

Cloud Computing

To enable the deployment of scalable and reliable services in the cloud, we utilize the [de.NBI](#) cloud based on OpenStack with containerization and Kubernetes as core components. We briefly describe them in the following and link to further reading material.

Cloud Platform



openstack[®]

OpenStack is a free and open-source cloud operating system which gives the user fine-grained control of computing and storage resources. Virtual machines, network and storage settings can be managed through a browser based dashboard or via an API. The cloud platform of the German network for bioinformatics infrastructure [de.NBI](#) is entirely based on OpenStack which serves as the basic infrastructure to run the services of the Research Data Commons (RDC). Containerized applications can be deployed within OpenStack, which offers the option of either deploying the Docker containers directly on the virtual machines or using a Kubernetes cluster (see below).

Read more: <https://www.openstack.org/>

Containerization

[blocked URL](#)

The goal of containerization is to package a software component together with its runtime dependencies such that it becomes portable and can run on any host. This is important for scalability, because we want to be able to dynamically deploy instances of a software component on demand and consistently on potentially different host systems. Containers also provide very fast startup time, compared to VMs, because they feature a light runtime environment sharing the operating system kernel with the host, further facilitating dynamic deployment. To deploy an application on a new host, all that needs to be done is its image must be built or a prebuilt image needs to be pulled (downloaded) and then it can be run as a container. Containers also offer isolation from the host system and other containers, which provides security. We strive to deploy all of our services as containers to make use of their advantages.

Read more: <https://www.docker.com/resources/what-container/>

Kubernetes

[blocked URL](#)

Kubernetes can be used to manage the deployment of containerized applications on our cloud infrastructure. A single Kubernetes cluster manages multiple applications consisting of one or many Docker containers. It schedules containers to run on certain hosts dynamically and as such enables automatic scaling based on resource demand with fine-grained control available on which hosts certain containers should be run preferably. Containers can be monitored to automatically detect failure and redeploy them. Kubernetes offers complex networking abilities to let containers communicate between each other and with hosts outside the cluster in a controlled and secure fashion. A powerful API is available to manage Kubernetes clusters with role-based access control.

Read more: <https://kubernetes.io/docs/concepts/overview/>