self-organized resource provisioning for cloud block storage, *Future Generation Computer Systems* (2018), <https://doi.org/10.1016/j.future.2018.06.045>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



A Self-Organized Resource Provisioning for Cloud Block Storage

Babak Ravandi

Department of Computer and Information Technology, Purdue University

Ioannis Papapanagiotou

Platform Engineering, Netflix

Abstract

Modern cloud deployments rely on multiple layers of microservices and a layer of one or more persistent data store systems. The latter has always received a lot of attention as it is the slowest part of the request of an online service. Its unpredictable performance due to workload interference, performance degradation, lack of control, and network partitions has made the cloud less attractive for some services, especially for the applications that require guaranteed throughput, latency, elasticity, and high availability. Motivated by those facts, we propose the Block Software Defined Storage (BSDS), a framework to automate the orchestration, deployment, and management of a storage system by separating the data layer from control. BSDS employs an SLA-aware (Service Level Agreement) and self-learning scheduler that treats the backend media as a black box and abstracts the resource allocation from the state of the storage cluster. Our proposed framework is a step toward the emergence of an intelligent cloud infrastructure that can adapt to unexpected failures, network traffic fluctuations, workload interference perturbations or even changes in the underlying infrastructure. We perform extensive simulations with workloads deduced from an enterprise block storage system and showcase that the proposed scheduler can provide guaranteed Quality of Service (QoS). We finally demonstrate

Email addresses: bravandi@purdue.edu (Babak Ravandi), ipapapa@ncsu.edu (Ioannis Papapanagiotou)

Download English Version:

<https://daneshyari.com/en/article/6872874>

Download Persian Version:

<https://daneshyari.com/article/6872874>

[Daneshyari.com](https://daneshyari.com)