A novel concept for the migration of workloads using virtual machines that can mitigate limitations of traditional service deployment concepts

Henry-Norbert Cocos, Christian Baun, and Martin Kappes

Frankfurt University of Applied Sciences
Frankfurt am Main D-60318, Germany
{cocos,christianbaun,kappes}@fb2.fra-uas.de

Abstract. The project aims to investigate the outsourcing of cloud services to clients by including heterogeneous, client-side resources and cooperation between cloud and client in the execution of services. For this purpose, an offline-first strategy shall be formulated, developed, and investigated to be used as a counterpart to the cloud-first strategy and enable the utilization of services without a permanent cloud connection. The scope of the offline-first strategy is to achieve autonomy for applications while cooperating with cloud services. To achieve the independence and simultaneous cooperation of applications in this context, methods of service migration are to be researched to enable the provision of services. In summary, this should increase resilience when using services, improve the stability of services along the cloud-to-edge supply, and allow transparent use for the end user. Furthermore, lab experiments were conducted to investigate the behavior of virtual machines during live migration in two setups, simulating a WAN link. The experiments show that migration over a WAN link takes five times longer for RAM-intensive tasks than in a LAN. This provides initial insights into selecting the technological basis for this project.

Keywords: Offline-First-Strategy · virtualization · virtual overlay-networks · resilience in Cloud Computing · computation offloading · resource allocation · resource provisioning · service migration

1 Introduction

At the Frankfurt University of Applied Sciences, the SKILL/VL project (Strategic Competence Platform - Innovative Learning and Teaching - Virtualization of Distributed Environments for Teaching) [8] has implemented a digital offer for teaching. The platform enables lecturers, students, and staff to create virtual landscapes to integrate them into their teaching. Figure 1 shows the architecture of the SKILL/VL platform.

The SKILL/VL project platform is operated on 12 high-performance servers, which run in a cluster. To provide virtual learning environments, the SKILL/VL

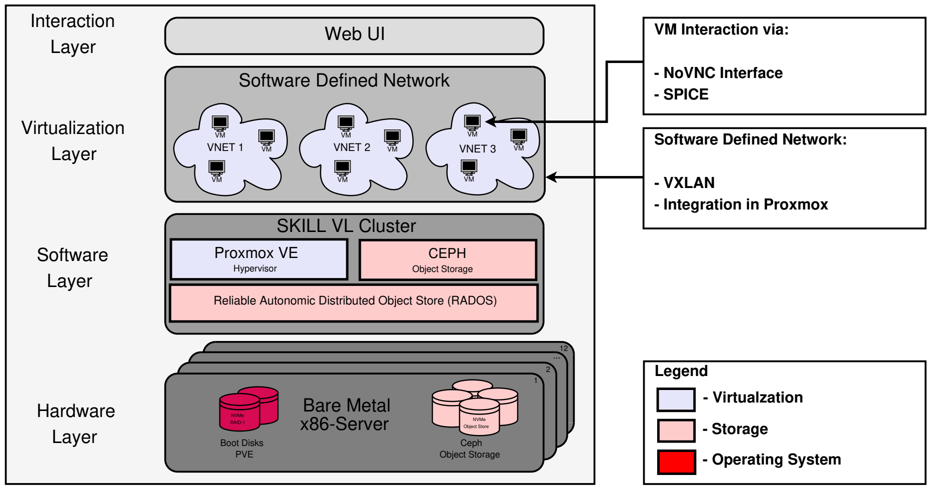

Fig. 1: Architecture SKILL/VL

platform uses the Proxmox [4] hypervisor, which is based on the free virtualization technology QEMU/KVM [5,2]. In addition, Proxmox offers the option of integrating the distributed storage service CEPH [6] to operate virtual machines. Proxmox allows the operation of virtual learning environments in software-defined virtual overlay networks. This allows the virtual machines to be configured in their virtual network segments without risking collisions with other networks so that users can configure the networks of their virtual learning environments freely without considering other participants.

The SKILL/VL platform is an archetypal application of the cloud computing concept and a reasonable basis for explaining the doctoral project. As with other cloud services, the SKILL/VL platform provides for online use of the service. It uses classic cloud computing methods, and many issues with cloud services are already apparent here. For example, the services must always be available online. Furthermore, the services only run on the server side and do not consider client resources. In addition, considerable latency is expected when using the service via a WAN (wide area network) connection. These problems lead to considering extending cloud services to the user's end devices and, simultaneously, the offline capability of services. In this way, the classic concept of cloud computing will be expanded, and a new dimension for the operation of cloud services will be explored. The following sections take a closer look at the problem definition of the doctoral project.

The paper aims to explain a novel offline-first strategy for migrating virtualized workloads over networking boundaries and reducing latency on the end-device level by bringing workloads closer to the end user and harnessing heterogeneous environments. The theory is established to demonstrate this prospect, and a field experiment on the example of the SKILL/VL platform is presented. In order to have a ground truth for the field experiments and understand the migration process and the limitations of this use case, laboratory experiments for the migration of virtual workloads are conducted. The WAN link is simulated using tools like `netem`.

2 State of the art

The project originates in distributed systems [15] and is based on the methods and technologies of cloud computing [9]. The National Institute of Standards and Technology (NIST) characterizes cloud computing by the five properties [12] on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured services. The resources offered are transparent and available to the customer in (seemingly) infinite quantities. Based on the service properties, cloud computing services are categorized according to their degree of abstraction into infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Infrastructure services give users complete control over operating systems, networks, storage connections, and applications.

Based on the hybrid cloud deployment model, multi-cloud environments have become established in the industry [13]. These have the advantage that you can avoid a vendor lock-in, where you become dependent on a single service provider, and can, therefore, pursue a flexible cloud strategy. In multi-cloud environments, services or parts are distributed or operated in parallel by different providers. This has economic and organizational advantages, as the distributed services can be used more flexibly and lead to increased availability of the service offering. However, this strategy also has a disadvantage, as the range of services offered by the CSP could be more transparent, and it is impossible to make them uniformly usable. Sky Computing addresses the lack of overview and inconsistent usability of the services offered by public CSPs [16]. Here, a further abstraction layer will be placed over the CSP's offering, enabling users to provide cloud services uniformly, regardless of their operating location.

The always-online principle is a general prerequisite for using cloud computing services. It has several advantages, such as consolidating resources through the centralized operation of services, independent provisioning, and use of services by the user. An offline offering should be created that allows users to use a service even without a constant internet connection while simultaneously providing the option of cooperating with an online cloud service. However, this requires the creation of technical and methodological foundations that enable the operation of services across the cloud infrastructure. The idea of a virtual PC environment is not new and has been considered in research for decades. For example, Steckermeier and Hauck [7,14] were already researching a concept for virtual PCs in 2000. However, the technical possibilities have developed further in recent decades, enabling the operation and distribution of applications across heterogeneous environments using suitable virtualization technologies.

3 Offline-First-Strategies in Cloud Computing

An offline-first strategy in cloud computing is about more than just improving usability and availability. It's about implementing the right (technical and methodological) options to achieve these goals. The five essential resources that are investigated in an offline-first-strategy are:

1. **computing capacity**
2. **network resources**
3. **storage resources**
4. **software**
5. **services**

The distribution of computing capacities plays an important role here, as suitable technology must be available for their isolation and operation. The options here are virtual machines and container runtime environments. Network resources also play an essential role in communication between the individual and distributed services. Virtualization technologies from Software Defined Networking (SDN) are also suitable. The distribution of services also leads to data distribution, which must be synchronized and orchestrated using a suitable option. This problem is frequently encountered in distributed systems [15,10].

The scaling of services can be done in multiple dimensions, which can be described in spatial dimensions. The vertical scaling of services (or scaling up) can be achieved by adding resources (CPU, RAM, etc.). Scaling on a horizontal level is possible by adding service instances (also known as scaling out). Geographical scaling, which places services on the client end device closer to the user and thus geographically further away from the core service in the cloud, receives little or no attention in practice. This dimension (referred to here as scaling away) is to be taken into account in this project. Service instances and their characteristics should benefit from their proximity to the user and be researched as part of this project.

4 Practical use case in SKILL/VL

The theories described in the previous paragraph should be implemented and tested in practice appropriately. The SKILL/VL platform mentioned above is ideal for this. As already described on page 2, the SKILL/VL project forms a platform for virtual teaching and uses a powerful hypervisor and virtual overlay networks based on the VXLAN (Virtual Extensible Local Area Network) [11] protocol to implement encapsulated, virtual learning spaces. VXLAN encapsulates Ethernet frames in UDP segments. This allows extending and operating logical Local Area Networks (LAN) across network boundaries (routers and gateways) so that hosts and virtual machines can be operated in a logical VXLAN but have a physical Layer 2 LAN connection, which, like a physical LAN, assigns a common broadcast domain to all participants in the local network.

Another exciting application scenario in the SKILL/VL project would be expanding virtual and logical resources to include physical ones. The Frankfurt University of Applied Sciences operates a laboratory for labs in the Computer Networks course with Twenty-four computers. The lab resources are, therefore, limited, and the number of students participating in this course is minimal. Figure 2 shows schematically the idea behind this field test.

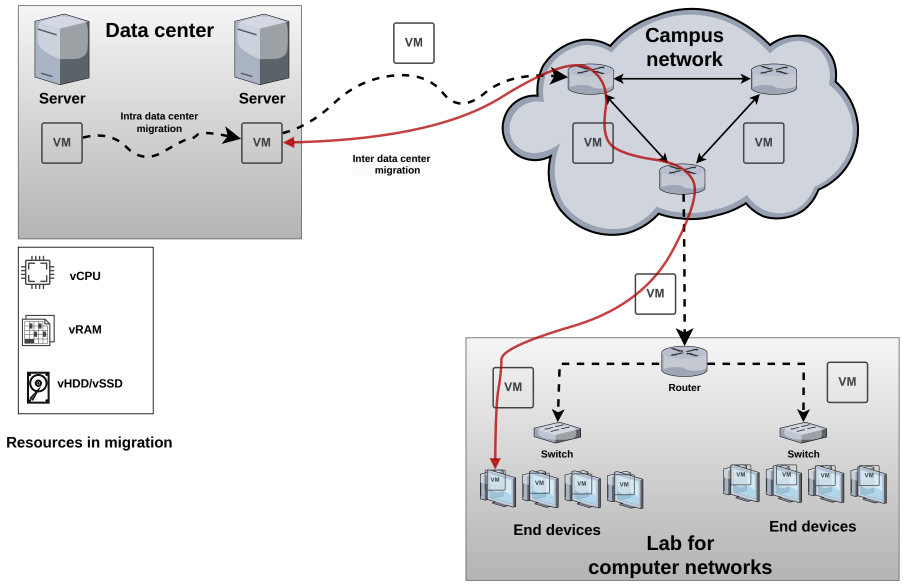

Fig. 2: Possible extension of SKILL/VL

5 Lab experiments

Lab experiments were conducted to get a first insight into the behavior of virtual machines and their characteristics while migrating machines over the network. As a basis for comparison and starting ground, these experiments are conducted to set a ground truth for field experiments as described in section 4. Three machines were equipped with the operating system Debian 12 for the lab experiments. On top of that, QEMU/KVM was installed to create a uniform and, therefore, a comparable environment for the first test. VMs were configured and migrated in a local area network (LAN) from a source server to a target server. A test application was used to simulate loads in different levels (e.g., 25%, 50%, 80%, 100%) for various system resources (CPU and RAM r/w). These resources were simulated using the `stress` tool.

6 Results

Figure 3 shows the measured times for migrating a virtual machine via a network with a bandwidth of 1 Gbit/s (net 938 Mbit/s) and as a comparison over a bandwidth of 37 Mbit/s. Figure 3a shows that the migration time in the idle state is 12,547 seconds. The times are similar to the CPU load variation, which can also be observed for a bandwidth of 37 Mbit/s. However, a clear correlation between RAM utilization and migration time can be observed for variations in that value. It can be seen that there is a factor of over five between a migration in an idle state and a migration with a RAM utilization of 80%. This shows a direct correlation between RAM utilization and VM live migration time.

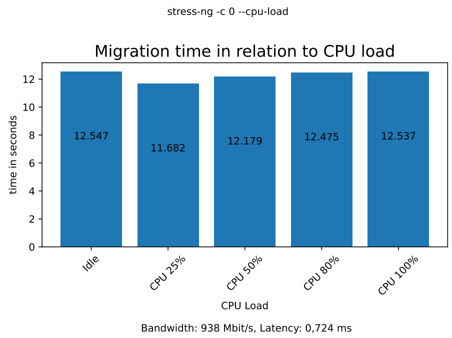

(a) Variation of CPU load (1 Gbit/s)

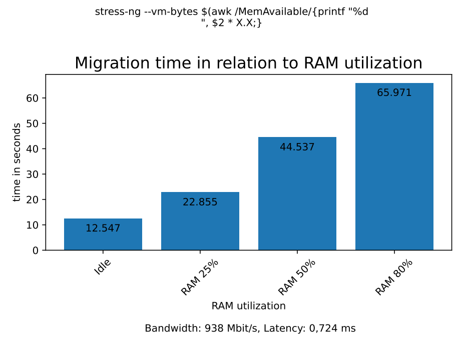

(b) Variation of RAM utilization (1 Gbit/s)

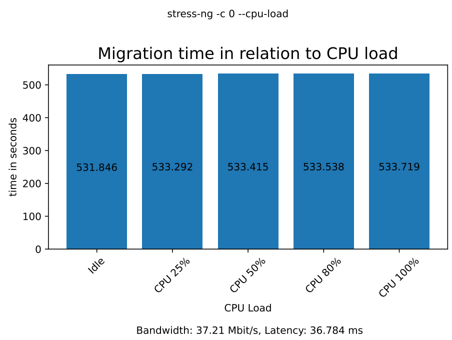

(c) Variation of CPU load (37 Mbit/s)

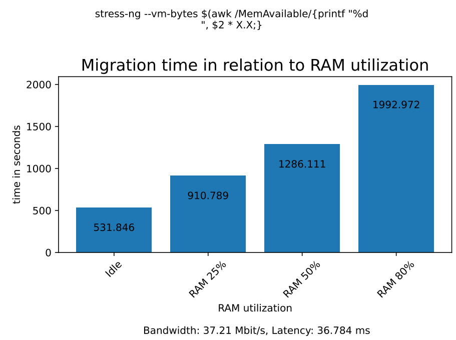

(d) Variation of RAM utilization (37 Mbit/s)

Fig. 3: Migration time bandwidth 1 Gbit/s and 37Mbit/s

Figure 3c and 3d shows the times for migrating a virtual machine via a network with a bandwidth of 37 Mbit/s. This value and the latency of 36,784 milliseconds were determined empirically from measurements and defined for comparison. It can be seen that the migration time of the VM remains virtually constant over the CPU load, which suggests that CPU-intensive applications have virtually no additional influence on the migration of virtual machines. The base value for a migration is ≈ 531 seconds. In contrast, the variation in RAM utilization apparently influences the migration time. For example, with a RAM utilization of 80%, a migration time of ≈ 1992 seconds can be expected, which is approx. 33 minutes. The factor here is 3.7 between a migration in an idle state and a migration with a RAM utilization of 80%. This confirms the assumption that a migration of VMs with RAM-intensive applications is enormous. A RAM utilization of 100% was deliberately avoided, as this has occasionally led to the phenomenon of an infinitely long migration. This can be attributed to the RAM's continuous "pollution".

7 Conclusion

This paper has proposed an offline-first strategy for distributed, heterogeneous, and virtualized resources. In the proposed concept, resources for terminal devices (client PC, laptop, etc.) shall be used to reduce server-side load and make services offline available on the client side. The proposed framework will benefit the operation of services, especially when network bandwidths are low. Furthermore, a practical application of the offline-first strategy for the project SKILL/VL has been presented. The application shall make server-side resources available to end devices in a physical lab and establish a hybrid setup of shared heterogeneous resources. Section 6 presented and discussed the results of lab experiments on migrating virtual machines over a network. A correlation between the RAM utilization and the migration time can be observed, resulting in the conclusion that the initial migration of workloads is feasible but time-consuming. Furthermore, the migration of services affects their operation.

The experiments brought many findings. Using QEMU/KVM is beneficial in a local, server-based setup, but implementing an offline-first strategy is not a good choice since it brings many limitations. For the live migration of virtual machines, the server- and client-side hypervisor need to be the same (in version as well!). Otherwise, a migration of workloads cannot take place. Also, the hardware type of the CPU for the server and client needs to be identical (CPU architecture and vendor). A live migration of virtual machines between an Intel and an AMD processor is not possible, posing strict limitations on the use of end devices. Another limiting factor is the necessity of shared storage for the virtual machines' virtual hard disks. The shared storage medium must be online and available throughout the migration process. On top of that, it is bound to the network bandwidth limitation, posing another bottleneck in using network resources. However, using QEMU/KVM can be promising for migrating workloads in legacy applications. In this scenario, the effort required to reconstruct applications could be reduced.

8 Future Work

In a next step the influence of virtual overlay networks on the migration of virtual machines shall be investigated. For that the lab experiments form section 5 shall be extended by the use of Open vSwitch and the construction of a VXLAN tunnel between the server and the client end device. Furthermore the hybrid setup as described in section 4 shall be implemented and investigated and set a first field experiment for an offline-first strategy in practice. Based on the lab experiments the use of container virtualization technologies like Docker [1] or OCI Containers [3] shall be investigated as an alternative to the use of heavyweight virtualization technologies investigated in the lab experiments presented in this paper.

Acknowledgments. The SKILL/VL project described in this paper is part of the SKILL project network (Strategic Competence Platform - Innovative Learning and Teaching), which is funded by the *"Initiative Hochschullehre durch Digitalisierung stärken der Stiftung Innovation in der Hochschullehre"*, at Frankfurt University of Applied Sciences.

References

1. Docker, https://www.docker.com/, accessed: 29.03.2024
2. KVM, https://linux-kvm.org/page/Main_Page, accessed: 29.03.2024
3. Open Conteiner Initiative, https://opencontainers.org/, accessed: 29.03.2024
4. Proxmox Virtual Environment, https://www.proxmox.com/en/proxmox-virtual-environment/overview, accessed: 29.03.2024
5. QEMU, https://www.qemu.org/, accessed: 29.03.2024
6. The Ceph stack: architectural overview, https://ceph.io/en/discover/technology/, accessed: 29.03.2024
7. Virtual Private Computing – A New Paradigm for Component-Based Distributed Applications, accessed: 31.03.2024
8. Baun, C., Kappes, M., Cocos, H.N., Koch, M.M.: Eine Plattform zur Erstellung und Verwendung komplexer virtualisierter IT-Strukturen in Lehre und Forschung mit Open Source Software. Informatik Spektrum (Feb 2024). https://doi.org/10.1007/s00287-024-01559-x, http://dx.doi.org/10.1007/s00287-024-01559-x
9. Baun, C., Kunze, M., Nimis, J., Tai, S.: Cloud computing. Springer, Berlin, Germany (jul 2011)
10. Gorton, I.: Foundations of scalable systems. O'Reilly Media, Sebastopol, CA (Jul 2022)
11. Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., Wright, C.: Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks. RFC 7348 (Aug 2014). https://doi.org/10.17487/RFC7348, https://www.rfc-editor.org/info/rfc7348
12. Mell, P., Grance, T.: The nist definition of cloud computing (2011-09-28 2011). https://doi.org/https://doi.org/10.6028/NIST.SP.800-145
13. Mulder, J.: Multi-Cloud Architecture and Governance. Packt Publishing, Birmingham, England (Dec 2020)
14. Steckermeier, M.: Virtual Private Computing A new paradigm for omnipresent, internet-based computing, https://www4.cs.fau.de/Projects/VPC/, accessed: 31.03.2024
15. Van Steen, M., Tanenbaum, A.S.: Distributed Systems. Createspace Independent Publishing Platform, North Charleston, SC (Feb 2017)
16. Yang, Z., Wu, Z., Luo, M., Chiang, W.L., Bhardwaj, R., Kwon, W., Zhuang, S., Luan, F.S., Mittal, G., Shenker, S., Stoica, I.: SkyPilot: An intercloud broker for sky computing. In: 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23). pp. 437–455. USENIX Association, Boston, MA (Apr 2023), https://www.usenix.org/conference/nsdi23/presentation/yang-zongheng