



# Analysis of a Network IO Bottleneck in Big Data Environments Based on Docker Containers <sup>☆</sup>



P. China Venkanna Varma <sup>\*</sup>, K. Venkata Kalyan Chakravarthy, V. Valli Kumari, S. Viswanadha Raju

## ARTICLE INFO

### Article history:

Received 1 June 2015

Received in revised form 29 November 2015

Accepted 1 December 2015

Available online 4 January 2016

### Keywords:

Containers  
Context switching  
Docker  
Hadoop  
Map reduce

## ABSTRACT

We live in a world increasingly driven by data with more information about individuals, companies and governments available than ever before. Now, every business is powered by Information Technology and generating Big data. Future Business Intelligence can be extracted from the big data. NoSQL [1] and Map-Reduce [2] technologies find an efficient way to store, organize and process the big data using Virtualization and Linux Container (a.k.a. Container) [3] technologies.

Provisioning containers on top of virtual machines is a better model for high resource utilization. As the more containers share the same CPU, the context switch latency for each container increases significantly. Such increase leads to a negative impact on the network IO throughput and creates a bottleneck in the big data environments.

As part of this paper, we studied container networking and various factors of context switch latency. We evaluate Hadoop benchmarks [5] against the number of containers and virtual machines. We observed a bottleneck where Hadoop [4] cluster throughput is not linear with the number of nodes sharing the same CPU. This bottleneck is due to virtual network layers which adds a significant delay to Round Trip Time (RTT) of data packets. Future work of this paper can be extended to analyze the practical implications of virtual network layers and a solution to improve the performance of big data environments based on containers.

© 2016 Elsevier Inc. All rights reserved.

## 1. Introduction

Data are unorganized facts that need to be processed. When data are processed, organized and structured in a given context so as to make them useful, they are called Information [10]. Now, every business is powered with Information Technology (IT) and generating big data. There is a potential demand to process big data for rapid and accurate business decisions. Data scientists divided big data into five dimensions: Volume, Value, Velocity, Variety and Veracity (simply Big data 5V's). Volume refers to vast amounts of data generated every second. To create Value for the business, volumes of data should be processed to determine the relevance within the data sets. Velocity refers to the data processing at data gathering speed. Variety refers to managing, merging and processing of structured and un-structured data. Veracity refers to the biases, noise and abnormality in data.

Virtualization technology became de facto standard for all public and private cloud requirements. Virtualization has consolidated all the hardware components and created software redundancy layers for elastic work loads. Docker framework introduced manageable Linux containers [3] called Docker Containers. Docker is an open platform for developers and system administrators to build, ship, and run distributed applications. Hadoop is an open-source software framework for storing and processing the big data in a distributed fashion on large clusters.

Nowadays, data centers are adopting Docker to provision Hadoop [4] Clusters (HC) for elastic work loads and multi-tenant service models. Hadoop cluster nodes requires good amount of CPU cycles, Random Access Memory (RAM) and Resource IO to execute Map-Reduce operations. Virtualization solved orchestration of the dynamic resources for data centers, but virtualization overhead is the biggest penalty. "Every single CPU cycle that goes to the Hypervisor is wasted," and "Likewise every byte of RAM". Containers solved orchestration of dynamic resources within an operating system by virtualizing necessary components. If a container runs inside a VM then the network IO path has two extra virtual layers: one layer between Hypervisor and VM and another layer between VM and container. These two extra network virtual-

<sup>☆</sup> This article belongs to Big Data Networking.

<sup>\*</sup> Corresponding author.

E-mail addresses: [pc.varma@gmail.com](mailto:pc.varma@gmail.com) (P. China Venkanna Varma), [tokalyankv@gmail.com](mailto:tokalyankv@gmail.com) (V. K.), [vallikumari@gmail.com](mailto:vallikumari@gmail.com) (V. Valli Kumari), [svraju.jntu@gmail.com](mailto:svraju.jntu@gmail.com) (S. Viswanadha Raju).

ization layers adds a significant delay to the RTT of data packets in the transmission.

We installed a Hadoop [4] cluster in our incubation lab and executed Hadoop benchmarks such as TestDFSIO-read [5], TestDFSIO-write [5], TeraSort [5] and TeraGen [5] against the number of the containers (Hadoop cluster nodes). We considered three use cases: Native, VMs and Containers while executing benchmarks. Each use case executed and analyzed the throughput of Hadoop cluster. TestDFSIO is a storage throughput test that is split into two parts: TestDFSIO-write writes 1 000 020 MB (about 1 TB) of data to HDFS, and TestDFSIO-read reads it back in. Because of the replication factor, the write test does twice as much I/O as the read test and generates substantial network traffic. These write tests spread across 140 map tasks. TeraSort sorts a large number of 100-byte records. It does considerable computation, networking, and storage I/O, and is often considered to be representative of real Hadoop workloads. It splits into three parts: generation, sorting, and validation. TeraGen creates the data and is similar to TestDFSIO-write except that significant computation is involved in creating the random data. The Map tasks write directly to HDFS so there is no reduce phase. TeraSort does the actual sorting and writes sorted data to HDFS in a number of partitioned files. A total of 280 Map tasks was used for TeraGen, 280 simultaneous (several thousand in all) Map tasks and 70 reduce [2] tasks for TeraSort, and 70 Map tasks for TeraValidate.

Based on test results, as the number of containers increases on a VM or a physical machine, there is a significant overhead on the big data operations. This overhead is due to a raise in the network IO RTT when the CPU is busy. If the CPU is busy, the network IO operation is in pending state and waiting for its scheduled slot to execute, this will increase the RTT of the network IO packet and leads to low network throughput. Ideally Hadoop throughput is linear, overall cluster throughput increases by adding new nodes. Based on our test results, we observed that the Hadoop cluster throughput is not linear with the nodes sharing the same CPU.

The outline of this paper is as follows: (1) Explore how container networking works? (Section 2). (2) Identify the important factors of the CPU context switch latency (Section 3). (3) Identify the effects of virtual network layers on the RTT (Section 3). (4) Experiment setup, big data environment and test cases (Section 4). (5) Analysis of virtual layer overhead (Section 5). (6) Analysis of Hadoop cluster performance (Section 6). (7) Conclusion (Section 7).

## 2. Docker networking

Docker has a software Ethernet bridge called “docker0” build on top of Linux “bridge-utils” [6]. When docker starts, it will create a default bridge “docker0” inside the Linux kernel and a virtual subnet to exchange the packets between containers and host. Docker randomly chooses the IP address range not used by host and defined by RFC 1918. When a container is created a pair of virtual interfaces will be initialized, those are similar to both the ends of a tube where a packet sent from one end will be delivered to another end. One virtual interface given to the container and another one hooked into docker0 to communicate with host and other containers. Fig. 1 describes the docker's docker0 and virtual interfaces. Communication between containers depends on the network topology and system firewall (iptables). By default docker0 bridge allows inter container communication. Docker never do changes to the iptables to allow a connection to docker0 bridge, it uses a flag “-iptables=false” when docker starts. Otherwise docker will add a default rule to the FORWARD chain with a blanket ACCEPT policy if you retain the default “-icc=true”, else will set the policy to DROP if “-icc=false”.

We observed that docker adds a network virtualization layer in the form of a software bridge which adds an extra hop in

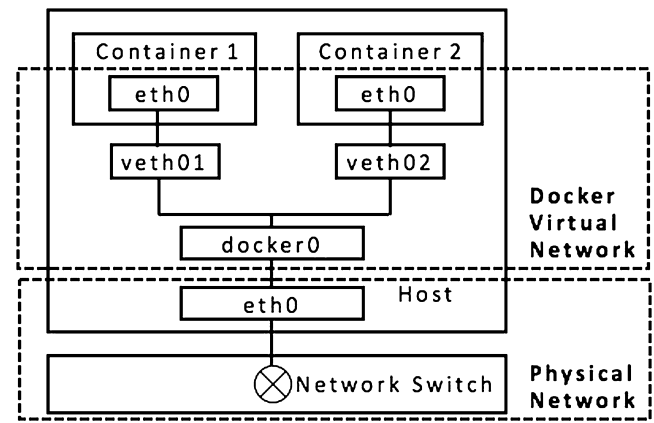


Fig. 1. Docker's docker0 bridge and virtual network layers.

the network path which adds some delay in the network packet transmission. We analyzed this overhead compared to native host network IO throughput.

## 3. Identify the important factors of the CPU context switch latency

“Context switch is the process of storing and restoring the state (context) of a process so that execution can be resumed from the same point at a later time” [9]. Context switch is applicable to a set of running processes in the same Central Processing Unit (CPU). There is a significant time required by the OS kernel to save execution state of the current running processes to memory, and then resume state of another process.

### 3.1. Process scheduling algorithms

Process Scheduling Algorithms plays a key role in the Context Switching. Each algorithm developed based on a use case. There is no ideal algorithm developed for all the known use cases. There is no choice given to end users to tune the Scheduling Algorithms. Improving the performance by tuning process scheduling algorithms is not in the scope of this paper.

### 3.2. Multiple-processor scheduling

Multiple-Processor Scheduling is a complicated concept. Main aim is to keep all the Processors busy to get the best throughput. But, if more and more context switch operations happened, that leads to less throughput. Improving the performance by tuning multiple-processor scheduling algorithms is also not in the scope of this paper.

### 3.3. Virtualization extensions

How to execute the various Virtual machine processes in the same CPU (Virtual CPU or Core), without compromising security? Is key for Virtualization extensions. To execute the process belongs to a Virtual Machine, hypervisor will add an extra layer on top of the existing Process Control Block (PCB) to identify the Virtual Machine and Process. Coding and decoding of this extra information will add a significant overhead in the process execution. Recent reports published that Round Trip Time (RTT) latency of a virtual machine process request is around 40 nm.

## 4. Experiment setup

Installed VMware ESXi v5.5 Operating System (OS) on a Dell PowerEdge server with Intel Core 2 Quad Processor Q9650, 128 GB

Download English Version:

<https://daneshyari.com/en/article/414509>

Download Persian Version:

<https://daneshyari.com/article/414509>

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