

A Thesis for the Degree of Ph.D. in  
Engineering

Study of Database Management  
System Performance and Isolation in  
Virtualization Environments

July 2021

Graduate School of Science and Technology  
Keio University

Asraa Abdulrazak Ali Mardan

## Thesis Abstract

No. \_\_\_\_\_

Registration Number	<input checked="" type="checkbox"/> "KOU" <input type="checkbox"/> "OTSU" No. _____ *Office use only	Name	ALI MARDAN, Asraa Abdulrazak
Thesis Title: Study of Database Management System Performance and Isolation in Virtualization Environments			
<b>Thesis Summary</b> <p>Database management system (DBMS) is one of the foundational and largest applications in the cloud. Major cloud service providers like Amazon web services, Microsoft Azure, Google clouds offer DBMS as a service. The cloud employs virtualization to consolidate DBMSes for efficient resource utilization and to isolate collocated workloads. There are two major virtualization technologies: hypervisor-based (virtual machines) and operating-system-level virtualization (containers). The underlying virtualization technologies in the cloud have a critical impact on performance and isolation, especially in disk I/O. To guarantee the service-level agreement (SLA), the disk I/O performance and its isolation are important in DBMSs because they are inherently disk I/O intensive.</p> <p>This dissertation investigates DBMS performance and isolation in containers and virtual machines. Containers have become widely used in clouds and preferred over virtual machines to consolidate DBMS due to their near-native performance and lightweight deployment. The key finding in this dissertation is virtual machines outperform containers in DBMS performance. Our analysis reveals that file-system journaling has negative effects on DBMS performance and isolation in containers. DBMS is an update-intensive application and causes a lot of file-system journaling. Journaling is very important to keep file-system consistency and for crash recovery. Hence, file-system journaling cannot be disabled especially with DBMS applications.</p> <p>The contribution of this dissertation is twofold. First, identifying the underlying causes behind file-system journaling problems in containers. Since containers share the same file-system, the sharing of journaling modules causes performance dependencies among containers. Also, file-system journaling interferes with disk I/O control of containers and violates isolation between them. Since the journaling module runs outside of controlled containers, their journaling I/Os are overlooked by the disk control mechanism and not accounted for by the container that initiates them. Second, proposing a configuration method to overcome the journaling problems in containers. The method achieves per-container journaling without re-designing the file-systems or modifying the existing kernel. The results show that DBMS performance improves up to 3.4x in containers with the proposed configuration. Eventually, containers get their performance advantage and outperform virtual machines by 1.4x, and show an identical disk I/O isolation.</p>			