

K8s L2 CNI for Containers and VMs

Students



Michael Brändli



Leandro Ceriani

Initial Situation: Containerization of applications has gained more and more growth in recent years and has become a standard in software development today. The container orchestration platform Kubernetes has therefore not without reason gained great popularity. It is therefore more understandable that companies and enterprises want to rely on this technology. The Institute for Network and Security (INS) offers students at the Eastern Switzerland University of Applied Sciences many laboratory exercises in network and security modules. The fundamentals of this infrastructure should now be migrated to Kubernetes. Individual architecture components only exist as virtual machines and not as containers as of today. The software KubeVirt enables the deployment of virtual machines in Kubernetes. Initial tests looked promising. However, it seems that while containers and virtual machines can work together, the communication between them can pose challenges when setting up a laboratory environment in the network field as required by the INS. For a standard enterprise operation, this works well. However, if a pure layer 2 is needed or even layer one connection between containers and virtual machines, it is not so easily implementable. Even more, not with the intent to have it in a multi-host setup running.

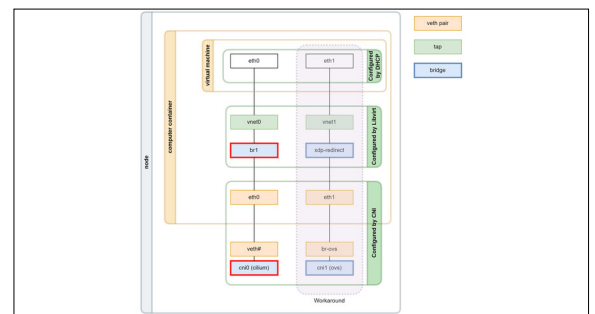
Approach: The goal was to show one or more proof of concepts of how the implementation of a layer 2 connection between virtual machines and containers could be possible, ideally, with the possibility of running this on multiple hosts. The analysis of various existing CNIs (container network interfaces) serves as the basis for the proof of concepts. These were checked for their function and it was tested how much the individual CNIs can contribute to an overall solution. Individual approaches and the technologies used by the CNIs were analyzed and documented.

Result: With the proof of concept one, it was shown that using Kube-OVN and macvtap as CNI and veth pair as virtual network cables, a layer 2 connection between virtual machines and containers can be established even on a multi-host setup. This setup still had some technological flaws, which made it impossible to run certain networking scenarios. The proof of concept one is quite extensive in terms of the manual effort required to set up a lab environment. With proof of concept two, the goal was to minimize this effort by eliminating a CNI and using veth pairs as virtual network cables. For this, a code improvement from the developer kvaps has been in the pipeline of KubeVirt for over a year but has yet to be integrated. Unfortunately, after dedicating a significant amount of time, the desired goal to implement this with the code improvement from the developer kvaps was not possible in the given time constraints of the semester

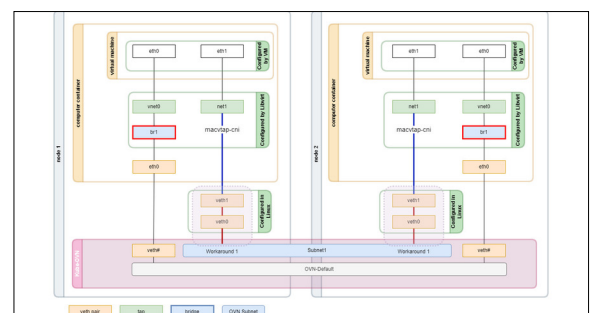
assignment.

On the whole, it has to be said that an own implementation of a CNI is the best solution for this INS project. What this custom CNI would have to fulfill and which technologies should be used by the existing CNIs is described in the chapter "Future Steps" in the technical documentation.

Issue with KubeVirt Own presentation



Proof of concept 1 Own presentation



Advisors

Urs Baumann, Yannick Zwicker

Subject Area
Miscellaneous

