1st Conference on Research Data Infrastructure

Poster presentations II

https://doi.org/10.52825/CoRDI.v1i.413

Published: 07 Sept. 2023

Improving the research desktop experience for OpenStack VDI

Integrating hardware accelerated rendering and remote transport

Dirk von Suchodoletz[1] [https://orcid.org/0000-0002-4382-5104], Yi Sun[2] [https://orcid.org/0000-0002-7636-0200], Jean-Karim Hériché[2] [https://orcid.org/0000-0001-6867-9425] , and Manuel J. Messner[1]

[1]Computer Center, University of Freiburg, Germany

[2]EMBL, Heidelberg, Germany

Abstract: Using virtual machines with dedicated rendering and remote access capabilities, virtual workplaces for various usecases can be created, and then accessed from anywhere at any time. If this is to happen on a large scale in the cloud, so-called VDI for the dynamic provision of virtual desktops play an increasingly important role. A sustainable VDI should be freely available to everyone for modification and redistribution at no cost, be scalable, and should support various desktop use cases with different resource requirements. Some use cases involve hundreds of similar VM running in parallel, which requires proper resource planning.

Keywords: Virtual Desktop Infrastructure, hardware accelerated rendering, remote access, NFDI4BioImage, DataPLANT

1 Motivation

The objective of this initiative in the **Enabling RDM track** is primarily to form a special interest group for hardware accelerated VDI on cloud infrastructure as we expect the uptake of the topic in other consortia as well. For various use cases in research, education and training Virtual Desktop Infrastructure (VDI) would provide an efficient means in the context of the NFDI. The available commercial frameworks are quite expensive (for large and scattered audiences) and do not necessarily cater to the core needs of researchers. Thus we are looking for Open Source alternatives based on well established software frameworks. We identified three relevant use cases:

1. Provide remote access to desktop environments to handle large scale high resolution imaging data in various research domains (near to the location of the data to avoid tedious transfers) and Remote visualization in High Performance Computing
2. Streamline training resources to allow more flexible remote teaching and working models for university staff in research
3. Provide controlled access to sensitive data in a protected environment

Many computer centers already offer OpenStack cloud infrastructure for various purposes. In Freiburg we are investigating the options to transfer our experience from our decentralized bwLehrpool VDI solution to OpenStack to offer an Open Source backed VDI (OSVDI).

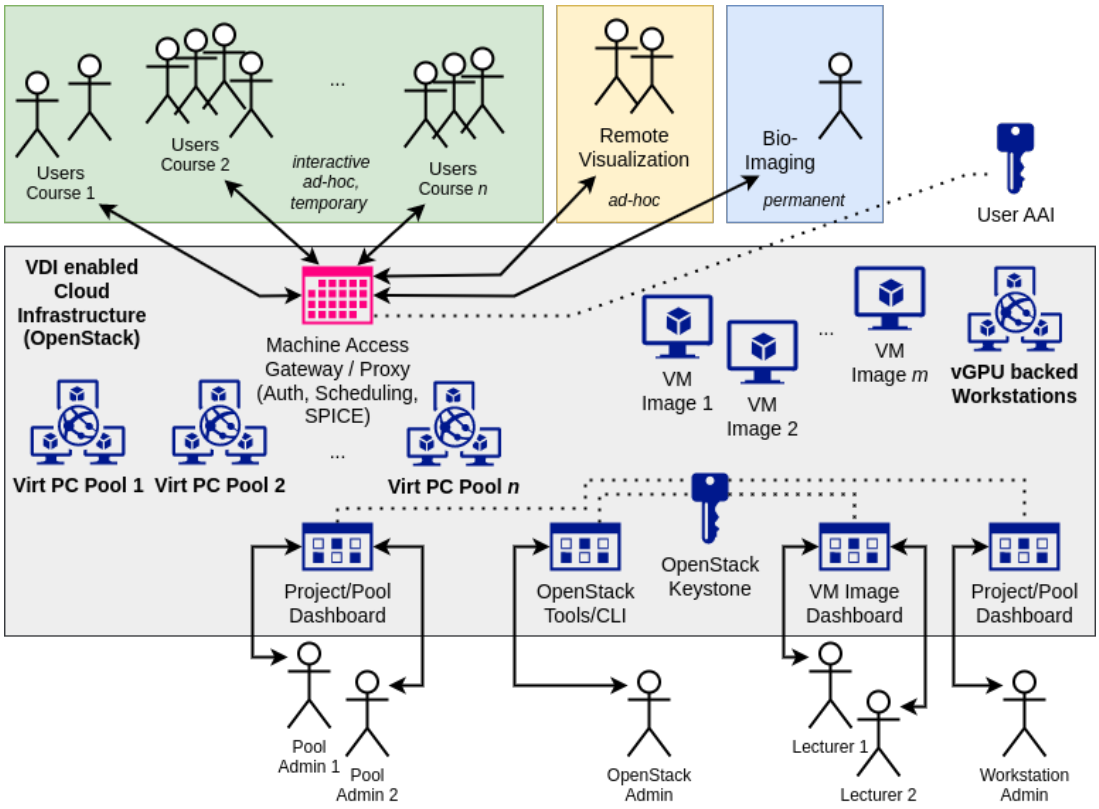

Figure 1. Various use cases and general infrastructure framework of the OSVDI framework.

OpenStack is a free cloud platform, most commonly deployed as Infrastructure-as-a-Service. The platform is composed of several components and services. Each component and service is responsible for a certain set of tasks and provides a RESTful API for communication. OpenStack's main focus is to provide an infrastructure for VMs, their storage and their network. Besides the mediated GPU passthrough support in OpenStack we need resource scheduling for cloud computing.

2 Objective

Our vision for a next generation approach is to have a public VDI web application where users log in, select a VM they want to use, and finally be scheduled to an appropriate cloud node. Additional logic could be added, e.g. limiting or skipping the VM selection for a researcher or training group depending on their location or time of day. Our focus is mainly on the mediated GPU passthrough for further development of an OSVDI because this approach combines the flexibility of emulation and paravirtualization with the performance boost of direct GPU passthrough. Linux as the host system already provides the *mdev* subsystem and tools for mediated devices (vGPUs) and their device drivers. Using this subsystem has the major advantage that the Linux host system can manage all vGPUs and mediate shared access. The direct access to a framebuffer of a vGPU means in terms of an OSVDI that the Linux host system can obtain the graphics output of any VM (virtual desktop) and can control those output for further processing and transfer to remote (thin) clients. Access to a VM session can then be

implemented via two methods: A browser-based approach, using modern technologies like WebAssembly, WebUSB and MediaDevices, resulting in immediate access from a wide range of devices. Still, the alternative approach of using a dedicated native application the user has to install first can offer even greater integration with the user's system, as well as yielding better performance depending on its use case.

Our goal is to structure the development of the missing bits and pieces which got identified by a paper published in 2021 [1]:

- Defining and implementing an access gateway to VDI on OpenStack clouds
- Evaluate and improve the existing remote access protocol SPICE and the associated web and native clients
- Adapt OpenStack scheduling to allow the hosting of the relevant usecases and in-advance reservation of resources for e.g. trainings [2]
- Coordinate these efforts with other projects especially in the context of EOSC interested in this domain

3 Envisioned milestones

For the imagined OSVDI we plan several development cycles and a minimum viable product approach. In a precursor the existing Guacamole bwLehrpool remote access should get improved through hardware rendering and stream encoding deploying the Intel GVT-g desktop graphic architecture together with the KVM infrastructure as a Linux-based hypervisor and produce an assessment of ease to use and stability. This will get implemented as an enhanced bwLehrpool service and demonstrate the capabilities of the existing kernel drivers regarding GPU virtualization and hardware partitioning. We will use the SPICE client and Looking Glass as a test and performance measure when accessing the virtual framebuffer for AVC/H.264 encoding and transport. Upon this we will explore how to encode with low latency, and how to send it to browsers and display the content there with low latency. This provides a possible baseline to check certain expectations and features before delivering similar services like those for an OpenStack cloud.

In a further milestone, we focus on a basic VNC model (leaving further improvements of remote access to parallel or later developments) in the cloud including orchestration of resources which covers the scope of our contribution to NFDI4BIOIMAGE. This milestone starts to extend the OpenStack framework for missing components and modules. This milestone deals with the challenges of a suitable access broker to distribute users requesting certain types of desktops onto a suitable VM. The access broker includes the provisioning of basic interaction channels starting from a the input of the various usecases. While the previous step focused on a basic integration and the outline of strategic components the next milestone focuses on the special hardware virtualization and integration parts both from the viewpoint of the guest systems and as encoding devices from the host perspective. The remote access should enjoy at least an enhanced hardware-backed video stream transport model for the remote visual cloud. Later milestones should deal with further remote interaction channels and further features and improvements for typical VDI setups like suspend and resume of interactive desktop sessions. Starting during the second milestone measures should be taken to form a sustainable community and financing concept around the proposed service. Both ongoing support, code maintenance and future development are to be supported through some stable organizational structure.

Competing interests

The authors declare that they have no competing interests.

Funding

is funded by the German Research Foundation (DFG) within the framework of the NFDI – project number: 442077441 and EOSC-Life under grant agreement H2020-EU.1.4.1.1. EOSC-Life 824087.

Acknowledgements

Part of the activities and insights presented in this paper were made possible through the collaboration in the PePP project (FBM2020-VA77-8-01241) funded by the *Foundation for Innovation in Higher Education*, EOSC-Life and preliminary work in the bwLehrpool project supported by the Baden-Württemberg Ministry of Science, Research, and the Arts.

References

References

[1] M. Bentele, D. von Suchodoletz, M. Messner, and S. Rettberg, "Towards a GPU-Accelerated open source VDI for OpenStack," in *Cloud Computing*, M. R. Khosravi, Q. He, and H. Dai, Eds., Cham: Springer International Publishing, 2022, pp. 149–164, ISBN: 978-3-030-99191-3. DOI: 10.1007/978-3-030-99191-3_12.

[2] M. Bentele, D. von Suchodoletz, M. Messner, and R. Piliszek, *A resource-aware scheduling concept for an openstack-based vdi*, Accepted for publication at CloudComp 2023 - 12th EAI International Conference on Cloud Computing, 2023.