Performance evaluation of the OSM orchestrator

Robert Kołakowski*◇, Lechosław Tomaszewski*, Sławomir Kukliński*◇
*Orange Polska, ◇Warsaw University of Technology
Warsaw, Poland

***Abstract*—The transformation of telecommunication networks towards fully softwarized solutions is ongoing. It requires efficient orchestration and management of in-software made networking solutions, including network slices. A widely accepted orchestration framework is already standardized by ETSI NFV, however, it is not yet validated by a commercial deployment and raises scalability concerns due to its highly centralized nature. The paper presents the results of experiments carried for performance evaluation of the Open Source MANO (OSM), which is the most popular ETSI-compliant orchestration solution. The performance assessment is focused on the lifecycle operations on multiple instances of Linux Foundation MAGMA Evolved Packet Core (vEPC) template that are treated as network slices. The experiments have shown excellent scalability of the OSM orchestrator.**

***Index Terms*—network slicing, orchestration, NFV, OSM, performance evaluation, management, KPI, scalability**

I. Introduction

The 5G network had brought a significant disruption to the telecommunication ecosystem. The concept uses network virtualization concept with the ability to create several, parallel and isolated networks satisfying specific service demands, i.e. network slices [1]. It is expected that the management and orchestration system of 5G will have to handle a vast number of slices with diverse lifecycle and runtime requirements, thus making effective management a crucial issue. Apart from the effective resource allocation to the network slices and satisfying the users' demands, the performance and scalability of the orchestration, which is responsible for the deployment, modification and termination of network slices, is of prime importance to slice tenants and network operators. A simple measure of the lifecycle operations performance can be slice deployment and termination time. Moreover, the behaviour of the system can be different under the light and heavy load of the orchestration platform. Slice deployment time can be critical in short-lived slices and the ones that are created to handle emergency situations. Long slice termination time or slow resource allocation to the slices' components can lead to inefficient usage of resources and may degrade the quality of services of solutions that are deployed as a slice.

This paper is an experimental study focused on the performance evaluation of lifecycle related procedures of a MANO compliant orchestrator. The performance assessments have been made using the ETSI-compliant Open Source Management and Orchestration (OSM) framework and relatively complex Network Slice Template (NST) that in the testing case is Magma Evolved Packet Core (a virtual EPC, vEPC) developed by the Linux Foundation [2]. To assess the scalability of orchestration, the slice lifecycle KPIs proposed in [3] have been used.

This work has been supported by the EU H2020 projects: 5G-DRIVE (Grant Agreement No. 814956) and MonB5G (Grant Agreement No. 871780).

II. Problem statement and related work

The Performance Management (PM) process is one of the fundamental activities of Network Operators. It serves both for Quality of Service (QoS) assurance and operators' environment expansion trends evaluation for proper anticipation of future demands to be answered in accordance with QoS targets. The hierarchy of the performance concerns reflects the hierarchy of service and network architecture as well as the underlying technologies. The numerous (thousands or more) Performance Indicators (PIs), composed either by direct or further reprocessed (recalculated, normalized, referred to the time of observation, etc.) readouts of specific counters, provide a quantitative view of the individual operation of equipment, functions, sub-systems and systems. To provide the practical and representative observability of the overall, end-to-end behavior of the systems, the PIs are aggregated to higher-level abstraction performance indicators, forming a short list of Key Performance Indicators (KPIs). They further contribute to the end-to-end communication service-level quality representation by Key Quality Indicators (KQIs).

Performance management is in the area of interest of various Standards Developing Organizations (SDO). A review of their approaches is provided in [3]. It has to be noted that the KPIs are used for the deployed solutions, but the management and orchestration related KPIs are so far rarely used. To that end, 3GPP focuses on performance measurements of specific network functions (PIs) [4] and 5G end-to-end KPIs [5]. Management and performance evaluation of orchestration has been initially ignored, but recently, as a part of Release 3 of ETSI NFV and several specifications related to orchestration management have been already published. In [6] the detailed specification of the interfaces designed for management of NFV-MANO framework has been provided, including interfaces for fault, configuration, performance and security management. Moreover, several performance metrics have been proposed (some being inline with the metrics already defined in [7]) that include physical resource usage such as CPU, RAM or storage, statistics regarding the objects managed by MANO or metrics related to lifecycle workflows (successful completions, failures, rollbacks etc.). These metrics can be exposed to external entities for further processing via Performance Management Interface defined in details in [6]. The orchestration of slices can also be controlled by the Priority parameter of the Network Service Deployment Flavor

[8]. The 5G System (5GS) uses ETSI Network Functions Virtualization (NFV) Management and Orchestration (MANO) framework [9], however, the 3GPP has extended the OSS/BSS part of the ETSI NFV architecture by the components specific for orchestration and runtime management of 5G network slices. Furthermore, some functions related to slice selection and authentication have also been added to the control plane of 5GC.

The performance of network slices' orchestration processes is crucial for slice provisioning in a timeframe accepted by verticals (slice tenants) and for efficient utilization of resources during slice lifetime, which is also a key feature in the context of energy savings. Long slice termination or deployment time contribute to inefficient usage of resources and can cause excessive energy loss. It has to be noted that at the moment of writing the paper, there is none commercially deployed 5G network supporting network slicing. Therefore, the performance of the orchestration of network slicing has not been yet validated in large scale deployments.

The orchestration KPIs are so far ignored by 3GPP. In [3] the set of network slicing-related KPIs that are agnostic to orchestrated slices has been proposed. They are divided into two categories: real-time KPIs, dedicated to indicate over- and underutilization of slice resources, and slice lifecycle-oriented ones to evaluate the end-to-end operational agility of the network slicing platform and communications networks implemented on it. The approach has been used to evaluate the OSM [10] orchestrator, with the generic CirrosOS VNFs [11] [12]. The scalability of the slice instantiation process has also been evaluated in the paper. It has to be noted that the mentioned KPIs are not absolute measures as the execution time of slice lifecycle operations depends on many factors, especially the number of VNFs of slice template, their footprint, the number of parameters ad procedures to be triggered in slice initialization phase or on the configuration and the performance of the infrastructure.

Some metrics for evaluation of the orchestration solutions as well as benchmarks for the OSM have also been described in [13] [14]. The two types of KPIs have been proposed functional (i.a. resource footprint, the maximum number of supported objects) and operational KPIs (i.a. on-boarding process delay, deployment process delay or a metric for a total quantification of MANO LCM performance called Quality of Decision). Moreover, the performance evaluation of OSM Release 4 has been done and compared with the Open Network Automation Platform (ONAP) orchestration solution.

III. Methodology

A. *Definition of the measured KPIs*

The list of the measured KPIs proposed in [3] includes:

- Slice Deployment Time (SDT) – a parameter that describes the interval between slice deployment request and the moment in which the slice is ready for operation. This interval depends on the performance of the orchestrator, the time related to the allocation of resources by the virtualized infrastructure. Both times are dependent on the slice template complexity, which may deal with the footprint size of VNFs, their inter-connection topology, and the number of configuration parameters. Therefore, it is impossible to define the required value of SDT in a generic case.
- Slice Deployment Time Scalability (SDTS) – a measure of scalability of slice deployment operations. To evaluate the scalability, it is possible to send N slice deployment requests of the same slice template and calculate SDTS as presented in the equation (1):

$$SDTS = \frac{GSDT}{N \cdot SDT} \tag{1}$$

  where *GSDT* is the overall time needed for the deployment of *N* identical slices and *SDT* is the deployment time of a single slice (as defined above).
- Slice Termination Time (STT) is a parameter that describes the interval between the slice termination request and the moment in which all slice allocated resources are released. If the time is long, it decreases the efficiency of the infrastructure resources usage (the resources are no more consumed).
- Slice Termination Time Scalability (STTS) is a measure of the scalability of slice termination operations. The calculation of STTS, analogically to SDTS, is defined by the equation (2):

$$STTS = \frac{GSTT}{N \cdot STT} \tag{2}$$

  where *GSTT* is the time needed for termination of *N* identical slices together.

Both scalability KPIs, when lower than 1, show scalability gain on slice deployment/termination and resources allocation/deallocation, respectively.

B. *Measurements of KPIs*

To calculate the KPIs defined above, it is proposed to monitor OSS/BSS-NFVO interactions of the NFV/MANO architecture [15]. The lifecycle KPIs can be obtained through the monitoring of the exchange of information between the OSS/BSS and NFVO: the OSS/BSS is able to determine both the beginning (request) and the end (completion report) of the requested procedure. There are two possible ways of calculation: (i) based on events logging in the on-board OSS/BSS log – each event is logged with a timestamp, and correlated search of beginning/finishing event for a specific procedure is sufficient; (ii) the OSS/BSS-MANO communication API will typically use the time-out mechanism and the time-out timer can be implemented – its value at the end of the procedure may be instantly passed to the PM engine of OSS/BSS.

For the tests, the first approach has been chosen. Moreover, the load of the infrastructure (CPU, RAM) during slices orchestration operations has been measured.

IV. Testbed description

A. Network slice templates used for the deployment

During the preliminary evaluation of the lifecycle KPIs [12], simple VNFs (i.e. different lightweight Linux distributions) were used. Hence, the obtained results might not be in line with the ones for more demanding tasks, i.e. orchestration of complex slices, composed of multiple VNFs, and requiring a high amount of resources.

As recently considerable effort had been put in providing compatibility of OSM and Magma [2] solution developed by the Linux Foundation project, we have decided to further investigate the OSM orchestration capabilities [16] by conducting measurements of the lifecycle KPIs for deployment and termination of the Magma vEPC network slices, i.e. by the orchestration of multiple Magma instances. Due to the size and complexity of the orchestrated Magma VNFs, the measurement results enable a relatively good assessment of the real-life OSM performance.

Magma is an open source software platform enabling the creation of the cloudified 3GPP core networks (2G–5G) for different access networks. The solution is composed of three main components:

- Access Gateway (AGW) – providing network services and policy enforcement (in case of LTE network, the AGW implements a single vEPC);
- Orchestrator (Orc8r) – cloud service to provide a way to configure and monitor the mobile network;
- Federation Gateway (FG) – a proxy to integrate AGWs and specific MNO core network entities that contain subscribers' identities or charging policies (e.g. Home Subscriber Server, Policy and Charging Rules Function, etc.).

The interconnected AGW and orchestrator form the single Magma vEPC network slice as illustrated in Fig. 1. The FG

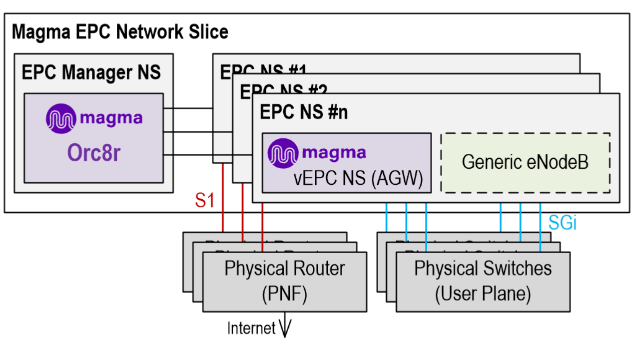

Fig. 1. High-level view of Magma EPC Network Slice (with multiple vEPC Network Slices)

is an optional component outside the Magma vEPC network slice template. Moreover, the number of FGs would also be largely dependent on the spatial distribution of MNO's core network entities (one gateway can proxy for multiple instances of Magma vEPC slices). Therefore, the measurements have been performed only for vEPC network slice (AGW).

B. Environment

The measurements of network slicing lifecycle KPIs have been conducted in the environment depicted in Fig. 2.

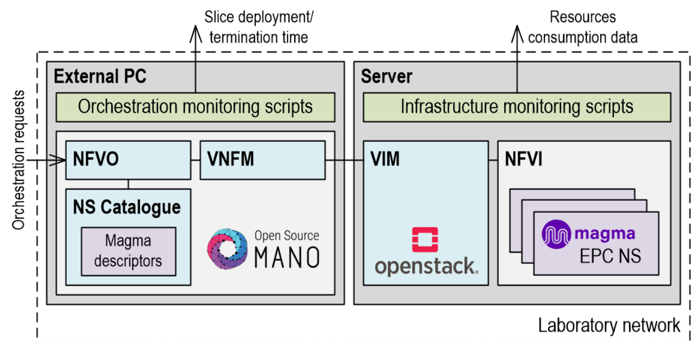

Fig. 2. The configuration of the measurement environment.

The test environment was composed of two machines: a PC hosting the OSM Release 9.1 and the server acting as VIM (OpenStack) as well as NFVI for the deployment of vEPC Network Slices. The parameters of the used hardware are presented in Table I.

TABLE I
SPECIFICATION OF THE HARDWARE USED DURING MEASUREMENTS

Hardware component	CPU	Memory	Disk
External PC	Intel® Core™ i7-8700 CPU @ 3.2 GHz	16 GB	256 GB
Server	Intel Xeon Silver 4216 2.1 GHz, 16C/32T, 9.6 GT/s	32 GB RDIMM 3200 MT/s	1 TB

The monitoring scripts were used on both machines during the orchestration operations. The scripts enabled the collection and extraction of the timestamps from OSM logs (PC case) and data related to resources consumption, i.e. RAM and CPU utilization (server case). The orchestration requests were generated by scripts using OSM Command Line Interface (CLI) commands. Prior to testing the deployment and termination of vEPC Network Slices, one Magma Orc8r instance has been instantiated in NFVI (not shown in Fig. 2). To deploy the vEPC Network Slice Instances (NSIs), the following packages have been used:

- NS package: *magma-agw_nsd*
- VNF package: *magma-agw_vnfd*
- Image: *magma101.qcow2*

Each vEPC NSI required 2 GB of RAM, 50 GB of storage and 2 CPU cores to operate. Due to the hardware resource limitations, the largest possible number of concurrently running slice instances was 18. Therefore, the conducted KPI measurements were performed by deployment and termination from 1 to 18 identical instances using the above-mentioned template (translated first into the required format [17]). To assess the scalability KPIs, the operations have been done both in a simultaneous (bundled) way, as well as a one-by-one approach (deploying each additional instance once the previous deployment is finished).

To minimize the impact of random and undesired fluctuations of parameters during measurements, each test has been conducted ten times, and the obtained results have been averaged. Prior to performing the measurement tasks, the RAM and CPU usage in the idle state of the system has been measured and averaged. The obtained result has been later on subtracted from the measurement results obtained during the orchestration tasks to separate the impact of orchestration processes on the resource consumption.

V. Results

A. Network slice deployment time measurements

The slice deployment process has been measured by deploying 1, 2, 3, 5, 7, 10, 14 and 18 AGW instances simultaneously (a bundle) and one-by-one, using the same template. Based on the measured SDT and GSDT values, the respective SDTS values for N given above have been calculated. The obtained results are presented in Fig. 3 and Fig. 4. The results of the AGW deployment measurements (cf. Fig. 3) show that the NSIs deployment bundling approach gives a clear advantage in comparison to the subsequent NSI deployment approach. The SDTS curve shows that the maximum gain on deployment bundling is achieved at a relatively small bundle size ($N = 5$), but for other N the SDTS values are still smaller than one.

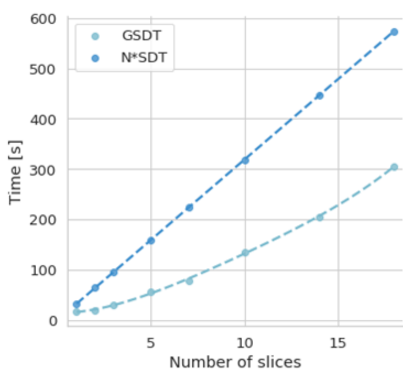

Fig. 3. Measured GSDT and N×SDT (for N = 1, 2, 3, 5, 7, 10, 14, 18).

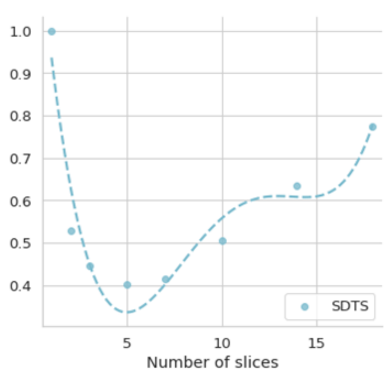

Fig. 4. Calculated SDTS (for N = 1, 2, 3, 5, 7, 10, 14, 18).

The averaged RAM and CPU consumption for a different number of slice instances during the slice deployment phase are presented in Fig. 5 and Fig. 6. These curves show trends that explain the behavior of the SDTS parameter and also show the load characteristics of the orchestration process. The deployment execution is preceded by a delay period, approximately independent of the bundle size, during which the resource orchestration processes prepare the environment (cf. Fig. 5). Then, the NSIs deployment process begins, which is manifested by a monotonic increase in memory usage without sudden and big changes. However, in the same period of the deployment execution preparation, the peak of the CPU load, which is associated with increased VIM computation activity hosted on the same server (cf. Fig. 6) can be observed. Later, the load drops to the level resulting from the computational demand of the NSIs being run.

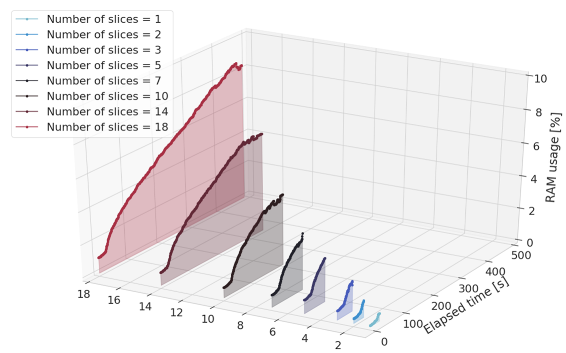

Fig. 5. Comparison of RAM usage measurement during the bundled deployment (for N = 1, 2, 3, 5, 7, 10, 14, 18).

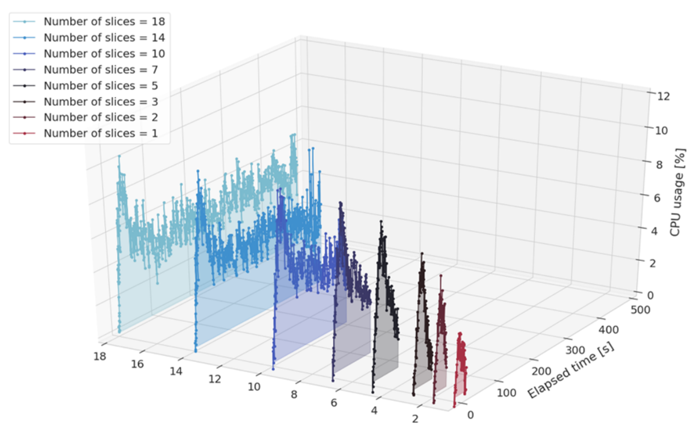

Fig. 6. Comparison of CPU usage measurement during the bundled deployment (for N = 1, 2, 3, 5, 7, 10, 14, 18).

B. Network slice termination time measurements

The slice termination measurements have been conducted for termination of 1, 2, 3, 5, 7, 10, 14 and 18 AGW instances in a simultaneous and one-by-one manner. It should be noted that during the series of measurements for $N = 18$ there was one termination process (out of 10), during which, GSTT was much longer than in other cases and its influence is clear on all measurement curves.

Based on the measured STT and GSTT values, the respective STTS values for N given above have been calculated. The obtained results are presented in Fig. 7 and Fig. 8. Similarly to the deployment processes, the NSIs termination bundling approach gives a time gain compared to the subsequent deployment approach. The gain, however, is not as big as in the case of deployment. Only for N = 18, the intersection of GSTT and N×STT curves is visible, but it is alleged (cf. Fig. 7), due to the mentioned disturbance in the mean time caused by one outstanding value in the termination process time. The maximum gain on termination bundling is here achieved at a very small bundle size (N = 2-3, cf. Fig. 8), but above this point, it is still time-efficient to combine deployment requests.

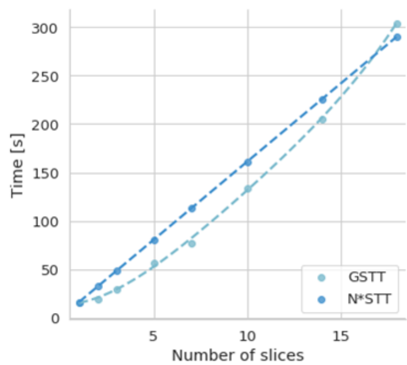

Fig. 7. Measured GSTT and N×STT (for N = 1, 2, 3, 5, 7, 10, 14, 18).

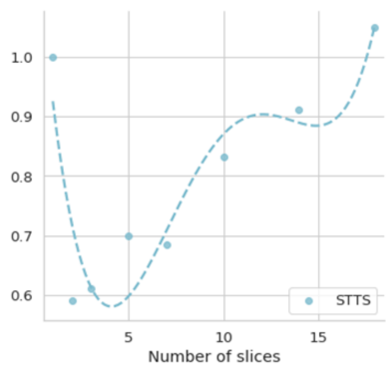

Fig. 8. Calculated STTS (for N = 1, 2, 3, 5, 7, 10, 14, 18).

The averaged results of RAM and CPU usage for a different number of slice instances during the termination phase are presented in Fig. 9 and Fig. 10. The memory usage curves (cf. Fig. 9) show a monotonic decrease (note: in the case of the curve for N = 18, the sudden jump up towards the end is related to the aforementioned disturbance in the average memory usage by one process out of 10, significantly longer than the others; the data series for the others ended earlier). For CPU usage curves during termination, the characteristic is a spike at the end of the process for all N values. Considering the timescale, it can be observed that the respective curves of both deployment and termination processes approximately mirror each other. It has to be noted that the NSI termination time is, on average only 50% shorter than NSI deployment time.

C. Discussion on results

The comparison of both, deployment and termination processes performed simultaneously (bundled) and one-by-one approach, in the most demanding case (N = 18), is shown in Fig. 11 and Fig. 12 respectively. In this case, the termination process takes approximately the same time as the deployment

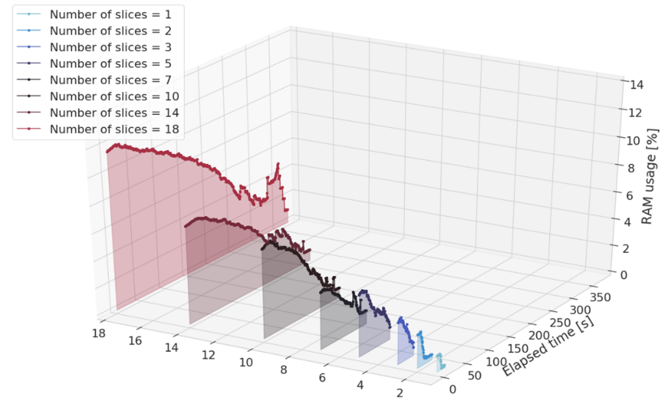

Fig. 9. Comparison of RAM usage measurement during the bundled termination (for N = 1, 2, 3, 5, 7, 10, 14, 18).

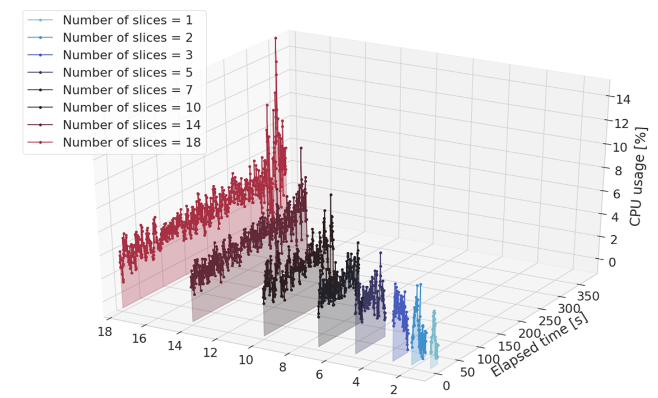

Fig. 10. Comparison of CPU usage measurement during the bundled termination (for N = 1, 2, 3, 5, 7, 10, 14, 18).

time, whereas there is still a significant time gain for bundled deployment in the case (cf. Fig. 11). Considering the demand for resources during both processes, there is no significant difference in CPU utilization. However, the one-by-one approach is clearly more memory-eating (cf. Fig. 11).

It should be noted that the overall time of both processes in the one-by-one approach was calculated excluding the intervals between confirmations of the NSI deployment and sending the next deployment request (only the request-confirmation intervals were considered). However, under the test conditions, the delay caused by OSM and its communication with VIM (order of ms) was negligible for the timescale of the entire process. The VIM level processes had an overwhelming effect, but verification of the total dynamics should also be performed for OSM under stress conditions.

Comparing the results for Magma vEPC with those obtained for lightweight Virtual Machines (VMs) [12], one can see that in the case of the latter ones, the SDTS curve had a knee shape, falling to the level of 0.2 and stabilizing at this level (slight increase to the right), while for the current measurements the SDTS curve has an elbow shape, showing a significant increase to the right at much lower values of N. There are two major differences between the environments: (i) 12-thread CPU previously versus 32-thread CPU currently (however, supporting respectively 50 and 18 NSIs), and (ii) one-layer virtualization previously (NFV only) versus two-layer one currently (NFV +

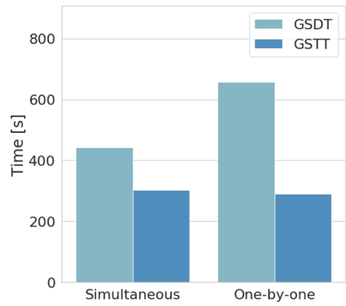

Fig. 11. Comparison of the overall time needed for simultaneous and one-by-one deployment and termination, N = 18.

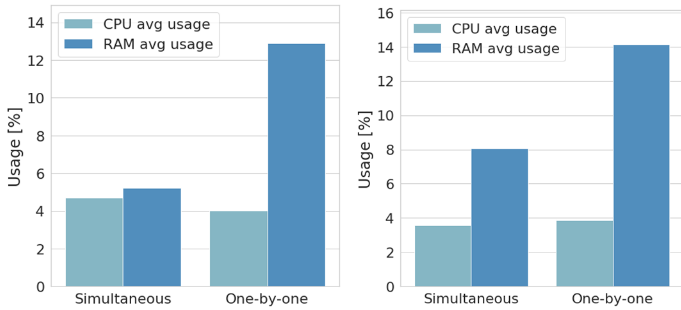

Fig. 12. Average values of CPU and RAM usage for both approaches to deployment (left) and termination (right), N = 18.

containerization). During the trials, the underlying inter-layer approach to the deployment of the worker nodes (VMs with Docker container runtime) for implementation of containers (embedding the network functions) was left unchanged. The impact at this level of either one-by-one strategy or bundled strategy on the orchestrator performance should be further studied and validated.

VI. Summary and Conclusion

The MANO orchestration is a complex process composed of multiple sub-processes. Its implementation is even more complicated as it has to deal with the deployed platform and operating system internal processes (garbage collection, etc.) that can be scheduled randomly by the operating system. Therefore, the full analytical approach to the orchestration process is problematic. It is, however, possible to make some assessment of the orchestration performance using experiments. This is the way which we have followed. The results of the experiments that we have made for performance evaluation of the OSM, the ETSI-compliant MANO orchestrator, are a bit surprising. Overall, the OSM orchestrator has proved to achieve very decent performance. The duration of a slice instantiation or termination phases are not negligible but at a satisfactory level. It is surprising that in the experiments the slice termination time was relatively high and, in all cases, was equal or exceeded 50% of slice deployment time. The deployment of a single Magma vEPC slice took $\sim$30 s, which is surprisingly fast. On the other hand, the typical termination time for a single instance, which took $\sim$20 s, is much longer than our expectations. With the increasing number of slices, the slice deployment and termination times increased less than linearly, which shows the excellent scalability of OSM. The "bundled mode" of slice orchestration has outperformed the one-by-one approach. In real life, however, such an approach will be hardly achievable – the requests are sent randomly and the scheduling of slice deployment or termination requests can be used only in large networks. Nonetheless, orchestration scheduling can be beneficial for network slice calendaring i.e. deploying or terminating a slice for the predetermined date.

The non-monotonic SDTS and STTS curves shape cannot be explained without the analysis of the performance of the internal blocks of the orchestration platform. Such an approach is hardly achievable, however, it is needed to find the performance bottleneck.

In future works, we plan to obtain more results for different slice templates and infrastructure configurations (more servers). Moreover, some experiments concerning the scalability of resource scaling (during slice runtime), omitted in the paper, are planned to be done.

Acknowledgment

The authors thank Michał Wierzbicki for his significant contribution to the preparation of the test environment and to the measurements.

References

[1] NGMN Alliance, "NGMN 5G White Paper", Jan. 2015.
[2] Magma, "An Open-Source Platform For Building Carrier-Grade Networks", [Online]. Available: https://www.magmacore.org/ (accessed Jul. 1, 2021).
[3] S. Kukliński and L. Tomaszewski, "Key Performance Indicators for 5G network slicing", 2019 IEEE Conference on Network Softwarization (NetSoft), 2019, pp. 464-471, doi:10.1109/NETSOFT.2019.8806692.
[4] 3GPP, "Management and orchestration; 5G performance measurements", 3GPP TS 28.552, ver. 17.3.1, Jun. 2021.
[5] 3GPP, "Management and orchestration; 5G end to end Key Performance Indicators (KPI)", 3GPP TS 28.554, ver. 17.3.0, Jun. 2021.
[6] ETSI, "Requirements and interfaces specification for management of NFV-MANO", ETSI GS NFV-IFA 031, V3.3.1, Sep. 2019.
[7] ETSI, "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Performance Measurements Specification", ETSI GS NFV-IFA 027, V4.2.1, May 2021.
[8] ETSI, "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Network Service Templates Specification", ETSI GS NFV-IFA 014, V4.2.1, May 2021.
[9] 3GPP, "Management of network slicing in mobile networks; Architecture framework", 3GPP TS 28.533, ver. 16.7.0, Apr. 2021.
[10] ETSI, "Open Source MANO", [Online]. Available: https://osm.etsi.org/ (accessed Jul. 1, 2021).
[11] ETSI, https://osm.etsi.org/wikipub/index.php/OSM_Release_FIVE
[12] M. A. Kourtis et al., "5G Network Slicing Enabling Edge Services", 2020 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), 2020, pp. 155160, doi:10.1109/NFV-SDN50289.2020.9289880.
[13] Yilma G. M., Yousaf F. Z. , Sciancalepore V., Costa-Perez X., "Benchmarking Open-Source NFV MANO Systems: OSM and ONAP". Computer Communications, vol. 161, pp. 86-98, ISSN 0140-3664, 2020. doi:10.1016/j.comcom.2020.07.013
[14] Yilma G. M., Yousaf F. Z. , Sciancalepore V., Costa-Perez X., "On the Challenges and KPIs for Benchmarking Open-Source NFV MANO Systems: OSM vs ONAP.", 2019, [Online]. Available: https://arxiv.org/abs/1904.10697 (accessed Jul. 19, 2021)
[15] ETSI, "Network Functions Virtualisation (NFV); Management and Orchestration", ETSI GS NFV-MAN 001, V1.1.1, Dec. 2014.
[16] OSM, "OSM architecture and functions", [Online]. Available: https://osm.etsi.org/docs/user-guide/02-osm-architecture-and-functions.html (accessed Jul. 1, 2021).
[17] ETSI, "Network Functions Virtualisation (NFV) Release 3; Protocols and Data Models; NFV descriptors based on YANG Specification", ETSI GS NFV-SOL 006, V3.3.1, Aug. 2020.