

ARTICLE TYPE

The State-of-the-Art in Container Technologies: Application, Orchestration and Security

Emiliano Casalicchio^{2,1} | Stefano Iannucci³

¹Department of Computer Science, Sapienza University of Rome, Rome, Italy

²Department of Computer Science, Blekinge Institute of Technology, Karlskrona, Sweden

³Department of Computer Science and Engineering, Mississippi State University, Mississippi, USA

Correspondence

Emiliano Casalicchio, Email:
emiliano.casalicchio@uniroma1.it

Stefano Iannucci, Email:
stefano@cse.msstate.edu

Abstract

Containerization is a lightweight virtualization technology enabling the deployment and execution of distributed applications on cloud, edge/fog and Internet-of-Things platforms. Container technologies are evolving at the speed of light and there are many open research challenges.

In this paper, an extensive literature review is presented that identifies the challenges related to the adoption of container technologies in High Performance Computing, Big Data analytics and geo-distributed (Edge, Fog, Internet-of-Things) applications. From our study, it emerges that performance, orchestration and cyber-security are the main issues. For each challenge, the state-of-the-art solutions are then analyzed.

Performance is related to the assessment of the performance footprint of containers and comparison with the footprint of virtual machines and bare metal deployments, the monitoring, the performance prediction, the I/O throughput improvement. Orchestration is related to the selection, the deployment, and the dynamic control of the configuration of multi-container packaged applications on distributed platforms. The focus of this work is on run-time adaptation. Cyber-security is about container isolation, confidentiality of containerized data, and network security.

From the analysis of 97 papers, it came out that the state-of-the-art is more mature in the area of performance evaluation and run-time adaptation rather than in security solutions. However, the main unsolved challenges are: I/O throughput optimization, performance prediction, multi-layer monitoring, isolation and data confidentiality (at rest and in transit).

KEYWORDS:

Run-time Adaptation, cyber-security, performance, Container, Docker, Cloud computing, Orchestration, Auto-scaling, Big Data, HPC, Distributed computing, Edge computing, Fog computing, Internet-of-Things

1 | INTRODUCTION

Nowadays, distributed applications and infrastructures are moving from being Virtual Machine (VM) centric to container centric. The technology of containers is strongly supported by PaaS providers^{1,2}, IaaS providers³ and Internet Service Providers⁴. Moreover, container technologies are used to deploy large scale applications in challenging fields, such as, big data analytics^{5,6}, scientific computing^{7,8}, edge computing^{9,10} and Internet-of-Things (IoT)^{11,12}.

Container's keys to success are many: an easier management of the life-cycle of distributed applications¹³; a negligible overhead for when they run either on bare-metal or virtual servers^{14,15,16,17}, and their time to start, re-start and stop, which is reduced up to an order of magnitude

with respect to VMs. Last but not least, containers enable application portability, which is where hypervisor-based virtualization failed. Docker, Kubernetes and Cloudify, the de-facto standard container technologies, are compliant with the cloud portability framework TOSCA^{18,19}. Moreover, portability is also supported by the Open Container Initiative (opencontainers.org), created to define a standard image format and a standard run-time environment for containers.

Because of these many advantages, the container technology landscape is developing and expanding at the speed of light, but there are still many challenges to be solved, for example: the reduction of networking and I/O overheads compared to hypervisors; the secure resource sharing and isolation to enable multi-tenancy; methodologies and tools for multi-layer monitoring, functional to sophisticated run-time adaptation algorithms; advanced run-time adaptation capabilities (e.g. scale-out/in, scale-up/down and migration) and high availability support.

In such fast evolving and challenging scenario, a wide spectrum review of the literature which identifies mature and early stage research results is missing. To fill that gap, we analyzed research work using the following methodology: firstly, we scrutinized the literature to identify the most demanding application fields adopting container technologies (i.e. High Performance computing, Big data analytics and geo-distributed applications) and their related challenges (i.e. performance, run-time adaptation, and security). Then, we collected and analyzed additional literature to identify the solutions proposed to solve these challenges, and to point out the open issues. As a result of the above methodology, this survey presents a classification and analysis of the literature along the following four categories:

- **Applications.** This category clusters research papers that use containers in High Performance Computing, Big Data analytics and geo-distributed (Fog, Edge, IoT) applications.
- **Performance comparison.** This category groups research works that evaluate the performance of containers under different workloads and on different host environments. All these papers investigated the advantages and drawbacks of using containers as an alternative to VMs or stacked on top of VMs.
- **Orchestration.** This category includes studies that address the challenge of run-time adaptation of containerized applications¹ and container platforms to guarantee performance and high availability.
- **Security.** This category contains research results addressing cyber-security challenges. The focus of the literature is on containers isolation, containers image layers security, storage volumes security, and network security.

To the best of our knowledge, and at the time of writing this paper, this is the first extensive literature review on the subject, and it largely extends the previous published survey on container orchestration²⁰. A research work quite close to this paper is a survey on the micro-service software architectures²¹. However, that study investigates architectural patterns and does not mention challenges related to the implementation and deployment of micro-services with containers.

The remaining of this paper is organized as in what follows. Section 2 provides background on application and system containers, container managers and container orchestrators. The literature's taxonomy, that delineate the above mentioned categories, is described in details in Section 3. Sections 4, 5, 6 and 7 analyze the state-of-the-art works for the categories introduced above. To conclude, Section 8 summarizes how the challenges in the application field are addressed by the literature and put in evidence the open issues.

2 | BACKGROUND

The idea of container dates back to 1992²². At the base of container technology there is the concept of *cgroup* and Linux *namespace*²³. To understand the landscape of containers technology, four main concepts are important: application and system containers, container managers and container orchestrators.

2.1 | Application and System Containers

Application and system containers build on the same technologies and concepts. The only difference is in their use: a system container runs a full operating system, while an application container is intended for the deployment of an application or an application component. For example, the Docker development guidelines and best practices recommend to run a microservice, rather than a full application, in a container. An application

¹A *containerized* application refers to an application packed in a container. It could be a monolithic application or an application composed of independent components. In the later case, each component could be packed in an independent container

TABLE 1 Container technologies considered in this work. The Type acronyms are defined as follow: Application Container (AC); System Container (SC), Container Manager (CM), Orchestration Framework (OF)

Technology	Type				Url
	SC	AC	CM	OF	
Linux Container (LXC)	Y	Y			https://linuxcontainers.org
OpenVZ	Y		Y		https://openvz.org
Windows Hyper-V Container (WHC)		Y			https://docs.microsoft.com/en-us/virtualization/windowscontainers
Docker		Y	Y		https://www.docker.com/
Windows Server Container (WSC)		Y			https://docs.microsoft.com/en-us/virtualization/windowscontainers
rkt			Y		https://coreos.com/rkt
LXD			Y		https://linuxcontainers.org
Oracle Solaris Container	Y	Y	Y	Y	http://www.oracle.com/technetwork/server-storage/solaris/containers-169727.html
Amazon EC2 Container Service (ECS)			Y	Y	https://aws.amazon.com/ecs
Google Container Engine (GCE)			Y	Y	https://cloud.google.com/container-engine
Microsoft Azure Container Service (ACS)			Y	Y	https://azure.microsoft.com/en-us/services/container-service
Joyent Triton			Y	Y	https://www.joyent.com/
Kubernetes				Y	https://kubernetes.io
Swarm				Y	https://www.docker.com
Marathon				Y	https://mesosphere.github.io/marathon
Cloudify				Y	http://cloudify.co

container is built up from a series of image layers leveraging the Union File System service²⁴ and the Copy-on-Write technique, that allows a filesystem to appear as writable, but without actually allowing writes to change the filesystem content. Usually, an application developer can start from a base image (e.g. the operating system kernel and default libraries). Afterwards, the application can be copied in a new layer (or possibly in two – a bottom layer with the source code and libraries and a top layer with the executable). In the same way, a system container incorporates a base operating system image in the bottom layer, which could be customized with the addition of libraries, tools, and data in new layers. After the multilayer image is built, all the layers are read-only except the last one, which is readable and writable, but not persistent. That means the content is lost when the container is deleted.

To persistently store data, a containerized application needs to mount a virtual disk (Docker volumes in the Docker jargon). Docker *volumes* can be mounted in multiple containers simultaneously, allowing concurrent reads and writes. Concurrency can be either managed at the application level, or with a distributed filesystem, such as, GlusterFS²⁵ and Ceph²⁶. Bind-mount is an alternative solution for persistent storage that allows to store data in the host filesystem. Its data is accessible both from the container and the host. Such solution is not portable and introduces security issues²⁷ but, on the other end, it could improve the I/O performance.

Table 1 lists the container technologies considered in this paper. The most widely used application container technology is Docker, a multi platform solution designed for Linux, OSX and Windows. Docker extends LXC with a kernel and application-level API to facilitate container management and the management of containerized applications. LXC also can be used as application container. rkt is a solution designed to run application containers in a cloud native environment, i.e. CoreOS (Container Linux). The Oracle Solaris container solution offers more isolation than Docker. Windows Server Container (WSC) is the Microsoft version of application containers. Concerning system containers, Linux Container and OpenVZ allow to share a linux kernel among containers built from the same base image. LXC uses a standard kernel while OpenVZ uses a specialized kernel. Windows Hyper-V Container (WHC) is the Microsoft version of system containers; however, in order to guarantee a strong environment isolation, WHCs do not share the kernel among containers.

2.2 | Container Manager

A container manager is a framework providing a set of API to easily manage all the life cycle of the container¹³ (cf. Fig. 1), which consists of the following phases: Acquire, Build, Deliver, Deploy, Run and Maintain. *Acquire* is related to select and download a container image from a container repository (e.g. Docker Hub), and to use it as the base layer on top of which the application will be containerized. In the *Build* phase, all the application components, libraries and possibly data are packaged in the container image. Afterwards, the image is published in a public/private

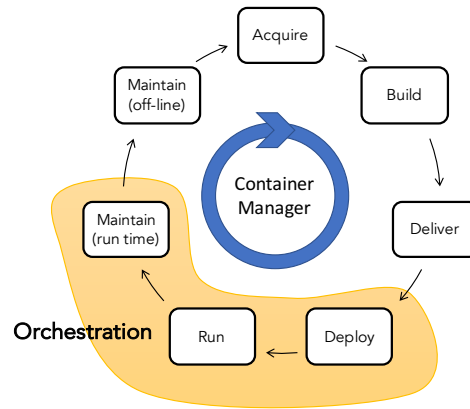


FIGURE 1 The container lifecycle. A container manager provides API to support at least from the acquire to the run phases, while the Orchestrator allows to automatize the deployment, run and maintain (run-time) phases.

repository. *Delivery* concerns to bring the application in production, and could also include a vulnerability analysis step. *Deploy* is about deploying the application and maintaining it up-to-date, for instance, by using a continuous delivery model. The *run* phase sets the management system and runtime environment (e.g. scaling policy, health check, recovery policies). The final step is to *maintain* the application at run-time and off-line. The application behaviour is monitored and, when failures are triggered, the system tries to manage them at run-time, e.g., by re-starting a container. Off-line, the application is de-bugged, to find the root-cause of fault, and fixed. Then, the process goes again through the Acquire or Build phases to create a new containerized version of the application that will be delivered, deployed and maintained.

Container managers can be classified as either on-premise or managed solutions. On-premise solutions need to be installed, configured and managed on private datacenters or on virtual machines running in the cloud or geo-distributed infrastructures. Managed solutions are instead offered by cloud providers as a service, and need only to be partially configured. Among the technologies listed in Table 1, Docker, LXD, OpenVZ and `rkt` are examples of on-premise solutions. Examples of managed solutions are: Google Container Engine, Microsoft Azure Container Service and Amazon ECS. Docker has been designed as a container management system and is becoming the de-facto standard. For example: Windows Server container and Hyper-V container both can be managed with Docker; GCE, ACS and ECS support Docker containers, and `rkt` offers APIs for easy application container management. Concerning system containers, LXD is the manager for LXC. OpenVZ also provide container management APIs.

2.3 | Container Orchestration

Container orchestration allows cloud and application providers to define how to select, deploy, monitor, and dynamically control the configuration of multi-container packaged applications in the cloud²⁸. Container orchestration is concerned with the management at runtime to support the Deploy, Run and Maintain phases (cf. Figure 1). Container orchestrators usually offer the following main features (cf. Table 2): resource limit control, scheduling, load balancing, health check, fault tolerance, auto-scaling.

Resource limit control allows to reserve a specific amount of CPU and memory for a container; these constraints can be used to make scheduling decisions and to limit the resource contention among containers. Resource limit control features leverage the equivalent mechanisms offered by the container manager. Indeed, while a container can use all the resource available in the underlying system, container managers provide APIs to limit the amount of memory and CPU used and the specific CPU used. However, more sophisticated mechanisms can be implemented. For example, Kubernetes offers the possibility to reserve a specific amount of a resource, e.g., the CPU, with `spec.containers[].resources.requests.cpu` and/or to limit the use of a reserved resource, e.g. the CPU with `spec.containers[].resources.limits.cpu`. If we consider the example of the CPU, `request` allows to allocate a specific amount of CPU (e.g. 0.5, 1, 2 cores) while the `limit` allows to specify the maximum amount of time the requested CPU portion is used over 100ms. The request quota is stringent: if a node is running two containers that specified a CPU request of 0.5 and 0.3, a new container can be allocated only if it has a CPU request less than or equal to 0.2. The limit quota is instead not stringent, that is a container can exceed, for some time, the quota.

Scheduling defines the policy used to place the desired amount of containers on the desired nodes at a give time instant. Scheduling can be done either on the basis of resource constraints, or node affinity, or both of them. More sophisticated schedulers can usually be integrated as external components (the custom feature in Tab. 2).

TABLE 2 Features implemented by different orchestration framework.

Orchestration Features	Orchestration Frameworks			
	Swarm	Kubernetes	Cloudify	Marathon
Resource limit control	CPU, Memory		N.A.	N.A.
Scheduling	Resource constrained, Node affinity, Custom			
Load balancing	Round-Robin, Custom		N.A.	Custom
Health check	Open port connections (UDP, TCP, SSL), HTTP req/res			
Fault tolerance	Replica control, High availability controller			N.A.
Auto-scaling	Custom	Custom, Threshold based (CPU)		Custom

The load balancer does the work of distributing the load among multiple container instances. Round-Robin is the default implemented policy. More complex policies can be provided by external load balancers (the custom feature in Tab 2).

Health-check is achieved controlling if a container is capable to answer requests. Usual implementations make use of TCP/UDP/SSH connection attempts to open ports, and HTTP request/response checking.

Fault tolerance can be implemented as replica control and/or high availability controller. Replica control allows to specify and maintain a desired number of containers. Health check is used to determine when a faulty container should be destroyed and a new one launched to maintain the target number of replicas. The high availability controller allows to configure multiple orchestration managers to have always control on the application in case an orchestrator node fails or is overloaded. The same technique used to create a high availability controller can be used to implement a scalable controller.

Auto-scaling allows to automatically add and remove containers. The implemented policies are usually based on thresholds (e.g., on CPU and memory utilization), but in some cases it is possible to plug-in more sophisticated autoscalers or to define custom autoscaling policies (the custom feature in Tab. 2).

In the landscape of container orchestration frameworks, the choice is between on-premise solutions and managed solutions (as for the container managers). The main on-premise solutions are: Docker Swarm, the native Docker orchestrator offering clustering functionality for Docker containers, which lets system administrator turn a group of Docker engines into a single, virtual Docker engine. Kubernetes is an orchestration system for Docker containers capable to handle scheduling and to manage workloads based on user-defined parameters. Mesosphere Marathon is a container orchestration framework for Apache Mesos. It offers key features for running applications in a clustered environment. Cloudify is a cloud orchestration framework that enables the modeling of applications and services and automates their entire life cycle. Cloudify is TOSCA compliant and could be used to deploy Docker containers, Docker Swarm clusters, or Kubernetes clusters. Container as a service solutions offer also orchestration services. Examples are: Google Container Engine, Amazon Elastic Container Service, Microsoft Azure Container Service.

3 | STATE-OF-THE-ART TAXONOMY

The focus of this survey is schematized in Figure 2. The literature is classified using a taxonomy composed of the four categories described in Section 1 and several sub-categories: *Application's* sub-categories are: High Performance Computing, Big Data analytics, and geo-distributed (Fog, Edge and Internet-of-Things) applications. Although any application can be deployed using containers, the above mentioned are the most challenging domains. The *Performance comparison* category doesn't need to be further detailed by sub categories because all the studies compare the performance of container solutions w.r.t. system virtualization and native environments. *Orchestration* sub-categories are: Performance Monitoring, Characterization and Prediction, Auto-scaling, High Availability, and Scheduling. Energy efficiency is another important aspect, and it is orthogonal to the Performance Evaluation and Orchestration categories. *Security* sub-categories collect research results about container's isolation, image layers and volumes encryption, and network security.

We decided to organize the literature as above because the analysis of application fields is valuable to understand challenges and open issues. And the study of the other categories, orthogonal to the first one, present the solution proposed to address the challenges and put in evidence what areas and results are mature, and what need more effort.

This survey presents the result of an extensive, although not omni-comprehensive, literature review that analyzes and summarizes the contribution of 97 verified sources selected among about 110 retrieved from the main research databases, like: ACM digital library, IEEE Explorer, Scopus,

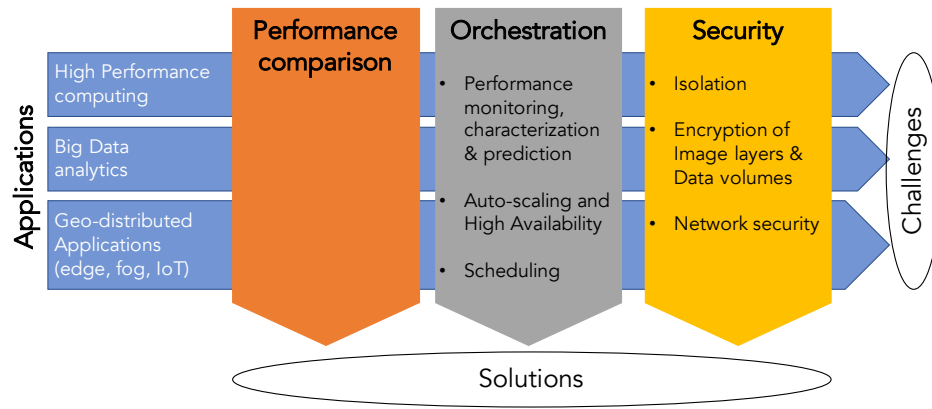


FIGURE 2 The proposed literature taxonomy

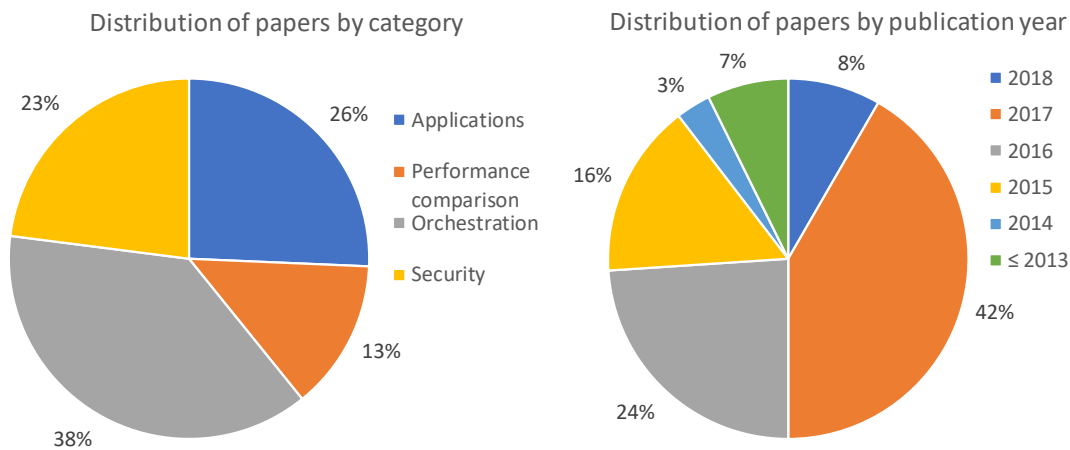


FIGURE 3 Distribution of the analyzed papers: by publication year (right) and by taxonomy categories (left)

arXiv. The following are the main keywords used in the search: docker, container, application virtualization, performance, monitoring, orchestration, self-adaptation, scheduling, energy efficiency, auto-scaling, security, encryption, availability, self-healing. The paper selection process took into consideration also the quantity of papers in each category. A reasonable balance between the applications category (about 1/4 of the papers) and the other categories (about 3/4 of the papers) has been maintained, as Figure 3 shows.

The distribution of the papers among the four main categories highlights that the majorities of research work is focused on orchestration (38%), while the 23% address security issues and 13% performance comparison. While it is understandable that there are few works on performance comparison, because the outcomes obtained until now converge to similar results (small footprint w.r.t. VMs, and limitation in network I/O performance), the number of papers addressing security issues is inadequate, considering also the importance of cyber security today. A more detailed analysis (c.f. Table 6) highlights that researchers' effort is mainly devoted to enhance isolation properties of containers, while there are few research works on image/volume encryption and network security.

Finally, it is also interesting to observe the novelty of research results: the 90% of the studies have been published between 2015 and 2018, as illustrated by the diagram in Figure 3; and the 50% of the papers analyzed have been published in the last two years, 2017-2018.

TABLE 3 Challenges in the HPC, Big Data analytics and geo-distributed (Fog, Edge, IoT) application fields

HPC	Big Data analytics	Fog, Edge, IoT applications
Security ^{8 29}	Security ^{30 31}	Security ³²
Deployment of parallel applications on CPU ^{33 34 29 35} and on GPU ^{34 35}	Scalable stream processing ^{36 37} and batch processing ^{38 5} platforms	Container migration ^{10 39}
HPC job Scheduling ^{40 29}	Scalable testing ^{5 37} and training ³¹ platforms	Resource and energy efficient service provisioning ^{41 39 42}
Optimization of performance ^{8 40 29 35}	Large data set management ^{30 43 44}	Deployment on etherogeneous nodes ^{45 32 42}

4 | APPLICATIONS

4.1 | High Performance Computing

Today, the users of super computing platforms are demanding for solutions to easily deploy their own defined software stack or virtual HPC clusters, and for leveraging the flexibility offered by containers in their real-life computational and data intensive jobs. To satisfy such demand, many challenges should be solved, for instance: support for MPI and CUDA libraries; support for execution of containerized parallel applications on GPU and FPGA; scheduling of containerized HPC jobs; mechanisms to guarantee, in the new container enabled HPC platforms, the same performance of native environments; solutions to guarantee the proper level of security (isolation of users' environment, data security, network security). Table 3 (first column) summarizes the challenges and the related work raising and/or proposing solutions to the challenges.

In⁸ the authors address the problem of deploying User Defined Software Stacks (UDSS) in supercomputing centers. The challenge is to provide a usable environment while minimizing the risks: security, support burden, missing functionality, and performance. The paper proposes Charliecloud, a framework based on Linux user and mount namespaces to run UDSS in Docker containers with no privileged operations or daemons. The adopted solution does not use network, host and domain name namespaces to gain in networking performance (critical for HPC). User namespace is used to limit the access to system calls, while security is managed using the Linux kernel capabilities (e.g. to enforce access control). Charliecloud has been validated against security attacks and performance. Socker²⁹ aims to leverage the flexibility offered by containers in HPC real-life computational and data intensive jobs. Socker is a solution to securely run Docker within cluster jobs and HPC job scripts (it is essentially a secure wrapper), and to limit the resource usage of a Docker job to the borders defined by the HPC queuing system. Socker is based on the Slurm queuing systems and does not require to modify the underlying Docker engine or to replace it. Fe2vCI2³⁵ is a framework for deploying HPC applications on common cloud infrastructures. Specifically, Fe2vCI2 allows the deployment of user defined virtual clusters for running HPC applications. The performance (FLOPS and network throughput) of the proposed solution has been assessed against a deployment based on VMs.

In³⁴ the authors propose a framework to efficiently deploy parallel C++ applications in a cloud environment by using Docker containers and Docker Swarm (for automatic node scaling). The proposed solution leverages OpenMPI (open-mpi.org) and CUDA (developer.nvidia.com/cuda-zone) container images to enable code execution on both CPU and GPU. In³³ the authors extend MPICH (www.mpich.org/) to support containers. The proposed solution allows to deploy MPI-based parallel applications with Dockers.

In⁴⁰ the authors propose a solution to deploy scientific applications and workflows on HPC cluster using Docker containers and the Adaptive Computing's Moab scheduler. The performance of the proposed solution has been evaluated in term of FLOPS, network throughput and filesystem throughput, and compared with the performance of a RedHat cluster. To increase container's network performance, a bridge network solution and the Linux's macvlan feature are used instead of Docker's NAT native mechanisms. The filesystem performance in case of NFS has been enhanced using a Docker's NFS plugin.

4.2 | Big Data Analytics Applications

Containers offer an easy way to deploy Big data analytics applications based on the map-reduce or stream processing paradigms. Moreover, as discussed in the previous section, solutions are emerging to support the execution of containers on GPU and FPGA, and to speedup specific analytics tasks. As summarized in Table 3 (second column), the deployment of Big Data analytics applications with containers poses new challenges: scalability and performance should be guaranteed for stream processing and batch processing applications, as well as for training and testing environments. For the latter, there is also the automation challenge. Finally, the large volume of data needs to be securely managed on-the-fly (ephemeral filesystems) and at rest (Docker volumes).

Data stream processing is a computing model that finds application in many fields, such as, big data analytics, security, IoT, performance monitoring, distributed event processing. In³⁶ the authors implement a Stream Processing Platform Service, deployed using Docker and Kubernetes, to test the data flow prediction algorithms they proposed. In³⁷ the authors propose a Docker-based architecture and a partitioning method to reduce the scalability problem of Apache Spark-based scale-up servers, caused by garbage collection and remote memory access overheads when the servers are equipped with significant number of cores and Non-Uniform Memory Access (NUMA). The paper shows that, using Docker container-based architecture, the problem could be minimized effectively by partitioning the original scale-up server into small logical servers.

In³⁸ the authors describe the development of Virtual Hadoop, a Docker-based framework that overcomes the scalability issues of HadoopCL, a variant of Hadoop that takes advantage of heterogeneous computing by offloading the computation kernels of a Hadoop application to the accelerators composed by multicore CPUs, GPUs, and/or even FPGA boards. HadoopCL does not support on-demand resource scaling. Docker is used to wrap-up the execution environment of Hadoop rather than the user application. Virtual Hadoop includes an extension of the methods of resource inference and allocation, and an auto-scaling mechanism to dynamically allocate resources on-demand when new workloads are submitted to the Hadoop cluster. The autoscaling is based on a performance model that predicts the needed resources for the new workload and estimates if the available resources (Hadoop virtual nodes) are enough to run the new workload, or new nodes are needed. In⁵ Docker is used to deploy an Hadoop virtual cluster on a single server node to run a distributed version of the Vp-Tree algorithm that leverages the MapReduce framework. The performance on the Docker cluster is compared with the performance of a traditional virtual machine based cluster (on a single server node).

In³¹ the authors present the Isolated, Scalable, and Lightweight Environment for Training (ISELT), targeting Massive Open Online Courses requiring interactive GNU/Linux commandline environment. ISELT has been deployed using docker containers and claims to overcome many of the distribution, scaling and security challenges for the specific training subject. While the authors provide a performance evaluation of ISELT, they do not provide any detail about how security and isolation are managed.

SciServer³⁰ is a big-data infrastructure project that is developing a modular and scalable infrastructure for the storage, access, query and processing of large, petabyte scale, scientific datasets. The SciServer Compute component leverages Docker containers to provide security and isolation, and to allow flexible configuration of computational contexts through domain specific images and mounting of domain specific data sets. Docker volumes are used to encapsulate dataset, which makes it easy to attach them to application containers as external storage, bringing data close to compute resources running analytic applications.

The storage of large datasets is also addressed in⁴⁴, who tackle the two main limitations of data deduplication based on Map-Reduce, namely: the inability to find duplicate data in adjacent blocks, and the possibility to miss duplicate data due to a fixed sliding window. Their approach has been validated using Docker containers to run 3 Hadoop instances on a physical node to execute the Map-Reduce jobs. Furthermore, given the highly distributed nature of container-based applications, and the need to process large datasets, the role of efficient caching becomes crucial to reduce the network utilization, which might become a bottleneck, and thus to reduce the processing and waiting time. Semantic cache allows the cache to be used even though the requested data is only partially present in the cache. Several solutions (e.g.,^{46,47}) have been proposed that implement semantic-caching; however, some challenges remain open, such as, the reduction of the query processing time and the cache hit ratio. To this end, in⁴³ the authors propose a novel semantic caching algorithm that outperforms regular semantic cache and decisional semantic cache⁴⁶.

4.3 | Edge, Fog and IoT Applications

Since containers are portable and have a small performance footprint, they are suitable to deploy applications on heterogeneous, resource limited and power limited computational units. Deploying containers on edge or fog computing nodes, or deploying containers on IoT boards like Raspberry Pi, is extremely advantageous, although it rises challenges in migration, service and resource provisioning and security, as outlined in Table 3 (Third column).

In¹⁰ the authors address the problem of mobile clients that need to offload services to an edge computing platform. The authors present a novel service handoff system based on container migration. In the paper, an important performance issue in Docker container migration has been identified and solved. The proposed solution leverages the layered storage system to reduce filesystem synchronization overhead, without dependence on the distributed filesystem.

The feasibility of using docker in Mobile Edge Computing scenarios, and specifically for video streaming and multiplayer gaming applications, is investigated in⁴⁸. The authors discover that, in case of video streaming, the overhead is independent from the number of clients and servers while, in case of multiplayer gaming, the overhead increases with the number of servers (running containers). A study about how to improve service provisioning flexibility in edge community cloud is proposed in⁴¹. The authors present the evolution of Guifi.net, integrating Docker and Cloudify to deploy software defined networks. Another example of how container could facilitate edge computing is provided in³⁹. In that paper, the authors describe how Docker is used: to deploy a learning environment on the Jetstream cloud infrastructure; and to enable portability of the students' learning environment from the cloud to their edge devices (e.g. laptop or tablet).

TABLE 4 A summary of the performance comparison studies

Performance Comparison
Docker vs KVM and Linux native ^{14 15}
Docker vs VMware and Linux native ⁴⁹
Docker vs LXC ¹⁶ and Joyent's Triton ⁵⁰
Docker configuration tuning: Memory ⁵¹ , Filesystems ^{52,53,54}
Docker vs native deployment on Raspberry Pi ⁵⁵
Energy consumption Docker/LXC vs KVM/Xen ⁵⁶

A container-based fog-oriented framework for Internet-of-Things applications based on Kura gateway and Docker is studied in⁴⁵. The paper demonstrates that Kura gateway can be containerized for an easy deployment on top of existing, off-the-shelf, and limited-cost gateway nodes. Specifically, it has been possible to implement a highly manageable and interoperable way to create fog nodes on-the-fly via the adoption of application containerization and container orchestration. In³² the authors address the heterogeneity problem in industrial IoT by proposing a Docker based architecture that allows the deployment of applications despite the specific device/protocol. Moreover, high availability and scalability, as well as fault tolerance, are guaranteed by the Docker engine and Docker Swarm.

Another interesting use of container is Cloudrone⁴²: a micro cloud deployed on Raspberry Pis mobile nodes transported by drones. The solution uses Docker to deploy services with a very small resource demand footprint and Swarm to manage clusters of Docker nodes.

5 | PERFORMANCE COMPARISON

One of the first goals of the research community has been to assess, under different workloads: how the performance of a deployment based on containers compare with the performance of virtual machines and/or bare-metal deployments; and how containers compares with VMs in term of energy consumption. The outcome of the analysis is the following: the performance footprint of containers is negligible; read and write performance heavily depends on the filesystem type selected and on its configuration; and network I/O is poor compared to VMs and native environments. Table 4 summarizes the comparison scenarios considered in literature.

The first seminal work on container performance evaluation¹⁴ provides an extensive comparison among a native Linux environment, Docker and KVM. In that work, the three environments are compared in presence of CPU intensive, I/O intensive, Network intensive, and NoSQL/SQL workloads. The main intention of the work is to assess the performance improvement of running workloads in containers rather than in VMs. The comparison is based on the performance metrics collected by the benchmarking tools, rather than on the workload footprint. In¹⁵ the authors present a similar study, aimed at comparing the performance of containers with hypervisors. The performance of Cassandra while running on bare-metal native cluster of nodes, VMware VMs and Docker containers is investigated in⁴⁹. The authors focus their investigation on the throughput and latency of Cassandra when different replication factors and consistency schemas are used, with the final goal of providing recommendations for Cassandra virtual cluster configuration and deployment.

Different studies investigate the performance of containers running on cloud infrastructures, like AWS EC2. A performance comparison of Docker versus Flockport (LXC) is presented in¹⁶, using the same benchmarks as in¹⁴. In this work, containers are deployed on top of the NeCTAR cloud. The comparison is intended to explore the performance of CPU, memory, network and disk. Docker versus Joyent's Triton is studied in⁵⁰. The authors show that on top of AWS EC2 virtual machines, Joyent's Triton performs better or almost the same than Docker, and that it is a valid alternative considering the added security features (cf. Sec. 7).

While standard benchmarks are used in the previous studies, the scientific workload is considered in⁵¹. The authors show that Docker memory configuration can be tuned to make container performance be slightly better than VMs.

The performance of containers running on Internet-of-Things devices are investigated in⁵⁵. Docker runs on a Single Board Computer device such as Raspberry Pi 21. Workload used are: system benchmarks to independently stress CPU, memory, network I/O, disk I/O; and application benchmarks reproducing MySQL and Apache workloads. The reference for comparison are the performance of the system without any virtualization.

How filesystem selection and configuration could impact the performance of Docker is investigated in^{52,53}. In⁵² the authors perform an extensive study to unravel the multi-faceted nature of Docker storage, for a wide range of filesystems, and to demonstrate its impact on system and workload

TABLE 5 A summary of the solutions proposed by the research community to address the container orchestration challenges: Performance Monitoring, Characterization and Prediction; Auto-scaling and High-availability; Scheduling.

Performance Monitoring, Characterization and Prediction	Auto-scaling and High-availability	Scheduling
Measurement methods and techniques ^{17 57 58}	Horizontal scaling ^{6 59 60 61 57 62}	Scheduling of containers on VMs ^{63 64 65} , on PMs ⁶⁶ , and in Geo-distributed environments ^{67 68 69}
Multi-layer Monitoring solutions ^{57 58 61}	Vertical scaling ^{70 59}	CPU and Memory-based optimization ^{64 66} , I/O-based optimization ⁶³
Characterization of resource contention ^{71 72}	Geo-distributed scaling ^{59 73}	Multi Objective optimization ⁶⁷
Enhancement of Docker monitoring: ^{72 57}	Application specific scaling: Time critical ⁶⁰ , Hadoop ³⁸ , Data stream processing ³⁶	Energy-based optimization ¹¹
Performance prediction ^{71 58 74 61}	High-availability solutions ^{75 76}	
Characterization of energy consumption ⁷⁷		

performance. An assessment of the performance of Docker when the Union filesystem and CoW are used to build image layers is provided in⁵³. The impact of high speed SSDs storage on Docker performance is analyzed in⁵⁴. The authors evaluate the performance of storing Docker images and data w.r.t. different filesystems type and storage drivers.

In⁵⁶ the authors compare the energy consumed by Docker and LXC with the energy consumed by KVM and Xen. In the comparison, the cases for memory, CPU and network intensive workload are considered, as well as for the idle state. Results show that container technologies consume less energy than hypervisors-based virtualization only for network intensive workload. In all the other cases, virtualization technologies behave the same.

6 | ORCHESTRATION

In the container life cycle, orchestration supports automation and run-time adaption (c.f. Figure 1). In this survey, we focus on the run-time adaptation challenges, that are: to design new methods and tools for performance monitoring, characterization and prediction; to design resource and energy efficient auto-scaling and scheduling algorithms; and to design new solutions for high-availability. The solutions proposed by the research community to address the above challenges are summarized in Table 5. An analysis of the literature is presented in what follows.

6.1 | Performance Monitoring, Characterization and Prediction

Performance monitoring is a topic of increasing interest for the containers' research community. In¹⁷ the authors assess the different measurement techniques used to collect performance counters for CPU and disk I/O intensive Docker workload. The study demonstrates that container performance counters introduce an under estimation ranging between the 5 and 10 percent w.r.t. system performance counters. Such a difference could hurt the behaviour of scheduling and autoscaling algorithms, as proved in⁶¹, where the impact of relative and absolute metrics on the Kubernetes autoscaling algorithm is investigated.

The importance of planning deployment and adaptation decisions on the basis of performance counter values collected at all the virtualization layers is addressed in⁵⁷ and ⁵⁸. In⁵⁷ the authors propose Elascalle, a cloud service that: automatically instruments the containerized application to collect performance counters at container and VMs layers; and then it uses the performance measure for autoscaling. The same approach is used in⁵⁸, where multi-layer monitored data are used to drive the run-time adaptation. In⁷² the authors propose an enhancement of the Docker's kernel to collect the maximum I/O bandwidth for the machine it is running on. That made available one more performance measure in Docker.

In⁷¹ the authors characterize the impact of Docker configuration and resource contention on the performance and scalability of big data workloads (four typical Apache Spark's applications). The resource contention among multiple I/O intensive applications sharing the same resources and running in Docker containers is characterized in⁷². A preliminary characterization of the energy consumed by Docker containers is provided in⁷⁷. The authors focus on correlating the energy consumed with the CPU and network pressure. The results show that CPU is the main source

of energy consumption. Intensive network load that stresses the OS results in a not negligible consumption of energy as well. Such results are important for designing energy efficient resource scheduling.

Performance prediction is key for capacity planning, optimal deployment and run-time adaptation. A performance prediction model based on the Support Vector Regression is proposed in⁷¹. The paper considers data stream processing Apache Spark applications, running in Docker containers, with different configurations and resource contention settings. In⁵⁸ the authors propose a Learning Classifier System that, on the basis of multi-layer monitored data, incrementally learns rules representing the behaviour of containerized applications. The extracted behaviour is used to predict the QoS level offered by the application. A Layered Queuing Network-based performance prediction model for multi-tier containerized applications is proposed in⁷⁴. The model allows to predict resource demand in development phase. In⁶¹ the author study the correlation between the container performance counters (counters from `/cgroup` filesystem) and the system performance counters (from `/proc` filesystem). Such correlation model is useful for workload prediction and for the design of scheduling and autoscaling strategies.

6.2 | Auto-Scaling and High-Availability

The adoption of container technologies call for new run-time adaptation solutions^{78,28}. Recently, many research works have proposed new solutions for horizontal and vertical scaling of Docker containers. The approaches used are generally inspired by the research results obtained in the last two decades in autonomic computing^{79,80,81,82}.

Among the auto-scaling solutions, the most promising mix vertical and horizontal scaling and/or migration. Usually the autoscaling algorithms were designed for locally-distributed adaptation, but solutions that consider the geo-distributed case, suitable for IoT applications, have been proposed as well. In high-availability, the most promising approach is to make the detection of faulty nodes more sophisticated, e.g., based on abnormal behavior rather than only on network level or application level probes.

An early study on container management⁸³ shows that Elastic Application Container-based resource management outperforms the VM-based approach.

In⁶ the authors provide an Integer Linear Programming (ILP) formulation of the elastic provisioning of virtual machines for container deployment. The general ILP problem formulation takes explicitly into consideration the heterogeneity of container requirements and virtual machine resources. Only QoS and cost are considered in the problem formulation. The same approach is used in⁵⁹ where the authors propose Adaptive Container Deployment (ACD), a general model of the deployment and adaptation of containerized applications. Besides acquiring and releasing geo-distributed computing resources, ACD can optimize multiple run-time deployment goals, by exploiting horizontal and vertical elasticity of containers. The ACD model is used as a benchmark to evaluate the behavior of several greedy heuristics for determining the container deployment.

An adaptive multi-instance container-based architecture targeting time-critical applications is proposed in⁶⁰. The solution has been implemented with Docker and Kubernetes.

The behaviour of the Kubernetes Horizontal Pods auto-scaling algorithm is studied in⁸⁴. The author proposes a new solution to make a more appropriate allocation of resources to fulfill application response time constraints. ElasticDocker, an autonomic controller powering vertical elasticity of Docker containers is presented in⁷⁰. ElasticDocker scales up and down both CPU and memory assigned to each container according to the application workload, and live-migrates containers when there is not enough resources on the hosting machine. The experiments show that ElasticDocker makes better resource utilization for container providers, and improves Quality of Experience for application end-users. An architecture of a SaaS autonomic application manager based on Docker and Kubernetes is proposed in⁷³. The high level architecture is based on three autonomic managers that should be capable to adapt the multi-cloud infrastructure and the multi technology data storage level with the goal of guaranteeing tenants' SLAs. While no details on adaptation algorithms are provided, the proposed architecture delineates a new research problem in the field. Elascle⁵⁷ scales micro services based applications deployed with Docker and Docker Swarm. The proposed solution applies a default threshold-based, reactive scaling algorithm for all the application's micro and macro services. In⁶² the authors propose a proactive autoscaling mechanism that scales in/out containers and distributes the load among instances on the base of the network traffic intensity.

Autoscaling algorithms targeting specific computing paradigms are discussed in³⁸ (Hadoop) and³⁶ (Data stream processing). In³⁸ the authors propose an auto-scaling algorithm for docker containers wrapping up the execution environments of Hadoop nodes. The autoscaling algorithm is integrated in a framework called Virtual Hadoop for deploying Hadoop clusters on heterogeneous nodes. The algorithm estimates the needed amount of containers to meet the time constraint of the Hadoop job and automatically allocates containers. In³⁶ the authors propose SSPS, an autoscaling mechanisms for Data Stream Processing applications deployed using Docker. The proposed autoscaling algorithm takes decisions based on the predicted data arrival rate. SSPS is designed to replace the default CPU threshold based scaling policy in the Horizontal Pod Autoscaler of Kubernetes.

Container orchestrators such as Docker Swarm and Kubernetes offer high-availability mechanisms based on monitoring the health state of a service and restarting the container in case the health check fails. Such basic mechanisms are enhanced in Serfnode⁷⁵, a lightweight platform

TABLE 6 A summary of the solution proposed by the research community to address the security challenges: isolation; volumes and images encryption, network security.

Isolation	Volumes and Images encryption	Network security
SGX-based isolation ^{85 86 87 88 89 90}	SGX-based volumes encryption ^{85 86 87 90}	SGX-based communication channel encryption ^{86 87 88 91}
Kernel-based isolation in Linux ^{8 29} and in Solaris ⁵⁰ (Joyent Triton)	SGX-based ephemeral filesystem encryption ^{86 87}	Host Identity Protocol based ⁹²
Linux hardening-techniques ^{93 94}	Selective image layers encryption at-rest ^{95 85} and on-the-fly ⁹⁵	Channel leakage mitigation techniques ⁹⁶
Application level isolation ⁹⁷		
Detection of escape attacks ⁹⁸		

agnostic and easy to integrate with an existing system of Docker containers. The proposed solution includes a monitoring and self-healing mechanism based on Supervisor (supervisord.org) for increased availability. In⁷⁶ the authors propose an intuitive approach based on Computational Intelligence (CI) for enhancing the availability of Docker Swarm. The proposed CI-based approach predicts the possible failure of the host of a manager node by observing its abnormal behaviour. Thus, this indication can automatically trigger the process of creating a new manager node or promoting an existing node as a manager for enhancing the orchestrator's availability.

6.3 | Scheduling

Scheduling is another brick of the run-time adaptation. Both static and dynamic scheduling solutions have been proposed.

In⁶³ the authors investigate the problem of scheduling network I/O intensive containerized applications on top of a cloud computing infrastructure (i.e. VMs). The proposed solution maximizes the number of allocated applications (i.e., containers) on a finite set of VMs with no outbreak of I/O bandwidth of physical nodes. At the same time, the optimal scheduling satisfies the application constraints on CPU, memory and disk demand. The proposed scheduler has been implemented in Diego, a Docker container orchestrator for Cloud Foundry.

In⁶⁴ the authors propose an Ant Colony Optimization (ACO) algorithm to schedule resource for Docker containers. The ACO algorithm has been implemented in Docker SwarmKit and the performance has been compared with the default greedy scheduling algorithm. A framework for Application Oriented Docker container (AODC) resource allocation is presented in⁶⁶. AODC minimizes the application deployment cost in datacenters and considers deployment of containers on PM. The deployment cost is related to available supporting libraries on PMs and required libraries of applications. AODC is compared against optimal VM placement algorithms.

C-Port⁶⁷ is the first example of orchestrator that makes it possible to deploy and manage containers across multiple clouds. The authors address the issues of container scheduling and placement, and dynamic adaptation. In term of orchestration policy, C-port uses a constraint-programming model for dynamic resource discovery, selection and scheduling. Decisions can be taken on the basis of availability, cost, performance, security, or power consumption constraints. The C-Port orchestrator is also used in⁶⁵ to realize a distributed software defined environment that deploys applications with docker over the CometCloud federated cloud infrastructure.

Scheduling containers on geo-distributed platform (e.g. federated clouds, edge, fog and IoT infrastructures) is challenging for a latency efficient download and deployment. In CoMICon⁶⁸ the authors propose a solution for reducing deployment time by co-operative management of Docker images. That enables their latency efficient scheduling algorithm to start container in a distributed or geo-distributed platform. In⁶⁹ the authors propose FID, a system for fast docker image distribution. FID uses a P2P BitTorrent based approach integrating a BitTorrent client in Docker registry and developing a FID agent that could work as a proxy for the Docker engine or as a REST API for developers (in that case the FID component uses Docker load).

In¹¹ the problem of properly scheduling resources to containers in a IoT scenario to reduce energy consumption is addressed. The authors propose DockerCap, a software-level power capping orchestrator for Docker containers that follows an Observe-Decide-Act loop structure: that allows to quickly react to changes that impact on the power consumption by managing resources of each container at run-time, to ensure the desired power cap. The paper shows that it is possible to obtain results comparable with the state-of-the-art power capping solution provided by Intel RAPL, while still being able to tune the performance of the containers and even guarantee SLA constraints.

7 | SECURITY

Docker security issues and challenges are discussed in⁹⁹. The authors give an overview of the vulnerabilities in terms of container isolation, host hardening and network security. Moreover, there are relevant security issues at the image distribution and container control levels. The Docker Content Trust Infrastructure relies on TUF¹⁰⁰ to enable image signature; Docker Hub integrates Diplomat¹⁰¹ to implement disambiguated trust delegations. However, other repositories could apply lower security mechanisms. Mechanisms and protocols to make trustable an image repository do not prevent the image to contain vulnerabilities that could be exploited. The authors of the Docker Image Vulnerability Analysis (DIVA) framework¹⁰² discovered that both official and community images contain more than 180 vulnerabilities on average when considering all versions. That challenge can be solved by means of automatic security updates. Frameworks for updates recommendation, such as RUDSEA¹⁰³, can be an approach to mitigate the problem.

A software engineering approach to enhance container security is considered in¹⁰⁴. The authors propose a reference architecture for the container ecosystem, claiming that it could facilitate who aims to ensure compliance, privacy, safety, reliability and/or governance. For example, security patterns can be used to build secure systems by describing ways to control specific threats, fix a vulnerability, or provide a security attribute.

Table 6 summarizes the solutions proposed by the research community to improve container isolation, image encryption and network security. A detailed analysis of the literature is presented in what follows.

7.1 | Isolation

The majority of work in container security addresses the isolation challenges. Isolation in accessing the resources is fairly guaranteed by the Linux namespace and *cgroup*, but there is a flaw in how containers share the same network bridge. Moreover, container isolation can be lowered at launch time playing with specific settings. In terms of host hardening, Linux operating system security policies offer a good protection for the host but they do not protect a container from other containers. The most promising solutions are based on Intel SGX enclaves, but the drawback is the dependence on the Intel architecture and the implementation complexity.

In⁹⁸ the authors focus on Docker escape attacks and exhaustively discuss the existing security mechanisms and security issues of Docker. The paper describes different methods to conduct Docker escape attacks and proposes a defense method based on status inspection of namespaces to detect anomalous processes. Covert channels could cause critical results like information leak between one container and another (or even the host). Covert channels attacks against Docker are extensively investigated in⁹³, where the authors analyze how to identify different types of covert channels and how current isolation mechanisms can be used and configured to prevent the attacks. Furthermore, the authors conclude that adopting full-fledged SELinux or AppArmor security policy is a key condition to protect the security perimeters of containers. Another study about the use of Linux hardening techniques to improve container security is presented in⁹⁴. Joyent Triton are containers running on Oracle/Solaris and offer the built-in zone, an isolated and secure environment for running containerized applications¹⁰⁵. Hence, Joyent Triton is a secure alternative to Docker when containers run on bare-metal. In⁵⁰ the authors compare the performance of Docker running on AWS EC2 with Joyent's Triton container showing that the latter performs better or almost the same, and that it is a valid alternative adding security to performance.

As mentioned before, Charliecloud⁸ is a secure framework based on Linux user and mount namespaces to run HPC applications on Docker containers. Charliecloud provides isolation using Linux kernel to enforce access control and other aspects of security. Moreover, Charliecloud prevents shenanigans such as creating device files or *setuid* binaries by means of doing operations on center-owned resources as the invoking unprivileged user. Charliecloud has been proved secure against hacks tentatives, such as, *chroot* escape, bypass of file and directory permissions, bind to privileged ports on all host IP addresses, *setuid* to an unmapped UID, and the like. As mentioned before, Socker²⁹ is a secure wrapper for running Docker containers on Slurm and similar HPC queuing systems. The execution of containers within Slurm jobs is done with the submitting user privileges instead of root privileges. Moreover, Socker enforces the inclusion of containers in the *cgroups* assigned by the queuing system to the parent jobs.

In⁸⁵ the authors present a generic concept to provide SGX security guarantees for operating system components, isolating kernel functionalities by means of Intel SGX enclaves. The solution can be useful to increase container security. SCONE⁸⁶ is a secure container mechanism for Docker that uses the Intel SGX to protect container processes from outside attacks. SCONE exposes a C standard library interface for transparently encrypting/decrypting application data on a per-file-descriptor basis. The proposed solution has a low overhead and is transparent to Docker because secure containers behave like regular containers. SCONE is also used by other solutions, such as, SERECA⁸⁷ and SecureCloud⁸⁸, to increase security of the stored data, and of the network communications. For example, SecureCloud aims to provide a secure remote computation platform running microservices in secure SCONE containers. Also in⁸⁹ the authors propose to use Intel SGX enclaves to isolate applications by means of encrypted memory. They assume to have applications which code need to be partitioned in single functionalities (step not needed if the application is designed using the microservice architecture). SecureStream⁹⁰ is a reactive middleware framework to deploy and process secure streams (e.g. Map-Reduce applications) at scale. SecureStream combines the high-level reactive dataflow programming paradigm with Intel SGX

enclaves in order to guarantee privacy and integrity of the processed data. All the containers run in enclaves and store encrypted data within enclaves, thus realizing high isolation.

DATS, a system to increase isolation at application level, is proposed in⁹⁷. The authors consider web applications that heavily access data in shared folders and that are deployed with containers. DATS introduces a templating language to compose data across data containers. Moreover, DATS introduces two primitives that act as robust declassifiers to enforce non-interference across containers, taking large applications out of the trusted computing base.

7.2 | Encryption of Image Layers and of Data Volume

Another issue with Docker is the lack of native support for: the encryption of image layers to protect sensitive data from malicious privileged users (e.g., registry administrator, cloud provider); and the encryption of the data volumes. Also for the data security challenge, the most encouraging approach is the use of the SGX support. For container images, the selective layer encryption is another possibility.

In⁹⁵ the authors propose and demonstrate a mechanism for secure Docker image manipulation throughout its life cycle: mechanisms for encryption/decryption at-rest and on-the-fly are used to protect confidentiality all the time, during creation, storage and usage of a Docker image. Distribution and migration of images has been enhanced with a mechanism that selectively encrypts only the layers of the filesystem that require strong confidentiality. SCONE⁸⁶ supports filesystem shields, preventing low-level attacks, such as the OS kernel controlling pointers and buffer sizes passed to the services and ensuring the confidentiality and integrity of the application data passed through the OS. With the filesystem shield, is possible to encrypt files or to authenticate the access to files. Moreover, SCONE provides a support for encryption and authentication of the ephemeral container's filesystem. SCONE is used in SERECA⁸⁷, which runs MongoDB and MySQL into SCONE secure containers. In that way the data is encrypted in memory and on the disk.

Another solution that leverages Intel SGX enclaves for disk encryption is presented in⁸⁵ and has proved to be robust to attacks to single-level and multi-level authentication. This solution could be used for encrypting data volumes and images. Unfortunately, it introduces a significant overhead w.r.t. traditional disk encryption.

7.3 | Network Security

Portability is one of the main advantages of containers, hence the mechanism for interconnecting containers should be independent from the cloud service provider (or platform). Moreover, in fog computing or federated cloud scenarios the application components could communicate inter-cloud. The inter-cloud connectivity highlights the need for NAT and firewall traversal support in the containers, and the need for secure end-to-end connectivity.

The network security challenges have been addressed in different ways, depending on the type of attack considered. The solutions independent from the specific attack and that seem to offer the stronger protection are based on SGX enclaves.

In⁹² the authors propose a solution, named SynAPTIC, to provide secure and persistent connectivity for containers. The proposed solution is based on the Host Identity Protocol (HIP), that enables strong authentication based on public key encryption. SynAPTIC integrates HIP in Docker containers. It does not require any changes in networking infrastructure and supports the use of Software Defined Networking. The Docker daemon is usually controlled through a Unix socket, but this can be changed to a TCP socket to grant attacker remote control to run any container in privileged mode. The only remedy is to enable by default TLS connections. SCONE⁸⁶, before mentioned, wraps all socket operations and redirects them to a network shield. The network shield, upon establishing a new connection, performs a TLS handshake and encrypts/decrypts any data transmitted through the socket. This approach is transparent to the client and server. The private key and certificate are read from the container's filesystem and can be protected by the filesystem shield (above mentioned).

SERECA⁸⁷ addresses the network security challenge by creating a secure event bus and encryption service leveraging SGX enclaves, and lets them communicate with the event bus. In SecureCloud⁸⁸, the secure communication among microservices (deployed into containers) is supported by Intel SGX enclaves. The proposed solution protects the exchanged messages by encrypting them inside of enclaves via an event bus. Encryption and decryption are automatically performed within enclaves. Decryption keys are stored securely within enclaves only. In⁹¹ SecureCloud is used to deploy a secure metering system.

The problem of mitigating network side channel leakage in stream processing applications is addressed in⁹⁶. The authors proposed a multicast and an any-cast technique to make the communication obvious and to mitigate leakage at inter-stage communication. That solution can be used to increase the security of any type of application, and specifically of container-based data-stream processing applications.

8 | DISCUSSION AND FINAL REMARKS

Since 2014, a considerable research effort has been devoted to improve container technologies. Many application areas found attractive the advantages offered by containers with respect to consolidated virtualization technologies, i.e. virtual machines. This survey reviewed the use of containers in High Performance Computing, Big Data analytics and geo-distributed (Fog, Edge, IoT) applications. These are different fields, but with common challenges: performance, scalability, high availability, isolation, data and network security. Moreover, there are also specific challenges, for example: in High Performance Computing there is the need to integrate parallel programming libraries with containers, such as Docker. Scheduling of containerized I/O intensive jobs is also a challenge. Furthermore, I/O performance limitation of containers should be mitigated or eliminated at all. Concerning Big Data applications, data stream and batch processing applications have performance as the key requirement and there is a need of run-time adaption algorithms to optimally deploy the required computing elements and to scale the ones with higher demand in case of surge of workload. Edge and fog computing applications are deployed on geo-distributed infrastructures, hence they demand for deployment strategies and run-time adaptation mechanisms that take into account latency and time to download images from repositories. In that case, there is also the problem of guaranteeing trusted geo-distributed image repositories. Internet of Things applications mainly demand for energy efficient deployment and run-time adaptation strategies.

The literature on *performance comparison* give an important contribution in demonstrating the small performance footprint of containers and the advantages with respect to virtual machines. Moreover, limitations in term of I/O overhead are found and remedies are proposed. Other performance aspects investigated in literature classify as *Performance Monitoring, Characterization and Prediction*. They focus on performance prediction models, multi-layer monitoring solutions, characterization of the performance contention for multi-tenant intensive I/O workloads, impact of different filesystems on the container performances. All this research work provide results and suggestions to fully exploit the advantages of containers.

Run-time adaptation challenges are investigated by research work on *auto-scaling, high-availability* and *scheduling*. The auto-scaling and high-availability solutions are mainly developed for big data and analytics applications running on cloud, edge and fog infrastructures (and more in general for fog and IoT applications). Among the auto-scaling solutions, the most promising mix vertical and horizontal scaling and/or migration. Usually the auto-scaling algorithms are designed for locally-distributed adaptation, but solutions that consider the geo-distributed case, suitable for IoT applications, are proposed as well. Different approaches aim to enhance the solutions implemented in tools like Docker Swarm and Kubernetes, like threshold-based algorithms, and algorithms based on mathematical programming or Ant Colony Optimization. The effort on high-availability focuses on enhancing the self-healing capabilities for container orchestrators like Docker Swarm and Kubernetes. The most promising approach is to make the detection of faulty nodes more sophisticated rather than based only on network level or application level probes.

The scheduling solutions are mainly intended for high performance computing applications and Internet-of-things applications. Scheduling challenges are addressed using different techniques to produce optimal schedules that take into account resource usage and energy consumption. Because of the poor network I/O performance of container, scheduling algorithms try to optimize the schedule of network I/O intensive workloads. Scheduling solutions try also to be energy efficient in the case of IoT applications. Solutions capping the energy at application level rather than at hardware level are promising.

Finally, *security* related research results largely focus on how to improve container isolation by means of hardening the configuration of existing mechanisms, or proposing new solutions based on Intel SGX. Resilience to different type of attacks (like covert channels and root escape) are also evaluated, and solutions to prevent and mitigate such attacks are proposed. Another security issue is to preserve the confidentiality of data storage in image layers. The proposed solutions are mainly techniques to encrypt/decrypt at rest and at run-time the image layers. Finally, the secure communication among container is evaluated and solutions or protocols to secure the communication are proposed. Results obtained leveraging the Intel-SGX enclaves are encouraging, also because they allow to address all the challenges: isolation, data encryption and network security. Hardening techniques are easier to implement, but weaker than SGX-based solutions.

To conclude, from the analysis of the literature presented in this paper, considering also how research works are distributed among topics (cf. Figure 3), it emerges that more effort should be devoted to solve I/O performance, multi-layer adaptation, energy efficiency and security challenges. Probably security is one of the most urgent needs, and security mechanisms should be integrated by design in the most popular container managers and orchestrators like Docker and Kubernetes. While security solutions should be application and platform agnostic, it makes sense that the energy efficiency, run-time adaptation and high availability solutions are tailored to specific deployments, e.g., geo-distributed versus locally-distributed, and workloads, e.g., batch versus stream processing.

ACKNOWLEDGMENT

The work of E.Casalicchio is funded by the following research grants: *Scalable resource-efficient systems for big data analytics* (grant: 20140032), Knowledge Foundation, Sweden; *PRIVacy-preserving, Security, and MACHine-learning techniques for healthcare applications (PRISMA)*, Ricerca d'Ateneo 2018, Sapienza University of Rome, Italy; *SmartDefense: Models, Algorithms and Mechanisms for Reducing Cyber Risks in Smart Industry* (RG11916B88C838E8 Ricerca Ateneo 2019) Sapienza University of Rome, Italy.

References

1. Dua R., Raja A. R., Kakadia D.. Virtualization vs Containerization to Support PaaS. In: Proc. of 2014 IEEE Int'l Conf. on Cloud Engineering; March, 2014.
2. Burns Brendan, Grant Brian, Oppenheimer David, Brewer Eric, Wilkes John. Borg, Omega, and Kubernetes. *ACM Queue*. 2016;14:70–93.
3. Vogels Werner. *Under the Hood of Amazon EC2 Container Service*. <http://www.allthingsdistributed.com/2015/07/under-the-hood-of-the-amazon-ec2-container-service.html>; 2015.
4. Natarajan S, Ghanwani A, Krishnaswamy D, Krishnan R, Willis P, Chaudhary A. *An Analysis of Container-based Platforms for NFV*. : IETF; April, 2016.
5. Nguyen DT-Tri, Yong Chan Ho, Pham Xuan-Qui, Nguyen Huu-Quoc, Loan Ton Thi Kim, Huh Eui-Nam. An Index Scheme for Similarity Search on Cloud Computing Using MapReduce over Docker Container. In: Proceedings of the 10th ACM International Conference on Ubiquitous Information Management and Communication; 2016; New York, NY, USA.
6. Nardelli Matteo, Hochreiner Christoph, Schulte Stefan. Elastic Provisioning of Virtual Machines for Container Deployment. In: Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering Companion; 2017; L'Aquila, Italy.
7. Gerlach Wolfgang, Tang Wei, Keegan Kevin, et al. Skyport: Container-based Execution Environment Management for Multi-cloud Scientific Workflows. In: Proceedings of the 5th IEEE International Workshop on Data-Intensive Computing in the Clouds; 2014; New Orleans, Louisiana.
8. Priedhorsky Reid, Randles Tim. Charliecloud: Unprivileged Containers for User-defined Software Stacks in HPC. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis; 2017; Denver, Colorado.
9. Ismail B. I., Goortani E. Mostajeran, Karim M. B. Ab, et al. Evaluation of Docker as Edge computing platform. In: 2015 IEEE Conference on Open Systems (ICOS):130-135; 2015.
10. Ma Lele, Yi Shanhe, Li Qun. Efficient Service Handoff Across Edge Servers via Docker Container Migration. In: Proceedings of the Second ACM/IEEE Symposium on Edge Computing; 2017; San Jose, California.
11. Asnaghi A., Ferroni M., Santambrogio M. D.. DockerCap: A Software-Level Power Capping Orchestrator for Docker Containers. In: 2016 IEEE Intl Conference on Computational Science and Engineering (CSE) and IEEE Intl Conference on Embedded and Ubiquitous Computing (EUC) and 15th Intl Symposium on Distributed Computing and Applications for Business Engineering (DCABES):90-97; 2016.
12. Celesti A., Fazio M., Giacobbe M., Puliafito A., Villari M.. Characterizing Cloud Federation in IoT. In: 2016 30th International Conference on Advanced Information Networking and Applications Workshops (WAINA):93-98; 2016.
13. McGee Jason. *The 6 steps of the container lifecycle*. <https://www.ibm.com/blogs/cloud-computing/2016/02/the-6-steps-of-the-container-lifecycle/>; 2016.
14. Felter W., Ferreira A., Rajamony R., Rubio J.. *An updated performance comparison of virtual machines and Linux containers*. : IBM Research Division, Austin Research Laboratory; July, 2014. RC25482(AUS1407-001).
15. Morabito R., Kjällman J., Komu M.. Hypervisors vs. Lightweight Virtualization: A Performance Comparison. In: 2015 IEEE International Conference on Cloud Engineering:386-393; 2015.

16. Kozhimbayev Zhanibek, Sinnott Richard O. A performance comparison of container-based technologies for the Cloud. *Future Generation Computer Systems*. 2017;68:175 - 182.
17. Casalicchio Emiliano, Perciballi Vanessa. Measuring Docker Performance: What a Mess!!!. In: Proc. of the 8th ACM/SPEC on International Conference on Performance Engineering Companion; 2017; L'Aquila, Italy.
18. OASIS . *Topology and Orchestration Specification for Cloud Applications*. Version 1.0: OASIS Standard; 2013.
19. Martino B. Di, Cretella G., Esposito A.. Advances in Applications Portability and Services Interoperability among Multiple Clouds. *IEEE Cloud Computing*. 2015;2(2):22-28.
20. E. Casalicchio. Container Orchestration: A Survey. Puliafito A., Trivedi K. (eds) *Systems Modeling: Methodologies and Tools*. EAI/Springer Innovations in Communication and Computing. Springer, Cham 2019.
21. Cerny Tomas, Donahoo Michael J., Trnka Michal. Contextual Understanding of Microservice Architecture: Current and Future Directions. *SIGAPP Appl. Comput. Rev.* 2018;17(4):29-45.
22. Pike Rob, Presotto Dave, Thompson Ken, Trickey Howard, Winterbottom Phil. The Use of Name Spaces in Plan 9. *SIGOPS Oper. Syst. Rev.* 1993;27(2):72-76.
23. Biederman Eric W.. Multiple Instances of the Global Linux Namespaces. In: 2006 Ottawa Linux Symposium; 2006.
24. Zadok E., Iyer R., Joukov N., Sivathanu G., Wright C. P.. On Incremental File System Development. *ACM Transactions on Storage (TOS)*. 2006;2(2):161-196.
25. Davies Alex, Orsaria Alessandro. Scale out with GlusterFS. *Linux Journal*. 2013;2013(235):1.
26. Maltzahn Carlos, Molina-Estolano Esteban, Khurana Amandeep, Nelson Alex J, Brandt Scott A, Weil Sage. Ceph as a scalable alternative to the hadoop distributed file system. *login: The USENIX Magazine*. 2010;35:38-49.
27. Laurén Samuel, Memarian M Reza, Conti Mauro, Leppänen Ville. Analysis of Security in Modern Container Platforms. In: Springer 2017 (pp. 351-369).
28. Casalicchio Emiliano. Autonomic Orchestration of Containers: Problem Definition and Research Challenges.. In: 10th EAI International Conference on Performance Evaluation Methodologies and Tools; 2016.
29. Azab A.. Enabling Docker Containers for High-Performance and Many-Task Computing. In: 2017 IEEE International Conference on Cloud Engineering (IC2E):279-285; 2017.
30. Medvedev Dmitry, Lemson Gerard, Rippin Mike. SciServer Compute: Bringing Analysis Close to the Data. In: Proceedings of the 28th International Conference on Scientific and Statistical Database Management; 2016; Budapest, Hungary.
31. Schipp Jonathan, Dopheide Jeannette, Slagell Adam. ISLET: An Isolated, Scalable, & Lightweight Environment for Training. In: Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure; 2015; St. Louis, Missouri.
32. Rufino J., Alam M., Ferreira J., Rehman A., Tsang K. F.. Orchestration of containerized microservices for IIoT using Docker. In: 2017 IEEE International Conference on Industrial Technology (ICIT):1532-1536; 2017.
33. Bayser M., Cerqueira R.. Integrating MPI with Docker for HPC. In: 2017 IEEE International Conference on Cloud Engineering (IC2E):259-265; 2017.
34. Wrede Fabian, Hof Vincent. Enabling Efficient Use of Algorithmic Skeletons in Cloud Environments: Container-based Virtualization for Hybrid CPU-GPU Execution of Data-parallel Skeletons. In: Proceedings of the ACM Symposium on Applied Computing; 2017; Marrakech, Morocco.
35. Ditter Alexander, Graf Gabriel, Fey Dietmar. Fe2vCI2: From Bare Metal to High Performance Computing on Virtual Clusters and Cloud Infrastructure. In: Proceedings of the 4th ACM Workshop on CrossCloud Infrastructures & Platforms; 2017; Belgrade, Serbia.
36. Wu Yongfeng, Rao Ruonan, Hong Pei, Ma Jin. FAS: A Flow Aware Scaling Mechanism for Stream Processing Platform Service Based on LMS. In: Proceedings of the 2017 ACM International Conference on Management Engineering, Software Engineering and Service Sciences; 2017; Wuhan, China.

37. Kyong Joohyun, Jeon Jinwoo, Lim Sung-Soo. Improving Scalability of Apache Spark-based Scale-up Server Through Docker Container-based Partitioning. In: Proceedings of the 6th ACM International Conference on Software and Computer Applications; 2017; Bangkok, Thailand.
38. Chen Yi-Wei, Hung Shih-Hao, Tu Chia-Heng, Yeh Chih Wei. Virtual Hadoop: MapReduce over Docker Containers with an Auto-Scaling Mechanism for Heterogeneous Environments. In: Proceedings of the International Conference on Research in Adaptive and Convergent Systems; 2016; Odense, Denmark.
39. Holdgraf Chris, Culich Aaron, Rokem Ariel, Deniz Fatma, Alegro Maryana, Ushizima Dani. Portable Learning Environments for Hands-On Computational Instruction: Using Container- and Cloud-Based Technology to Teach Data Science. In: Proceedings of the ACM Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact; 2017; New Orleans, LA, USA.
40. Julian Spencer, Shuey Michael, Cook Seth. Containers in Research: Initial Experiences with Lightweight Infrastructure. In: Proceedings of the ACM Conference on Diversity, Big Data, and Science at Scale (XSEDE16); 2016; Miami, USA.
41. Baig R., Freitag F., Navarro L.. Fostering Collaborative Edge Service Provision in Community Clouds with Docker. In: 2016 IEEE International Conference on Computer and Information Technology (CIT):560-567; 2016.
42. Sathiseelan Arjuna, Lertsinsruttavee Adisorn, Jagan Adarsh, Baskaran Prakash, Crowcroft Jon. Cloudrone: Micro Clouds in the Sky. In: Proceedings of the 2nd ACM Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use; 2016; Singapore, Singapore.
43. Ma Kun, Yang Bo, Yang Zhe, Yu Ziqiang. Segment access-aware dynamic semantic cache in cloud computing environment. *Journal of Parallel and Distributed Computing*. 2017;110:42–51.
44. Ma Kun, Dong Fusen, Yang Bo. Large-scale schema-free data deduplication approach with adaptive sliding window using mapreduce. *The Computer Journal*. 2015;58(11):3187–3201.
45. Bellavista Paolo, Zanni Alessandro. Feasibility of Fog Computing Deployment Based on Docker Containerization over RaspberryPi. In: Proceedings of the 18th ACM International Conference on Distributed Computing and Networking; 2017; Hyderabad, India.
46. Perrin Mikael, Mullen Jonathan, Helff Florian, Gruenwald Le, d'Orazio Laurent. Time-, Energy-, and Monetary Cost-Aware Cache Design for a Mobile-Cloud Database System. In: Springer 2015 (pp. 71–85).
47. Ren Qun, Dunham Margaret H, Kumar Vijay. Semantic caching and query processing. *IEEE transactions on knowledge and data engineering*. 2003;15(1):192–210.
48. Avino G., Malinverno M., Malandrino F., Casetti C., Chiasserini C. F.. Characterizing Docker Overhead in Mobile Edge Computing Scenarios. In: Proceedings of the Workshop on Hot Topics in Container Networking and Networked Systems; 2017; Los Angeles, CA, USA.
49. Shirinbab Sogand, Lundberg Lars, Casalicchio Emiliano. Performance Evaluation of Container and Virtual Machine Running Cassandra Workload. In: The 3rd IEEE International Conference on Cloud Computing Technologies and Applications (IEEE CloudTech17); 2017; Rabat, Morocco.
50. Balasubramanian Sekar Vishaka, Patil Vrushali, Giusti Michael, Bhide Aishwarya, Gupta Abhishek. AWS EC2 vs. Joyent's Triton: A Comparison of Docker Container-hosting Platforms. In: Proceedings of the 8th ACM Workshop on Scientific Cloud Computing; 2017; Washington, DC, USA.
51. Adufu Theodora, Choi Jieun, Kim Yoonhee. Is container-based technology a winner for high performance scientific applications?. In: 17th IEEE Asia-Pacific Network Operations and Management Symposium, APNOMS 2015, Busan, South Korea, August 19-21, 2015; 2015.
52. Tarasov V., Rupprecht L., Skourtis D., et al. In Search of the Ideal Storage Configuration for Docker Containers. In: 2017 IEEE 2nd International Workshops on Foundations and Applications of Self* Systems (FAS*W):199-206; 2017.
53. Dua R., Kohli V., Patil S., Patil S.. Performance analysis of Union and CoW File Systems with Docker. In: 2016 International Conference on Computing, Analytics and Security Trends (CAST):550-555; 2016.
54. Xu Q., Awasthi M., Malladi K. T., Bhimani J., Yang J., Annavam M.. Docker characterization on high performance SSDs. In: 2017 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS):133-134; 2017.

55. Morabito R.. A performance evaluation of container technologies on Internet of Things devices. In: 2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS):999-1000; 2016.
56. Morabito R.. Power Consumption of Virtualization Technologies: An Empirical Investigation. In: 2015 IEEE/ACM 8th International Conference on Utility and Cloud Computing (UCC):522-527; 2015.
57. Khazaei Hamzeh, Ravichandiran Rajsimman, Park Byungchul, Bannazadeh Hadi, Tizghadam Ali, Leon-Garcia Alberto. Elascala: Autoscaling and Monitoring As a Service. In: Proceedings of the 27th Annual International Conference on Computer Science and Software Engineering (CASCON17)IBM Corp.; 2017; Markham, Ontario, Canada.
58. Taherizadeh Salman, Stankovski Vlado. Incremental Learning from Multi-level Monitoring Data and Its Application to Component Based Software Engineering. In: 41st IEEE Annual Computer Software and Applications Conference (COMPSAC); 2017.
59. Nardelli M., Cardellini V., Casalicchio E.. Multi-Level Elastic Deployment of Containerized Applications in Geo-Distributed Environments. In: 2018 IEEE 6th International Conference on Future Internet of Things and Cloud (FiCloud); 2018.
60. Stankovski Vlado, Trnkoczy Jernej, Taherizadeh Salman, Cigale Matej. Implementing Time-critical Functionalities with a Distributed Adaptive Container Architecture. In: Proceedings of the 18th ACM International Conference on Information Integration and Web-based Applications and Services; 2016; Singapore, Singapore.
61. Casalicchio E., Perciballi V.. Auto-Scaling of Containers: The Impact of Relative and Absolute Metrics. In: 2017 IEEE 2nd International Workshops on Foundations and Applications of Self* Systems (FAS*W):207-214; 2017.
62. Kim Won-Yong, Lee Jin-Seop, Huh Eui-Nam. Study on Proactive Auto Scaling for Instance Through the Prediction of Network Traffic on the Container Environment. In: Proceedings of the 11th ACM International Conference on Ubiquitous Information Management and Communication; 2017; Beppu, Japan.
63. Zhao D., Mohamed M., Ludwig H.. Locality-aware Scheduling for Containers in Cloud Computing. *IEEE Transactions on Cloud Computing*. 2018;(99).
64. Kaewkasi C., Chuenmuneewong K.. Improvement of container scheduling for Docker using Ant Colony Optimization. In: 2017 9th International Conference on Knowledge and Smart Technology (KST):254-259; 2017.
65. Abdelbaky Moustafa, Diaz-Montes Javier, Parashar Manish. Towards Distributed Software-Defined Environments. In: Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing; 2017; Madrid, Spain.
66. Guan X., Wan X., Choi B. Y., Song S., Zhu J.. Application Oriented Dynamic Resource Allocation for Data Centers Using Docker Containers. *IEEE Communications Letters*. 2017;21(3):504-507.
67. Abdelbaky M., Diaz-Montes J., Parashar M., Unuvar M., Steinder M.. Docker Containers across Multiple Clouds and Data Centers. In: 2015 IEEE/ACM 8th International Conference on Utility and Cloud Computing (UCC):368-371; 2015.
68. Nathan S., Ghosh R., Mukherjee T., Narayanan K.. CoMICon: A Co-Operative Management System for Docker Container Images. In: 2017 IEEE International Conference on Cloud Engineering (IC2E):116-126; 2017.
69. Kangjin W., Yong Y., Ying L., Hanmei L., Lin M.. FID: A Faster Image Distribution System for Docker Platform. In: 2017 IEEE 2nd International Workshops on Foundations and Applications of Self* Systems (FAS*W):191-198; 2017.
70. Al-Dhuraibi Y., Paraiso F., Djarallah N., Merle P.. Autonomic Vertical Elasticity of Docker Containers with ELASTICDOCKER. In: 2017 IEEE 10th International Conference on Cloud Computing (CLOUD):472-479; 2017.
71. Ye K., Ji Y.. Performance Tuning and Modeling for Big Data Applications in Docker Containers. In: 2017 International Conference on Networking, Architecture, and Storage (NAS):1-6; 2017.
72. McDaniel S., Herbein S., Taufer M.. A Two-Tiered Approach to I/O Quality of Service in Docker Containers. In: 2015 IEEE International Conference on Cluster Computing:490-491; 2015.
73. Truyen Eddy, Van Landuyt Dimitri, Reniers Vincent, Rafique Ansar, Lagaisse Bert, Joosen Wouter. Towards a Container-based Architecture for Multi-tenant SaaS Applications. In: Proceedings of the 15th ACM International Workshop on Adaptive and Reflective Middleware; 2016; Trento, Italy.

74. Barna Cornel, Khazaei Hamzeh, Fokaefs Marios, Litoiu Marin. Delivering Elastic Containerized Cloud Applications to Enable DevOps. In: 12th IEEE/ACM International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS); 2017.
75. Stubbs J., Moreira W., Dooley R.. Distributed Systems of Microservices Using Docker and Serfnode. In: 2015 7th International Workshop on Science Gateways:34-39; 2015.
76. Naik N.. Applying Computational Intelligence for enhancing the dependability of multi-cloud systems using Docker Swarm. In: 2016 IEEE Symposium Series on Computational Intelligence (SSCI):1-7; 2016.
77. Tadesse S. S., Malandrino F., Chiasserini C. F.. Energy Consumption Measurements in Docker. In: IEEE 41st Annual Computer Software and Applications Conference (COMPSAC); 2017.
78. Pahl C.. Containerization and the PaaS Cloud. *IEEE Cloud Computing*. 2015;2(3):24-31.
79. Huebscher Markus C., McCann Julie A.. A Survey of Autonomic Computing - Degrees, Models, and Applications. *ACM Comput. Surv.* 2008;40(3):7:1-7:28.
80. Lemos Angel Lagares, Daniel Florian, Benatallah Boualem. Web Service Composition: A Survey of Techniques and Tools. *ACM Comput. Surv.* 2015;48(3):33:1-33:41.
81. Faniyi Funmilade, Bahsoon Rami. A Systematic Review of Service Level Management in the Cloud. *ACM Comput. Surv.* 2015;48(3):43:1-43:27.
82. Singh Sukhpal, Chana Inderveer. QoS-Aware Autonomic Resource Management in Cloud Computing: A Systematic Review. *ACM Comput. Surv.* 2015;48(3):42:1-42:46.
83. He S., Guo L., Guo Y., Wu C., Ghanem M., Han R.. Elastic Application Container: A Lightweight Approach for Cloud Resource Provisioning. In: 2012 IEEE 26th International Conference on Advanced Information Networking and Applications:15-22; 2012.
84. E. Casalicchio. A Study on Performance Measures for Auto-scaling CPU-intensive Containerized Applications. *Cluster Computing Springer*. 2019;.
85. Richter Lars, Götzfried Johannes, Müller Tilo. Isolating Operating System Components with Intel SGX. In: Proceedings of the 1st ACM Workshop on System Software for Trusted Execution; 2016; Trento, Italy.
86. Arnautov Sergei, Trach Bohdan, Gregor Franz, et al. SCONE: Secure Linux Containers with Intel SGX. In: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation; 2016; Savannah, GA, USA.
87. Fetzer Christof, Mazzeo Giovanni, Oliver John, Romano Luigi, Verburg Martijn. Integrating Reactive Cloud Applications in SERECA. In: Proceedings of the 12th ACM International Conference on Availability, Reliability and Security; 2017; Reggio Calabria, Italy.
88. Kelbert Florian, Gregor Franz, Pires Rafael, et al. SecureCloud: Secure Big Data Processing in Untrusted Clouds. In: Proceedings of the Conference on Design, Automation & Test in Europe, European Design and Automation Association; 2017; Lausanne, Switzerland.
89. Krieter Sebastian, Krüger Jacob, Weichbrodt Nico, Sartakov Vasily A., Kapitza Rüdiger, Leich Thomas. Towards Secure Dynamic Product Lines in the Cloud. In: Proceedings of the 40th ACM International Conference on Software Engineering: New Ideas and Emerging Results; 2018; Gothenburg, Sweden.
90. Havet Aurélien, Pires Rafael, Felber Pascal, Pasin Marcelo, Rouvoy Romain, Schiavoni Valerio. SecureStreams: A Reactive Middleware Framework for Secure Data Stream Processing. In: Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems (DEBS '17); 2017; Barcelona, Spain.
91. Riella Rodrigo J., Iantorno Luciana M., Junior Laerte C. R., et al. Securing Smart Metering Applications in Untrusted Clouds with the SecureCloud Platform. In: Proceedings of the 1st ACM Workshop on Privacy by Design in Distributed Systems; 2018; Porto, Portugal.
92. Ranjbar Alireza, Komu Miika, Salmela Patrik, Aura Tuomas. SynAPTIC: Secure And Persistent connectIvity for Containers. In: Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing; 2017; Madrid, Spain.
93. Luo Y., Luo W., Sun X., Shen Q., Ruan A., Wu Z.. Whispers between the Containers: High-Capacity Covert Channel Attacks in Docker. In: 2016 IEEE Trustcom/BigDataSE/ISPA:630-637; 2016.

94. MP Amith Raj, Kumar A., Pai S. J., Gopal A.. Enhancing security of Docker using Linux hardening techniques. In: 2016 2nd International Conference on Applied and Theoretical Computing and Communication Technology (iCATcCT):94-99; 2016.
95. Giannakopoulos I., Papazafeiropoulos K., Doka K., Koziris N.. Isolation in Docker through Layer Encryption. In: 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS); 2017.
96. Bilal Muhammad, Alsibyani Hassan, Canini Marco. Mitigating Network Side Channel Leakage for Stream Processing Systems in Trusted Execution Environments. In: Proceedings of the 12th ACM International Conference on Distributed and Event-based Systems; 2018; Hamilton, New Zealand.
97. Hunger Casen, Vilanova Lluís, Papamanthou Charalampos, Etsion Yoav, Tiwari Mohit. DATS - Data Containers for Web Applications. *ACM SIGPLAN Not.*. 2018;53(2):722-736.
98. Jian Zhiqiang, Chen Long. A Defense Method Against Docker Escape Attack. In: Proceedings of the 2017 International Conference on Cryptography, Security and Privacy; 2017; Wuhan, China.
99. Combe T., Martin A., Pietro R. Di. To Docker or Not to Docker: A Security Perspective. *IEEE Cloud Computing*. 2016;3(5):54-62.
100. Samuel Justin, Mathewson Nick, Cappos Justin, Dingleline Roger. Survivable Key Compromise in Software Update Systems. In: Proceedings of the 17th ACM Conference on Computer and Communications Security; 2010; Chicago, Illinois, USA.
101. Kuppusamy Trishank Karthik, Torres-Arias Santiago, Diaz Vladimir, Cappos Justin. Diplomat: Using Delegations to Protect Community Repositories. In: 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16):567-581; 2016; Santa Clara, CA.
102. Shu Rui, Gu Xiaohui, Enck William. A Study of Security Vulnerabilities on Docker Hub. In: Proceedings of the 7th ACM on Conference on Data and Application Security and Privacy; 2017; Scottsdale, Arizona, USA.
103. Hassan Foyzul, Rodriguez Rodney, Wang Xiaoyin. RUDSEA: Recommending Updates of Dockerfiles via Software Environment Analysis. In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering; 2018; Montpellier, France.
104. Syed Madiha H., Fernandez Eduardo B.. A Reference Architecture for the Container Ecosystem. In: Proceedings of the 13th ACM International Conference on Availability, Reliability and Security; 2018; Hamburg, Germany.
105. Oracle . System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones. Oracle Corporation 500 Oracle Parkway Redwood City, CA 94065 U.S.A.

How to cite this article: E. Casalicchio, S. Iannucci (2018), The State-of-the-Art in Container Technologies: Application, Orchestration and Security, XXX, XXX.