



HAL
open science

Docker Container Deployment in Fog Computing Infrastructures

Arif Ahmed, Guillaume Pierre

► **To cite this version:**

Arif Ahmed, Guillaume Pierre. Docker Container Deployment in Fog Computing Infrastructures. IEEE EDGE 2018 - IEEE International Conference on Edge Computing, Jul 2018, San Francisco, CA, United States. pp.1-8, 10.1109/EDGE.2018.00008 . hal-01775105

HAL Id: hal-01775105

<https://inria.hal.science/hal-01775105>

Submitted on 24 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Docker Container Deployment in Fog Computing Infrastructures

Arif Ahmed
Univ Rennes, Inria, CNRS, IRISA
F-35000 Rennes

Guillaume Pierre
Univ Rennes, Inria, CNRS, IRISA
F-35000 Rennes

Abstract—The transition from virtual machine-based infrastructures to container-based ones brings the promise of swift and efficient software deployment in large-scale computing infrastructures. However, in fog computing environments which are often made of very small computers such as Raspberry Pis, deploying even a very simple Docker container may take multiple minutes. We demonstrate that Docker makes inefficient usage of the available hardware resources, essentially using different hardware subsystems (network bandwidth, CPU, disk I/O) sequentially rather than simultaneously. We therefore propose three optimizations which, once combined, reduce container deployment times by a factor up to 4. These optimizations also speed up deployment time by about 30% in datacenter-grade servers.

Keywords-Docker, Container, Edge Cloud, Fog Computing.

I. INTRODUCTION

Fog computing extends datacenter-based cloud platforms with additional resources located in the immediate vicinity of the end users. By bringing computation where the input data was produced and the resulting output data will be consumed, fog computing is expected to support new types of applications which either require very low network latency to their end users (e.g., augmented reality applications) or produce large volumes of data which are relevant only locally (e.g., IoT-based data analytics).

Fog computing architectures are fundamentally different from classical cloud platforms: to provide computing resources in the physical vicinity of any end user, fog computing platforms must necessarily rely on very large numbers of small Points-of-Presence connected to each other with commodity networks, whereas clouds are typically organized with a handful of extremely powerful data centers connected by dedicated ultra-high-speed networks. This geographical spread also implies that the machines used in any Point-of-Presence may not be datacenter-grade servers but much weaker commodity machines. As a matter of fact, one option which is being explored is to use single-board computers such as Raspberry Pis for this purpose. Despite their obvious hardware limitations, Raspberry Pis offer excellent performance/cost/energy ratios and are well-suited to scenarios where the device’s physical size and energy consumption are important enablers for actual deployment [1], [2].

However, building a high-performance fog platform based on tiny single-board computers is a difficult challenge: in particular these machines have very limited I/O performance.

In this paper, we focus on the issue of downloading and deploying Docker containers in single-board computers. We assume that server machines have limited storage capacity and therefore cannot be expected to keep in cache the container images of many applications that may be used simultaneously in a public fog computing infrastructure.

Deploying container images can be painfully slow, in the order of multiple minutes depending on the container’s image size and network condition. However, such delays are unacceptable in scenarios such as a fog-assisted augmented reality application where the end users are mobile and new containers must be dynamically created when a user enters a new geographical area. Reducing deployment times as much as possible is therefore instrumental in providing a satisfactory user experience.

We show that this poor performance is not only due to hardware limitations. In fact it results from the way Docker implements the container’s image download operation: Docker exploits different hardware subsystems (network bandwidth, CPU, disk I/O) sequentially rather than simultaneously. We therefore propose three optimization techniques which aim to improve the level of parallelism of the deployment process. Each technique reduces deployment times by 10-50% depending on the content and structure of the container’s image and the available network bandwidth. When combined together, the resulting “Docker-pi” implementation makes container deployment up to 4 times faster than the vanilla Docker implementation, while remaining totally compatible with unmodified Docker images.

Interestingly, although we designed Docker-pi in the context of single-board computers, it also provides 23–36% performance improvements on high-end servers as well, depending on the image size and organization.

This paper is organized as follows. Section II presents the background and related work. Section III analyzes the deployment process and points out its inefficiencies. Section IV proposes and evaluates three optimizations. Finally, Section V discusses practicalities, and Section VI concludes.

II. BACKGROUND

A. Docker background

Docker is a popular framework to build, package, and run applications inside containers [3]. Applications are packaged in the form of *images* which contain a part of a file system

with the required libraries, executables, configuration files, etc. Images are stored in centralized repositories where they are accessible from any compute server. To deploy a container, Docker therefore first downloads the image from the repository and locally installs it, unless the image is already cached in the compute node. Starting a container from a locally-installed image is as quick as starting the processes which constitute the container’s application. The deployment time of any container is therefore dictated by the time it takes to download, decompress, verify, and locally install the image before starting the container itself.

1) *Image structure*: Docker images are composed of multiple layers stacked upon one another: every layer may add, remove, or overwrite files present in the layers below itself. This enables developers to build new images very easily by simply specializing pre-existing images.

The same layering strategy is also used to store file system updates performed by the applications after a container has started: upon every container deployment, Docker creates an additional writable top-level layer which stores all updates following a Copy-on-Write (CoW) policy. The container’s image layers themselves remain read-only. Table I shows the structures of the three images used in this paper.

2) *Container deployment process*: Docker images are identified with a name and a tag representing a specific version of the image. Docker users can start any container by simply giving its name and tag using the command:
`docker run IMAGE:TAG [parameters]`

Docker keeps a copy of the latest deployed images in a local cache. When a user starts a container, Docker checks its cache and pulls the missing layers from the docker registry before starting the container.

This work aims to better understand the hardware resource usage of the Docker container deployment process, and to propose alternative techniques to speed up the download and installation of the required image layers. We assume that the image cache is empty at the time of the deployment request: fog computing servers will most likely have very limited storage capacity so in this context we expect that cache misses will be the norm rather than the exception.

B. Related work

Many research efforts have recognized the potential of single-board devices for building fog computing infrastructures and have evaluated their suitability for handling cloud-like types of workloads. For instance, Bellavista *et al* demonstrated that even extremely constrained devices such as Raspberry Pis may be successfully used to build IoT cloud gateways. [4]. With proper configuration, these devices can achieve scalable performance with minimal overhead. However, the study assumes that the Docker container images are already cached in the local nodes. In contrast, we focus on the download-and-install part of the container deployment,

and show that simple modifications can significantly improve the performance of this operation.

A number of approaches propose to improve the design of Docker registries [5]. CoMIcon is a distributed Docker registry which distributes layers of an image among multiple nodes to increase availability and reduce the container provisioning time [6]. Distribution allows one to pull an image from multiple registries simultaneously, which reduces the average layer’s download times. Similar approaches rely on peer-to-peer protocols instead [7], [8]. However, distributed downloading relies on the assumption that multiple powerful servers are interconnected with a high-speed local-area network, and therefore that the main performance bottleneck is the long-distance network to a remote centralized repository. In the case of fog computing platforms, servers will be geographically distributed to maximize proximity to the end users, and they will rarely be connected to one another using high-capacity networks. As we discuss in the next section, the main bottleneck in fog computing nodes is created by hardware limitations of every individual node.

Another way to improve the container deployment time is to propose a new Docker storage driver. Slacker proposes to rely on a centralized NFS file system to share the images between all the nodes and registries [9]. The lazy pulling of the container image in the proposed model significantly improves the overall container deployment time. However, Slacker expects that the container image is already present in the local multi-server cluster environment; in contrast, a fog computing environment is made of large numbers of nodes located far from each other, and the limited storage capacity of each node implies that few images can be stored locally for future use. Besides, Slacker requires flattening the Docker images in a single layer. This makes it easier to support snapshot and clone operations, but it deviates from the standard Docker philosophy which promotes the layering system as a way to simplify image creation and updates. Slacker therefore requires the use of a modified Docker storage driver (with de-duplication features) while our work keeps the image structure unmodified and does not constraint the choice of a storage driver. We discuss the topic of flattening Docker images in Section V-A.

III. UNDERSTANDING THE DOCKER CONTAINER DEPLOYMENT PROCESS

To understand the Docker container deployment process in full details we analyzed the hardware resource usage during the download, installation and deployment of a number of Docker images on a Raspberry PI-based infrastructure.

A. Experimental setup

We monitored the Docker deployment process on a testbed which consists of three Raspberry Pi 3 machines connected to each other and to the rest of the Internet with 10 Gbps Ethernet [10]. The testbed was installed with

Table I
STRUCTURE OF THE DOCKER IMAGES

	Ubuntu	Mubuntu	BigLayers
6th layer	–	51 MB	–
5th layer	<1 MB	<1 MB	–
4th layer	<1 MB	<1 MB	62 MB
3rd layer	<1 MB	<1 MB	54 MB
2nd layer	<1 MB	<1 MB	64 MB
1st layer	46 MB	46 MB	52 MB
Total size	50 MB	101 MB	232 MB

Docker version 17.06. This setup also allowed us to emulate slower network connections – which are arguably representative of real fog computing scenarios – by throttling network traffic at the network interface level. We used the `tc` command to run experiments either with unlimited bandwidth, or with limits of 1 Mbps, 512 kbps or 256 kbps.

Table I depicts the images we used for this study. The first one simply conveys a standard *Ubuntu* operating system: it is composed of one layer containing most of the content, and four small additional layers which contain various updates of the base layer. We created a so-called *Mubuntu* image by adding an extra 51 MB layer which represents a typical application’s code and data which rely on the base Ubuntu image, following the incremental approach promoted by the Docker system. Finally, the *BigLayers* image is composed of four big layers which allow us to highlight the effect of the layering system on container deployment performance. We stored these images in the public Docker repository [11] so every experiment includes realistic download performance from a highly-utilized public repository.

We instrumented the testbed nodes to monitor the overall deployment time as well as the utilization of important resources during the container deployment process:

- *Deployment time*: We measured deployment times from the moment the deployment command is issued, to the time when Docker reports that the container is started.
- *Network activities*: We monitored the traffic from/to the Docker daemon (excluding other unrelated processes) on the Ethernet interface using `NetHogs` tool at a 1-second granularity.
- *Disk throughput*: We monitored the disk activity with the `iostat` Linux command which monitors the number of bytes written to or read from disk at a 1-second granularity.
- *CPU usage*: We monitored CPU utilization by watching the `/proc/stat` file at a 1-second granularity.

Every container deployment experiment was issued on an otherwise idle node, and with an empty image cache.

B. Monitoring the Docker container deployment process

Figure 1 depicts the results when deploying the three images using regular Docker. Figure 1(a) shows the deployment time of our three images in different network conditions: deploying the Ubuntu, Mubuntu and Biglayers

images with unlimited network bandwidth respectively takes 240, 333 and 615 seconds. Clearly, the overall container deployment time is roughly proportional to the size of the image. When throttling the network capacity, deployment times grow steadily as the network capacity is reduced to 1 Mbps, 512 kbps, and 256 kbps. For instance, deploying the Ubuntu container takes 6 minutes when the network capacity is reduced to 512 kbps. This is considerable with regards to the deployment efficiency one would expect from a container-based infrastructure. However, the interesting information for us is the *reason* why deployment takes so long, as we discuss next.

Figure 1(b) depicts the utilization of different hardware resources from the host machine during the deployment of the standard Ubuntu image. The red line shows incoming network bandwidth utilization, while the blue curve represents the number of bytes written to the disk and the black line shows the CPU utilization. The first phase after the container creation command is issued involves intensive network activities, which indicates that Docker is downloading the image layers from the remote image registry. By default Docker downloads up to three image layers in parallel. The duration of downloads clearly depend on the image size and the available network capacity: between 55 s and 110 s for the Ubuntu and Mubuntu images. During this phase, we observe no significant disk activity in the host machine, which indicates that the downloaded file is kept in main memory.

After the download phase, Docker extracts the downloaded image layers to the disk before building the final image of the application. The extraction of a layer involves two operations: decompression (which is CPU-intensive) and writing to the disk (which is disk-intensive). We observe that the resource utilization alternates between periods during which the CPU is busy (~40% utilization) while few disk activities are performed, and periods during which disk writes are the only notable activity of the system. We conclude that, after the image layers have been downloaded, Docker sequentially decompresses the image and writes the decompressed data to disk. When the image data is big, Docker alternates between partial decompressions and disk writes, while maintaining the same sequential behavior.

We see exactly the same phenomenon in Figure 1(c). However, here, the downloading of the first layer terminates before the other layers have finished downloading. The extraction of the first layer can therefore start before the end of the download phase, creating a small overlap between the downloading and extraction phases.

C. Critical observations

1) *Overall deployment time*: Container deployment involves three operations: checking the cache, pulling the image from the registry and starting the container. The longest operation is the `docker pull` operation, whereas

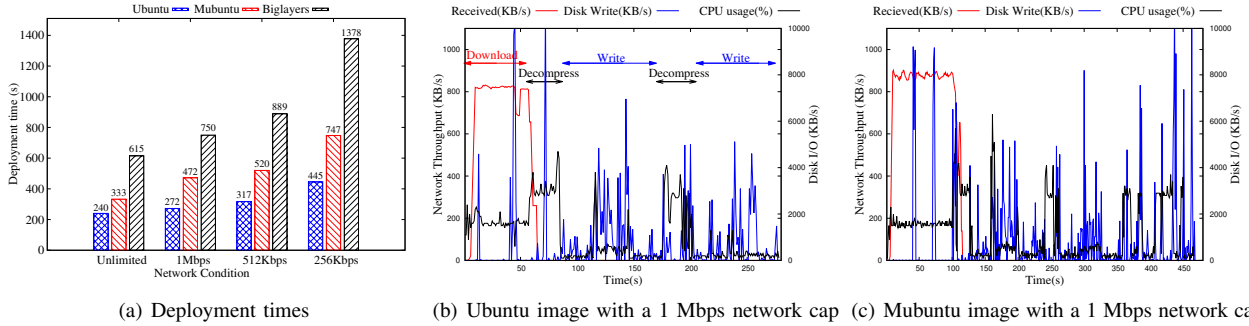


Figure 1. Deployment times and resource usage using standard Docker

the other two take a negligible amount of time. In this paper, we therefore focus on optimizing the pull operation.

2) *Pulling image layers in parallel*: By default, Docker downloads image layers in parallel with a maximum parallelism level of three. These layers are then decompressed and extracted to disk sequentially starting from the first layer. However, when the available network bandwidth is limited, downloading multiple layers in parallel will delay the download completion of the first layer, and therefore will postpone the moment when the decompression and extraction process can start. Therefore, delaying the downloading of the first layer ultimately leads to slowing down the extraction phase.

3) *Single-threaded decompression*: Docker always ships the image layers in compressed form, usually implemented as a *gzipped* tar file. This reduces the transmission cost of the image layers but it increases the CPU demand on the client node to decompress the images before extracting the image to disk. Docker decompresses the images via a call to the standard *gunzip.go* function, which happens to be single-threaded. However, even very limited machines usually have several CPU cores available (4 cores in the case of a Raspberry Pi 3). The whole process is therefore bottlenecked by the single-threaded decompression. As a result the CPU utilization never grows beyond $\sim 40\%$ of the four cores of the machine, wasting precious computation resources which may be exploited to speed up image decompression.

4) *Resource under-utilization*: The standard Docker container deployment process under-utilizes the available hardware resources. Essentially, deploying a container begins with a network-intensive phase during which the CPU and disk are mostly idle. It then alternates between CPU-intensive decompression operations (during which the network and disk are mostly idle) and I/O-intensive image extraction operations (during which the network and CPU are mostly idle). The only case where these operations slightly overlap are images such as Mubuntu and BigLayers when the decompress and extraction process of the first layer can start while the last images are still being downloaded.

This resource under-utilization is one of the main reason for the poor performance of the overall container deployment

process. The main contribution of this paper is to show how one may reorganize the Docker deployment process to maximize resource utilization during deployment.

IV. OPTIMIZING THE CONTAINER DEPLOYMENT PROCESS

To address the inefficiencies presented in the previous section we propose and evaluate three optimizations which deal with different issues in the deployment process. We can therefore combine them all together, which brings significant performance improvement.

A. Sequential image layer downloading

Simultaneous downloading of multiple images obviously aims to maximize the overall network throughput. However, it has negative effects because the decompress and extraction phases of each image layer must take place sequentially to preserve the Copy-on-Write policy of Docker storage drivers. The decompress & extract phase can start only after the first layer has been downloaded. Downloading multiple image layers in parallel will therefore delay the download completion of the first layer because this download must share network resources with other image layer downloads, and will therefore also delay the moment when the first layer can start its decompress & extract phase. We therefore propose to download image layers sequentially rather than concurrently.

Figure 2 illustrates the effect of downloading layers sequentially rather than in parallel. In both cases, three threads are created to handle the three image layers. However, in the first option the downloads take place in parallel whereas the only required inter-thread synchronization requires that the decompression and extraction of layer n can start only after the decompression and extraction of layer $n - 1$ has completed. In sequential downloading, the second layer starts downloading only when the first download has completed, which means that it takes place while the first layer is being decompressed and extracted to disk. This allows the first-layer extraction to start sooner and it also increases resource utilization because the download and

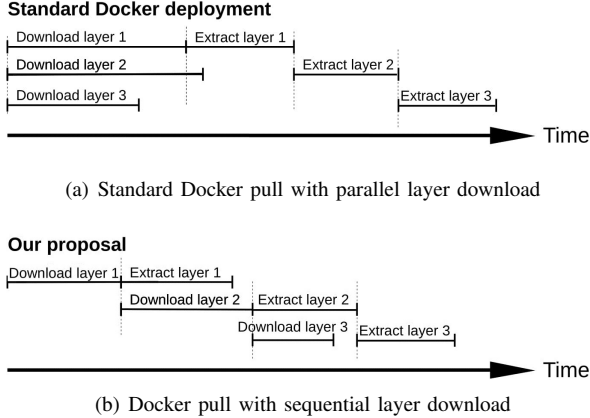


Figure 2. Standard and sequential layer pull operations

the decompress & extract operations make intensive use of different part of the machine’s hardware.

Implementing sequential downloading can be done easily by setting the “max-concurrent-downloads” parameter to 1 in the `/etc/docker/daemon.json` configuration file.

Figure 3(a) depicts the resource usage of the host machines when deploying the Mubuntu image with sequential downloading and a 1 Mbps network capacity. We observe that after the downloading of layer 1 has completed, the utilization of hardware resources is much greater, with in particular a clear overlap between periods of intensive network, CPU and I/O resources. Also we can observe that the decompression of the first layer (visible as the first spike of CPU utilization) takes place sooner than in Figure 1(c).

Figure 3(b) compares the overall container deployment times with parallel and sequential downloads in various network conditions. When the network capacity is unlimited the performance gains in the deployment of the Ubuntu, Mubuntu and BigLayers images are 3%, 4.2% and 6% respectively.

However, the performance gains grow steadily as the available network bandwidth gets reduced. With a bandwidth cap of 256 kbps, sequential downloading brings improvements of 6% for the Ubuntu image, 10% for Mubuntu and 12% for BigLayers. This is due to the fact that slower network capacities exacerbate the duration of the download phases and increases the delaying effect of parallel layer downloading.

B. Multi-threaded layer decompression

By default, Docker uses the `gunzip.go` library to decompress the downloaded image layers before extracting to the disk. However, this function is single-threaded, which implies that the CPU utilization during decompression never exceeds 40% of the four available cores in the Raspberry Pi machine. We therefore propose to replace the single-threaded `gunzip.go` library with a multi-threaded implementation so

that all the available CPU resources may be used to speed up this part of the container deployment process.

We use `pgzip`, which is a multi-threaded implementation of the standard `gzip/gunzip` functions [12]. Its functionalities are exactly the same as those of the standard `gzip`, however it splits the work between multiple independent threads. When applied to large files of at least 1 MB, this can significantly speed up decompression.

Figure 4(a) depicts the deployment time of a single-layered image while using various numbers of threads for decompression. When `pgzip` uses a single thread, the performance and CPU utilization during decompression are very similar to the standard `gunzip` implementation. However, when we increase the number of threads, the overall container deployment time decreases from 154 s to 136 s. At the same time, the CPU utilization during decompression steadily increases from 40% to 71% of the four available CPU cores. If we push beyond 12 threads, no additional gains are observed. We clearly see that the parallel decompression does not scale linearly, as it is not able to exploit the full capacity of the overall CPU: this is due to the fact that `gzip` decompression must process data blocks of variable size so the decompression operation itself is inherently single-threaded [13]. The benefit of multi-threading decompression is that other necessary operations during decompression such as data buffering and CRC verification can be delegated to other threads and moved out of the critical path.

Figure 4(b) shows the effect of using parallel decompression when deploying Mubuntu images with 12 threads. We observe that the CPU utilization is greater during the decompression phases than with standard Docker, in the order of 70% utilization instead of 40%. Also, the decompression phase is notably shorter.

Figure 4(c) compares the overall container deployment times with parallel decompression against that of the standard Docker. The network performance does not influence the time of the decompression phase so we conducted the evaluation only with an unlimited network capacity. The performance gain from multi-threaded decompression is similar for all three images, in the order of 17% of the overall deployment time.

C. I/O pipelining

Despite the sequential downloading and the multi-threaded decompression techniques, the container deployment process still under-utilizes the hardware resources. This is due to the sequential nature of the workflow which is applied to each individual layer: each layer is first downloaded in its entirety, then it is decompressed entirely, then it is extracted to disk. This requires Docker to keep the entire decompressed layer in memory, which can be significant considering that a Raspberry Pi 3 has only 1 GB of main memory. Also, it means that the first significant disk activity can start only after the first layer has been fully




Figure 3. Resource usage and deployment time with sequential layers downloading

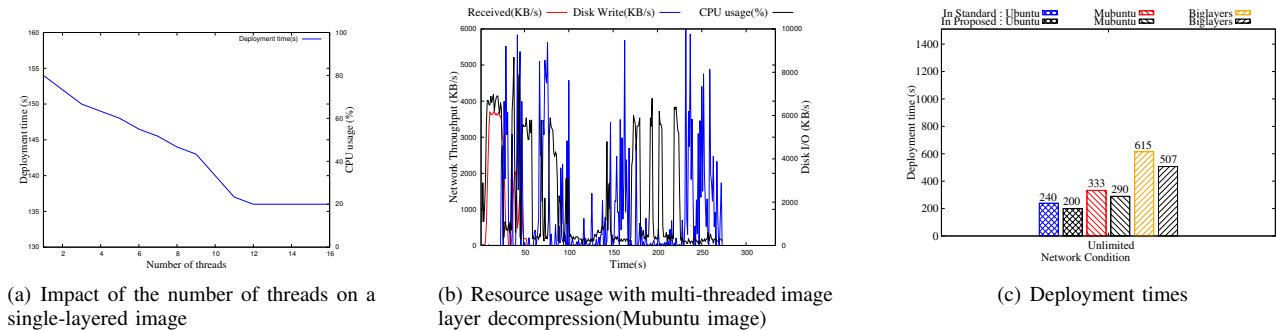


Figure 4. Evaluation of multi-threaded image decompression

downloaded and decompressed. Similarly, Docker necessarily decompresses and extracts the last layer to disk while the networking device is mostly inactive.

However, there is no strict requirement for the download, decompress and extraction of a single layer to take place sequentially. For example, decompression may start right after the first bytes of the compressed layer have been downloaded. Similarly, extracting the layer may start immediately after the beginning of the layer image has been decompressed.

We therefore propose to reorganize the download, decompression and extraction of a single layer in three separate threads where each thread pipelines data to the next as soon as some data is available. In Unix shell syntax this essentially replaces the sequential “download; decompress; crc-check; extract” command with the concurrent “download | decompress | crc-check | extract” command. Since we stream the incoming downloaded data without buffering the entire layer, the extraction can start writing content to disk long before the download process has completed.

We implemented pipelining using the *io.pipe()* GO API [14], which creates a synchronized in-memory pipe between an *io.reader(pgzip/decompress)* and an *io.writer(network/download)* without internal buffering. However, we must be careful about synchronizing this

process between multiple image layers: for example, if we created an independent pipeline for each layer separately, the result would violate the Docker policy that layers must be extracted to disk sequentially, as one layer may overwrite a file which is present in a lower layer. If we extracted multiple layers simultaneously we could end up with the wrong version of the file being given to the container. Rather than building complex synchronization mechanisms, we instead decided to rely on Docker’s sequential downloading feature already discussed in Section IV-A. When a multi-layer image is deployed, this imposes that layers are downloaded and extracted one after the other, while using the I/O pipelining technique within each layer.

Figure 5 evaluates the I/O pipelining technique using a single-layer image. We can see that the pipelined version is roughly 50% faster than its standard counterpart: in the standard deployment, resources are used one after the other: first network-intensive download, then CPU-intensive decompression, then finally disk-intensive image creation. In the pipelined version all operations take place simultaneously, which better utilizes the available hardware and significantly reduces the container deployment time.

D. Docker-pi

The three techniques presented previously address different issues. Sequential downloading of the image layers speeds up the downloading of the first layer in slow network

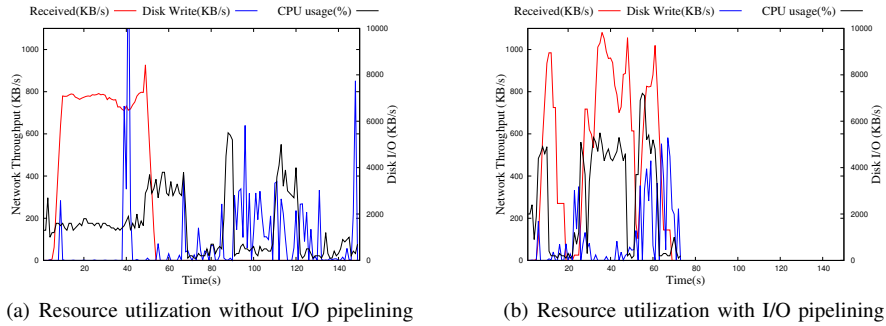


Figure 5. Evaluation of I/O pipelining

environments. Multi-threaded decompression speeds up the layer decompression by utilizing multiple CPU cores. Finally, I/O pipelining speeds up the deployment of each layer by conducting the download, decompress and extraction processes simultaneously, while avoiding having to keep large amounts of data in memory during the deployment process. We therefore propose *Docker-pi*, an optimized version of Docker which combines the three techniques to optimize container deployment on single-board machines.

Figure 6(a) shows the resource usage while deploying the Mubuntu image using Docker-pi. We can clearly see that the networking, CPU and disk resources are used simultaneously and have a much greater utilization than with the standard Docker implementation. In particular, the CPU and disk activities start very early after the first few blocks of data have been downloaded.

Figure 6(b) highlights significant speedups compared to vanilla Docker: with no network cap, Docker-pi is 73% faster than Docker for the Ubuntu image, 65% faster for Mubuntu and 58% faster for BigLayer. When we impose bandwidth caps the overall deployment time becomes constraint by the download times, while the decompression and extraction operations take place while the download is taking place. In such bandwidth-limited environments the deployment time therefore cannot be reduced any further other than by pre-fetching images before the container deployment command is issued.

The reason why the gains are slightly lower for the Mubuntu and BigLayers images is that the default Docker download concurrency degree of 3 already makes them benefit from some of the improvements we proposed in Docker-pi. If we increase the concurrency degree of Docker to 4, the BigLayers image deploys in 644 s whereas Docker-pi needs only 207 s, which represents 68% improvement.

V. DISCUSSION

A. Should we flatten all Docker images?

Flattening Docker images may arguably provide performance improvement in the deployment process. Indeed, multiple image layers may contain successive versions of

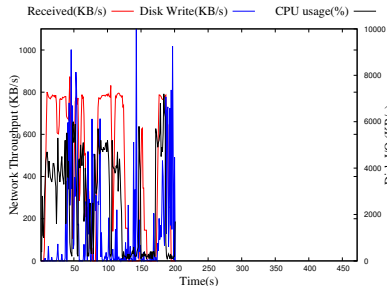
the same file whereas a flattened image contains only a single version of every file. A flattened image is therefore slightly smaller than its multi-layered counterpart. Systems such as Slacker actually rely on the fact that images have been flattened [9]. On the other hand, Docker-pi supports both flattened images and unmodified multi-layer images. We however do not believe that flattening all images would bring significant benefits.

Docker does not provide any standard tool to flatten images. This operation must be done manually by exporting an image with all its layers, and re-importing the result as a single layer while re-introducing the startup commands from all the initial layers. The operation must be redone every time any update is made in any of the layers. Although this process could be integrated in an image build workflow, it contradicts the Docker philosophy which promotes incremental development based on image layer reusability.

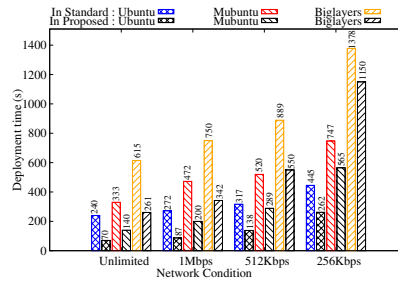
In a system where many applications execute concurrently, one may reasonably expect many images to share at least the same base layers (e.g., Ubuntu) which produce a standard execution environment. If all images were flattened this would result in large amounts of redundancy between different images, creating the need for sophisticated de-duplication techniques [9]. On the other hand, we believe that the layering system can be seen as a domain-specific form of de-duplication which naturally integrates in a developer’s devops workflow. We therefore prefer keeping docker images unmodified, and demonstrated that container deployment can be made extremely efficient without the need for flattening images.

B. Does Docker-pi work also for powerful server machines?


Although we designed Docker-pi for single-board machines, the inefficiencies of vanilla Docker also exist in powerful server environments. We therefore evaluated the respective performance of Docker and Docker-pi in the Grid’5000 testbed which is commonly used for research on parallel and distributed computing including Cloud, HPC and Big Data [15]. We specifically used a Dell PowerEdge C6220 server equipped with two 10-core Intel Xeon E5-



(a) Resource usage when deploying Mubuntu using Docker-pi on a RPI machine



(b) Deployment times in RPI machines



(c) Deployment times in a server-grade machine

Figure 6. Evaluation of Docker-pi in RPI and powerful server machines

2660v2 processors running at 2.2GHz, 128 GB of main memory and two 10 Gbps network connections.

Figure 6(c) compares the deployment times of Docker and Docker-pi with our three standard images. Obviously container deployment is much faster in this environment than in Raspberry PIs. However, here as well Docker-pi provides respectable performance improvement in the order of 23–36%. In this powerful server the network and CPU resources cannot be considered as bottlenecks so the sequential layer downloading and multi-threaded decompression techniques bring little improvement compared to the standard Docker. On the other hand, the sequential nature of the download/decompress/extract process is still present regardless of the hardware architecture, so the I/O pipelining technique brings similar performance gains as with the Raspberry PI.

VI. CONCLUSION

The transition from virtual machine-based infrastructures to container-based ones brings the promise of swift and efficient software deployment in large-scale computing infrastructures. However, this promise is not being held in fog computing platforms which are often made of very small computers such as Raspberry PIs, and where deploying even a very simple Docker container may take multiple minutes.

We identified three sources of inefficiency in the Docker deployment process and proposed three optimization techniques which, once combined together, speed up container deployment roughly by a factor 4. Last but not least, we demonstrated that these optimizations also bring significant benefits in regular server environments.

This work eliminates the unnecessary delays that take place during container deployment. Depending on the hardware, deployment time is now dictated only by the slowest of the three main resources: network bandwidth, CPU, or disk I/O. As hardware will evolve in the future the bottleneck may shift from one to the other. But, regardless of the specificities of any particular machine, Docker-pi will exploit the available hardware to its fullest extent.

REFERENCES

- [1] A. van Kempen *et al*, “MEC-ConPaaS: An experimental single-board based mobile edge cloud,” in *IEEE Mobile Cloud Conference*, 2017.
- [2] W. Hajji and F. P. Tso, “Understanding the performance of low power Raspberry Pi cloud for big data,” *Electronics*, vol. 5, no. 2, 2016.
- [3] Docker Inc., “Docker: Build, ship, and run any app, anywhere,” <https://www.docker.com/>.
- [4] P. Bellavista *et al*, “Feasibility of fog computing deployment based on Docker containerization over RaspberryPi,” in *Proc. ICDCN*, 2017.
- [5] A. Anwar *et al*, “Improving Docker registry design based on production workload analysis,” in *Proc. FAST*, 2018.
- [6] S. Nathan *et al*, “CoMICon: A co-operative management system for docker container images,” in *Proc. IC2E*, 2017.
- [7] I. Babrou, “Docker registry bay,” <https://github.com/bobrik/bay>, 2016.
- [8] M. Ma *et al*, “Dockyard – container and artifact repository,” <https://github.com/Huawei/dockyard>, 2017.
- [9] T. Harter *et al*, “Slacker: Fast distribution with lazy Docker containers,” in *Proc. FAST*, 2016.
- [10] “The mobile edge cloud testbed at IRISA Myriads team,” YouTube video, Jan. 2017, <https://www.youtube.com/watch?v=7uLkLitiSPo>.
- [11] Docker Inc., “Docker hub,” <https://hub.docker.com/>.
- [12] K. Post, “klauspost/pgzip: Go parallel gzip (de)compression,” <https://github.com/klauspost/pgzip>, Oct. 2017.
- [13] Evangelia Sitaridi *et al*, “Massively-parallel lossless data decompression,” in *Proc. ICPP*, 2016.
- [14] The Go Authors, “Source file src/io/pipe.go,” <https://golang.org/pkg/io/#Pipe>, 2018.
- [15] D. Balouek *et al*, “Adding virtualization capabilities to the Grid’5000 testbed,” in *Cloud Computing and Services Science*, I. I. Ivanov *et al*, Ed. Springer, 2013, vol. 367.