



ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Impact of Softwarization in Microservices-based SDN Controller

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Impact of Softwarization in Microservices-based SDN Controller / Domenico Scotece; Sisay Tadesse Arzo; Riccardo Bassoli; Luca Foschini; Michael Devetsikiotis; Frank H. P. Fitzek. - ELETTRONICO. - (2022), pp. 1-7. (Intervento presentato al convegno European Wireless 2022; 27th European Wireless Conference tenutosi a Dresden, Germany nel 19-21 September 2022).

This version is available at: <https://hdl.handle.net/11585/936576> since: 2023-07-26

Published:

DOI: <http://doi.org/>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

(Article begins on next page)

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

This is the final peer-reviewed accepted manuscript of:

D. Scotece, S. T. Arzo, R. Bassoli, L. Foschini, M. Devetsikiotis and F. H. P. Fitzek, "Impact of Softwarization in Microservices-based SDN Controller," *European Wireless 2022; 27th European Wireless Conference*, Dresden, Germany, 2022, pp. 1-6.

The final published version is available online at:
<https://ieeexplore.ieee.org/abstract/document/10071932>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Impact of Softwarization in Microservices-based SDN Controller

Domenico Scotece[†], Sisay Tadesse Arzo^{*}, Riccardo Bassoli^{◇‡}, Luca Foschini[†]

Michael Devetsikiotis^{*}, Frank H.P. Fitzek^{¶‡}

[†]Department of Computer Science and Engineering, University of Bologna, Bologna, Italy

^{*}Department of Electrical and Computer Engineering, University of New Mexico, USA

[¶]Deutsche Telekom Chair of Communication Networks, Technische Universität Dresden, Dresden, Germany

[◇]Quantum Communication Networks (QCNets) Research Group, Technische Universität Dresden, Dresden, Germany

[‡]Centre for Tactile Internet with Human-in-the-Loop (CeTI), Cluster of Excellence, Dresden, Germany.

Email:[†]{domenico.scotece, luca.foschini}@unibo.it

^{*}{sarzo, mdevets}@unm.edu,

^{◇‡}{riccardo.bassoli, frank.fitzek}@tu-dresden.de

Abstract—Software-defined networking decouples control and data plane in *softwarized* networks. This allows for centralized management of the network. However, complete centralization of the controller's functions raises the issue of the central point of failure, latency, and scalability. Distributed controller deployment is adopted to optimize scalability and latency problems. However, the existing controllers are monolithic, resulting in code inefficiency for distributed deployment. Recently, microservices-based SDN solutions have been started and deployed as virtual network functions, enabling flexible deployment. Nonetheless, the *softwarization* of network functionalities introduces network I/O performance degradation both considering deployments based on virtual machines and containers. This paper first introduces a microservices-based SDN solution based on Ryu SDN Framework. Then investigates whether running Ryu's network functionalities in a softwarized environment (e.g., virtual machine and container) would have a significant network I/O performance degradation. In particular, this work examines more in deep the Docker Container technology and analyzes its network setups. Multiple measurements were performed locally in a single machine. Our results show a comparison between microservices-based non-virtualized SDN Controller and the virtualized one, and moreover, show a comparison between different Docker Container network setups.

Index Terms—Software-Defined Networking, Microservices, Docker Container, Virtual Machine, 5G

I. INTRODUCTION

Software-defined networking (SDN) architecture has three layers [1]: application, control and forwarding layer. The *application layer* contains software applications to provide network services. Such layer performs ranges of functionalities such as security and routing. The *control layer* is the central agent, which interfaces the application and forwarding layer, to implement the applications' network requirements. It communicates through northbound interface to the applications and via southbound interface to the forwarding devices. The *forwarding plane*, is responsible for handling and forwarding packets. It contains a group of data plane resources that can forward and manipulate packets. Forwarding devices have physical/logical interfaces to receive the incoming packets

and forward them to outgoing interface(s). The controller communicates with forwarding devices using OpenFlow [2] or other protocols.

The existing controllers are implemented as a monolithic entity, which creates multiple problems in terms of scalability, fault isolation, and latency. The legacy definition of the SDN architecture does not stipulate the internal composition, implementation, and design of an SDN controller. Thus, the SDN controller can be decomposed and implemented as a set of software components, running in a distributed manner. Specifically, it is possible to design the SDN controller as a composition of logical sub-functions, sharing the network service load and creating robust system against failures. These sub-functions are loosely-coupled units that can be executed in a different and distributed computing platforms [3].

The benefits provided by the decomposition of SDN Controller into microservices are manifold including: i) *scalability*, that allow microservices-based network functionalities to scale either horizontally and vertically; ii) *reliability*, the fault of a single microservices-based network functionality does not impact on the entire system; iii) *reusability*, microservices-based functionalities can be reused to avoid unnecessary duplication of code and effort. In our previous work, we presented an implementation based on Ryu SDN Framework following the aforementioned principles [4]. In that paper, we introduced MSN (Microservices-based SDN Controller) framework to design and evaluate a microservices-based SDN Controller. In particular, we provided an extensive evaluation of MSN in terms of latency, reliability, and scalability. Moreover, a first release is available to the community at the link: https://gitlab.com/dscotece/ryu_sdn_decomposition/.

Accordingly, in this paper, we consider two different softwarization technologies for implementing the microservices network functionalities proposed in the MSN. In particular, we examine virtual machines and container technologies with a more focus on the Docker Container technology. In this context, due to the widespread deployment nature of edge

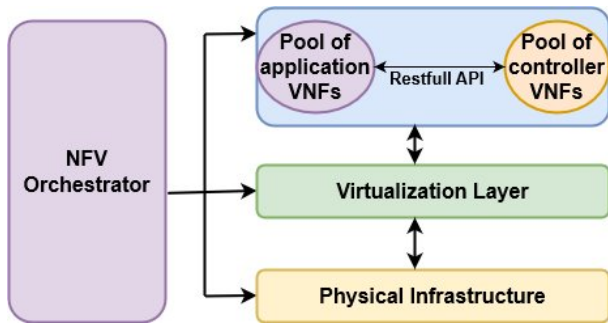


Fig. 1. MSN Decomposition Architecture

computing units, microservices and Docker container-based virtualization become an indispensable approach [5]–[7]. Microservices are a means of creating loosely-coupled sub-functions or sub-services, replacing a large software system. In addition, container-based virtualization is one of the key enabler to the success of 5G edge-enabled networks, which is supported by network infrastructure suppliers that are keen to bring containerization into edge computing [8]. However, networking performance is important to the user experience in containerized applications, it is important to investigate containers networking. Generally, container networks can be divided into two categories: single-host networks and multi-host networks. Since communications within the same host are done in a shared memory fashion, a single-host network is mainly to provide a networking interface for containerized applications. Multi-host network centers on providing IP addressing services, such as network address translation (NAT), overlay network, or routing, to interconnect containers on different hosts. The paper investigates whether running MSN in a softwarized environment would have a significant network I/O performance degradation. Specifically, we considered three different deployment scenarios of MSN including no virtualization, virtual machines deployment, and containers deployment.

In a nutshell, the flow of the paper is as follows. First, we provide some background materials for MSN implementation and related works (Section II). Then, we present our softwarized-based implementation of MSN functionalities both inside virtual machines and Docker containers (Section III). Furthermore, we benchmark the different softwarization implementations of MSN in terms of network latency (Section IV). Finally, we draw conclusions (Section V).

II. BACKGROUND AND RELATED WORK

In this section, we briefly introduce our MSN framework and we provide some details of the SDN Controller decomposition architecture. Finally, we briefly introduce the research directions of the literature in the areas of decomposing the SDN Controller into microservices.

A. Background

The main principle we adopted in the MSN decomposition implementation is that the network information and state

should be synchronized and self-consistent providing a global view of the network. This allows an independent implementation and components reuse. MSN considers a decomposition of SDN Controller, as depicted in Figure 1. The figure shows a decomposed three-layer SDN architecture reflected in an NFV architecture [9]. The control layer is decomposed into subfunctions which are implemented as softwarized network functions. The subfunctions could be orchestrated by standard orchestrators, such as ETSI MANO or Kubernetes, creating a service function chain to equivalently perform the legacy SDN controller’s functions. In the figure, the upper layer is a pool of independently implemented SDN components. Each component performs a specific function such as topology management and routing. These functions could be categorized as basic SDN controller functions and additional functions or applications. Basic SDN controller’s functions are mandatory functions, which are required to emulate the minimum possible function of SDN controller. Additional functions could be considered as applications such as firewall, and QoS monitoring. The basic functions include:

- **Event handling function:** this function receives events and distributes them to the appropriate function for further processing. This could be a link-failure notification from the forwarding layer, which needs to be communicated to both Network State-Management Function and Topology Management Function for updates. Moreover, the event handler should communicate with the management and orchestration unit for function orchestration.
- **Network Resource Management Function:** this function is an inventory system for the available network resources such as links, ports, switches, and routers.
- **Topology-Management Function:** this function is a set of valid associations between network resources and objects such as ports, devices, and links.
- **Data-Plan Control Functions:** these could be multiple functions implemented in multiple containers depending on the size of the network. The data plan control functions include: forwarding of packets, flow rule installation to the forwarding device, and resource reservation.
- **Network State-Management Function:** this database manages the global network information such as link and device status. A centralized database system with a backup per-controller domain is required to have global information.

In addition to the above basic functions, applications could be incorporated to extend the controller’s basic functionality. Depending on the network to be controlled, various types of applications could be implemented in the application layer such as firewall, deep packet inspection, and traffic predictions.

B. Related Work

The first work showing an externalization of packet processing in SDN is presented in [10]. As an extension of this work the author in [11] provided steps that are required to migrate from a monolithic to a microservice-based architecture. The functional components are externalized and a gRPC is

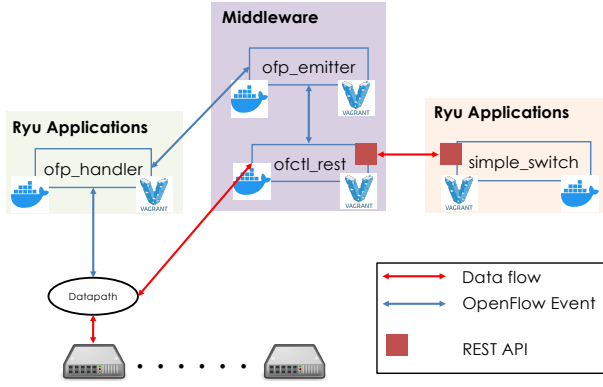


Fig. 2. Ryu-based MSN Implementation Architecture

used to communicate between the core modules and external components or applications.

μ ONOS or (micro Open Network Operating System) is the latest solution proposed towards standard architecture for distributed and split control plane [12]. The μ ONOS project aims at creating new generation of SDN architecture based on ONOS, splitting it into a set of microservices. These splitted functionalities are deployed as Docker containers and managed by Kubernetes orchestrator. Moreover, the μ ONOS project is relatively young (it started in October 2019). In particular, it is based on a new generation control protocol such as P4/P4runtime [13] that guarantees more flexibility compared to OpenFlow protocol, can also be used to emulate the behavior of OpenFlow. Furthermore, the communication between functionalities is via gRPC-based protocols, including gNMI for network management interface configuration and gNOI for network command operations.

However, even if the μ ONOS work is leading the research on decomposition of SDN controller, there is still no available implementation to test its performance. Moreover, the communication is based on gRPC, and not REST API, which has resulted in the following major limitations:

- the μ ONOS implementation has limited isolation mechanism; which means the core functions and applications share the same resources or process;
- the μ ONOS cannot have on-platform tenant-specific applications but only tenant-aware ones: tenant-specific apps must be off-platform and it should use REST APIs;
- the on-platform applications are limited to Java based languages: applications developed using other languages have to be off-platform and need to use REST APIs;
- horizontal service scaling is difficult;
- it has limited integration with and support for NFV that do not adhere to either an openflow abstraction of that of a legacy network element.

In general, to the best of the authors' knowledge, there are very few works, implementing a complete SDN controller's decomposition as VNFs, in a *containerised* environment.

III. SOFTWAREZED MSN IMPLEMENTATION

The following section provides the proposed microservices-based implementation of the SDN Controller. We used the Ryu controller framework [14] as the baseline SDN Controller implementation due to its component-based architecture that blends well with the microservices-based SDN controller perspective. Then, we describe our proposed implementation model and we show two distinct deployment models specifically based on Virtual Machine and on Docker Container technologies.

A. Decomposing Ryu SDN Controller

To start splitting an SDN Controller into microservices is important to identify the core part that allows network communications between microservices both for control and data plane functionalities. This also grants complete access to microservices from external processes to provide new functionalities. The MSN prototype is built around the objective to demonstrate the use of a middleware that allows interaction between external processes with the underlying infrastructure. Figure 2 shows the basic implementation architecture of the MSN prototype. The blue block in the center represents the middleware that is composed of two different microservices such as *ofp_emitter* and *ofctl_rest*. The green block is the internal Ryu application named *ofp_handler* that acts as the event handler for the OpenFlow requests. External Ryu applications (yellow block) can communicate with the Ryu framework via REST APIs through the middleware. In particular, when the *ofp_handler* microservices propagates an OpenFlow request to the middleware (through the *ofp_emitter* microservice), the same request goes to external applications via REST APIs (thanks to the *ofctl_rest* microservice). In this way, we transform SDN functionalities into microservices releasing them from the whole SDN Framework.

B. Docker-based MSN implementation

In order to implement MSN in a Dockerized environment, we produced two different Dockerfile one for the middleware and the other for the external Ryu applications. We chose to put the *ofp_handler* inside the same container of the middleware always as a separated service. This is because the Ryu framework uses OpenFlow protocol to communicate with the *ofp_emitter* module. Instead, external Ryu applications are in separated containers and communicate with the middleware via REST APIs.

On the one hand, the middleware Dockerfile starts from *python:3* base images [15]. In particular, it contains the two internal applications such as *ofp_emitter* and *ofctl_rest*. Moreover, the *ofp_handler* is installed by the *pip install* command and allows the middleware to catch OpenFlow events. Finally, the command *ryu-manager* starts the Ryu environment with the two applications as arguments. See the example of the Dockerfile in the following code.

```

1 FROM python:3
2 WORKDIR /usr/src/app
3
4 COPY ofp_emitter.py .
5 COPY ofctl_rest.py .
6
7 RUN pip install ryu
8
9 EXPOSE 8080
10 EXPOSE 6633
11
12 CMD ["ryu-manager", "ofp_emitter.py", "ofctl_rest.py"]

```

On the other hand, the external Ryu application (in our case the *simple_switch* application) uses the same *python:3* base image. This application is in charge of analyzing incoming packets and creating the flow rules for the underlying switches. Moreover, it creates a Flask web server [16] in order to communicate with REST APIs. In the following the example of the Dockerfile.

```

1 FROM python:3
2 WORKDIR /usr/src/app
3
4 COPY requirements.txt ./
5 RUN pip install --no-cache-dir -r requirements.txt
6
7 COPY lib ./lib
8 COPY simple_switch_rest.py .
9
10 EXPOSE 8090
11
12 CMD ["python", "./simple_switch_rest.py"]

```

C. Vagrant-based MSN implementation

Vagrant [17] is open-source software that allows the creation of virtual machines independently from specific hypervisors. For this reason, Vagrant provides Vagrantfile which is the same concept as Dockerfile. Moreover, we followed the same implementation guidelines that we explained in the previous subsection.

Different from Docker, the Vagrantfile of the MSN middleware starts from *ubuntu/bionic64* base image [18]. The middleware VM provides the two internal services of Ryu and installs the Ryu environment with *pip install* command. See an example of Vagrantfile for the MSN middleware in the following code.

```

1 Vagrant.configure("2") do |config|
2   config.vm.box = "ubuntu/bionic64"
3   config.vm.network :forwarded_port, guest: 8080, host:
4     ↪ 8080
5   config.vm.network :forwarded_port, guest: 6633, host:
6     ↪ 6633
7
8   config.vm.provision "shell", inline: <<-SHELL
9     apt-get update
10    apt-get install -y python3-pip
11    pip3 install ryu
12  SHELL
13
14  ryu-manager /vagrant/ofp_emitter.py
15  ↪ /vagrant/ofctl_rest.py
16 end

```

Similarly, we implement our *simple_switch* application in a Vagrant box. Also here, we provided a Flask server to use

REST for the communication purpose. In the following code, we reported an example of MSN *simple_switch* Vagrantfile.

```

1 Vagrant.configure("2") do |config|
2   config.vm.box = "ubuntu/bionic64"
3   config.vm.network :forwarded_port, guest: 8090, host:
4     ↪ 8090
5
6   config.vm.provision "shell", inline: <<-SHELL
7     apt-get update
8     apt-get install -y python3-pip
9     pip3 install --no-cache-dir -r
10    ↪ /vagrant/requirements.txt
11  SHELL
12
13  python3 /vagrant/simple_switch_rest.py
14 end

```

IV. TESTBED AND PERFORMANCE EVALUATION

This section provides performance evaluation of the proposed MSN solution in softwarized environment discussed in Section III. We first discuss the experimental setup and next we discuss the set of comprehensive results.

A. Experimental Environment

Our testbed consists of a Linux workstation (Ubuntu 20.04) equipped with an AMD Opteron(tm) Processor 6376 processor and 32 GB 1600MHz DDR3 memory. The machine hosts the docker community edition version 20.10.12, Vagrant version 2.2.19, Ryu SDN Framework, Python version 3, and Mininet (for creating virtual networks). We evaluated the feasibility and performance of the MSN solution for three different deployment scenarios. Figure 3 shows the three different deployment scenarios that we have tested in the experimental evaluation. This evaluation aims at measuring the network latency introduced by the softwarization in the MSN solution.

In order to create a test network topology, we leveraged Mininet software and its python extensions. The topology we chose to test MSN is composed of 5 switches (S_i) and 10 hosts (H_n) as described in Fig. 3. Moreover, to easily test our MSN solution we leveraged the *ping* tool that sends periodic *ICMP request* and replies with *ICMP reply*. This allows us to recognize the total end-to-end time. Furthermore, we are only interested in the first packet, because subsequent packets do not pass through the SDN Controller. Finally, all results obtained in the testbed are an average of 30 runs that exhibited a limited variance of under 5%.

B. Experimental results

With the experiments, we analyze the network impact of the softwarization on our MSN solution. Specifically, we measured the latency time for the first packet, normal flow, and updates in three different deployment scenarios as described in Fig. 3. Furthermore, we repeated tests for different hosts in the network topology specifically for H1-H3 and H1-H10.

Figure 4 shows the overall results obtained in the three different deployment scenarios for three different categories of packets including the first packet, normal flow, and update packet. The MSN deployment scenario (Fig. 3a) is used as the baseline for comparison in this experimental evaluation.

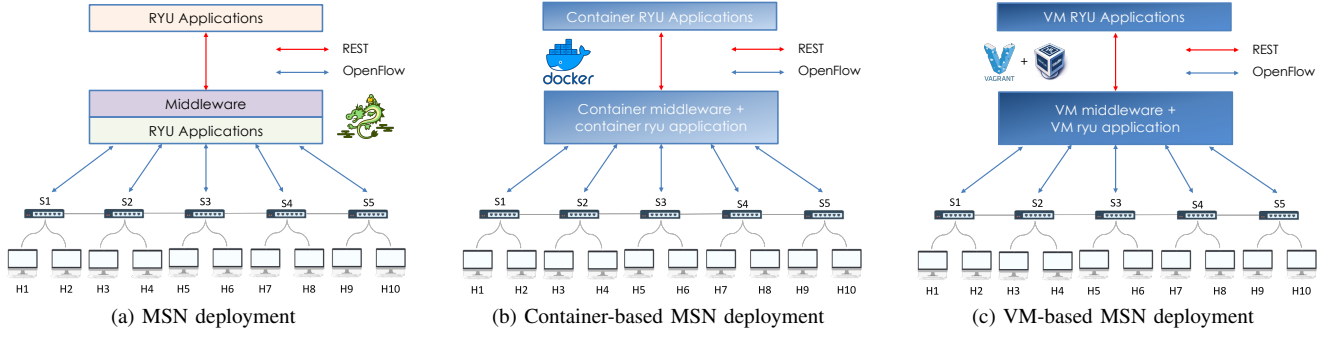


Fig. 3. Different MSN implementation deployment scenarios

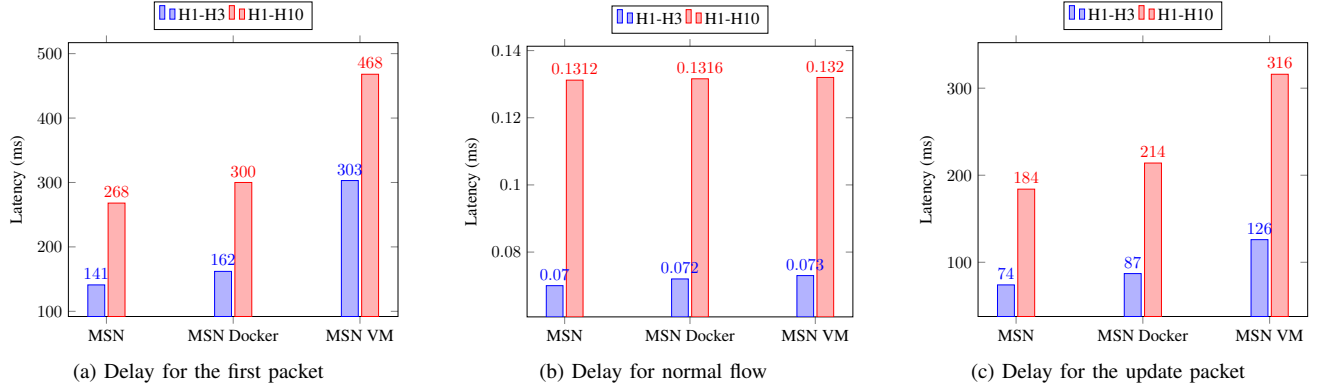


Fig. 4. Experimental Results

The first round of experiments is about the delay introduced by the first packet (Fig. 4a) both in the case of H1-H3 and H1-H10 communications. Note that the first packet goes to the SDN Controller before reaching the involved nodes. The MSN Docker deployment introduces a bit of delay compared to the standard MSN deployment about close to 20 ms for H1-H3 communications and 30 ms for H1-H10 communications. On the contrary, the MSN virtual machine deployment increases the delay by around twice compared to the standard MSN deployment for both H1-H3 and H1-H10 communications.

Then the second round of experiments is about the normal flow. Note that, once the flow rule is installed on the switches incoming packets do not go through the SDN Controller anymore. This is not completely true, because flow rules may have an expiry time, which is the focus of the next experiments. As shown in Fig. 4b, the delay observed in the normal flow is the same for all MSN deployment scenarios. This is because, incoming packets do not require SDN Controller interactions. Furthermore, this is true independently from the network topology (both for H1-H3 and for H1-H10).

Last, we benchmark the delay introduced by the flow rule updates. As mentioned before, flow rules may have an expiry time. After that time, the SDN Controller needs to update the validity of flow rules. This process takes time similar to the first packet delay. Figure 4c shows the trend of the delay for the update packet for all MSN deployment scenarios. The behavior is almost the same that we observed in the first packet analysis.

C. Docker networking test

Since the focus of this part of experiments is to examine the impact of dockerized version of MSN on network I/O performance, it is crucial to take a look at the network setup of containers in Docker [19]. Let us introducing two basic networking mode for Docker containers in a single host. For the purpose of this research, we taking into account only host and bridge networking mode. This is justified, as bridge is the default mode of Docker and host is the simple way that Docker uses to share the network interfaces with the host machine.

- **Bridge:** this is the default networking setup when a Docker container is created in a single host. In this way, Docker creates a virtual bridge named *docker0* in the host when the *dockerd* is started. Once a new container is started, a pair of vethernet interfaces are created in order to connect the container to the *docker0* interface. All containers connecting to the *docker0* bridge can connect each other using a private IP addresses. By default, bridge mode does not connect containers to external network, therefore, allows containers to create an isolated network namespace and IP addresses and all inter-container communications need to go through the *docker0* bridge.
- **Host:** this networking allows all containers on the same host to share the network namespace of the host. In this way, all containers are visible to each other and inter-container communications are based on inter-process communication (IPC). The host mode provides the lowest

level of security among the described networking modes as all users share the same IP address as well as the namespace of the host machine.

In these experiments, we run our MSN solution in the scenario shown in Fig 3b by varying the network configuration of Docker containers from *bridge* to *host*. Note that we used the same network topology as the previous experiments and ping tool as well. Results from Docker bridge network configuration are shown in Table I (middle column). Successively, we repeat the same experiments by switching the Docker network configuration from bridge to host. Table I (right column) shows the results obtained for Docker host network configuration.

TABLE I
MSN DOCKER DEPLOYMENT RESULTS

H1 ping H3		
	Docker Bridge	Docker Host
First packet	162 ms	157 ms
Normal Flow	0.0765 ms	0.0722 ms
Update packet	87 ms	85 ms
H1 ping H5		
First packet	226 ms	215 ms
Normal Flow	0.083 ms	0.0841 ms
Update packet	136 ms	129 ms
H1 ping H7		
First packet	265 ms	244 ms
Normal Flow	0.107 ms	0.0923 ms
Update packet	174 ms	174 ms
H1 ping H9		
First packet	300 ms	292 ms
Normal Flow	0.116 ms	0.106 ms
Update packet	214 ms	214 ms

Obtained results show that latency times are a bit lower in the case of Docker host network configuration. This is because, as mentioned before, in the case of Docker bridge network configuration there is a virtual bridge named *docker0* between containers and the host that introduces latency. This difference is constant and remain unchanged in the case of different network topology as demonstrated in the results. Finally, Docker network latency affects both the first packet and the update packet; normal flow is not affected by the Docker network configuration.

V. CONCLUSION

The paper investigated the network impact of softwarization in our microservices-based SDN controller (MSN) solution. First, we reviewed our MSN solution and we motivated our guidelines for the SDN decomposition. In particular, we proposed two different softwarized implementations of the MSN including a virtual machine-based implementation and a container-based implementation. The results show that the container-based implementation can reduce the network latency by around 50% compared to the virtual machine-based implementation. Furthermore, the evaluation also shows that the Docker host network configuration outperforms a bit the Docker bridge mode in terms of latency.

Future research directions include new mathematical models to manage efficiently the network impact of the softwarization.

On the other hand, the evaluation of our MSN solution in wide-scale scenarios by carrying out experiments in a real 5G campus networks is also important.

ACKNOWLEDGEMENTS

This work has been partially funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy – EXC2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden. The authors also acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of “Souverän. Digital. Vernetzt.”. Joint project 6G-life, project identification number: 16KISK001K.

REFERENCES

- [1] O. N. Foundation, “Software-Defined Networking: The New Norm for Networks,” <http://opennetworking.wpengine.com/wp-content/uploads/2011/09/wp-sdn-newnorm.pdf>, 2022, [Online; accessed 30-April-2022].
- [2] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, “Openflow: Enabling innovation in campus networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, p. 69–74, Mar. 2008.
- [3] I. Afolabi, T. Taleb, K. Samdanis, A. Ksentini, and H. Flinck, “Network slicing and softwarization: A survey on principles, enabling technologies, and solutions,” *IEEE Commu. Surveys Tutorials*, vol. 20, no. 3, pp. 2429–2453, thirdquarter 2018.
- [4] Arzo, Sisay Tadesse, Scotece, Domenico, Bassoli, Riccardo, Barattini, Daniel, Granelli, Fabrizio, Foschini, Luca, and Fitzek, Frank H. P., “Msn: A playground framework for design and evaluation of microservices-based sdn controller,” *Journal of Network and Systems Management*, vol. 30, no. 1, pp. 1573–7705, 2021.
- [5] Microsoft. Azure iot edge. [Online]. Available: <https://docs.microsoft.com/en-us/azure/iot-edge/?view=iotedge-2020-11>
- [6] Amazon. Aws iot greengrass. [Online]. Available: <https://aws.amazon.com/it/greengrass/>
- [7] IBM. Ibm edge application manager. [Online]. Available: <https://developer.ibm.com/components/ibm-edge-application-manager/>
- [8] Ericsson. Edge computing and deployment strategies for communication service providers. [Online]. Available: <https://www.ericsson.com/en/reports-and-papers/white-papers/edge-computing-and-deployment-strategies-for-communication-service-providers>
- [9] “NFV Technology,” <https://www.etsi.org/technologies/nfv>, [Online; accessed 30-April-2022].
- [10] D. Comer and A. Rastegarnia, “Externalization of packet processing in software defined networking,” *IEEE Networking Letters*, vol. 1, no. 3, pp. 124–127, Sep. 2019.
- [11] D. Comer and A. Rastegarnia, “Towards disaggregating the SDN control plane,” *CoRR*, vol. abs/1902.00581, 2019. [Online]. Available: <http://arxiv.org/abs/1902.00581>
- [12] ONOS, “micro ONOS,” <https://docs.onosproject.org/>, 2022, [Online; accessed 30-April-2022].
- [13] P4, “P4 Consortium,” <https://p4.org/>, 2022, [Online; accessed 30-April-2022].
- [14] Ryu, “Ryu SDN Framework,” <https://ryu-sdn.org/>, 2022, [Online; accessed 30-April-2022].
- [15] Python, “Python - Docker official image,” https://hub.docker.com/_/python, 2022, [Online; accessed 30-April-2022].
- [16] Flask, “Flask - Python web server,” <https://flask.palletsprojects.com/en/2.1.x/>, 2022, [Online; accessed 30-April-2022].
- [17] Vagrant, “Vagrant by HashiCorp,” <https://www.vagrantup.com/>, 2022, [Online; accessed 30-April-2022].
- [18] Ubuntu, “ubuntu/bionic64 Vagrant box,” <https://app.vagrantup.com/ubuntu/boxes/bionic64>, 2022, [Online; accessed 30-April-2022].
- [19] Docker, “Networking overview,” <https://docs.docker.com/network/>, 2022, [Online; accessed 30-April-2022].