LIVE VIRTUAL MACHINE MIGRATION IN BANDWIDTH-CONSTRAINED WAN ENVIRONMENTS

The Multi Conference on Computer Science and Information Systems
25th - 27th July 2026

Henry-Norbert Cocos, Christian Baun
{cocos,christianbaun}@fra-uas.de

Computer Science
Department of Computer Science and Engineering
Frankfurt University of Applied Sciences

Contents

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Project Virtualization of distributed environments for teaching[1]

SKILL/VL

Key features of SKILL/VL plattform[2]:

- Hyper converged infrastructure (cluster of 12 servers)
- Software Defined Storage
- Software Defined Networking
- QEMU/KVM as virtualization engine
- Strict use of open source software:
  - *Proxmox VE* as hypervisor (*KVM-based*)
  - *Ceph* as distributed object storage
  - *QEMU-Guest Agent* integration for configuration of VMs
  - *VXLAN* (Virtual Extensible LAN) integration for (virtual) network configuration

[1] Original: Strategische Kompetenzplattform - Innovativ Lernen und Lehren – Virtualisierung verteilter Umgebungen für die Lehre

[2] `http://www.virtuellelehre.de/`

SKILL/VL Architecture

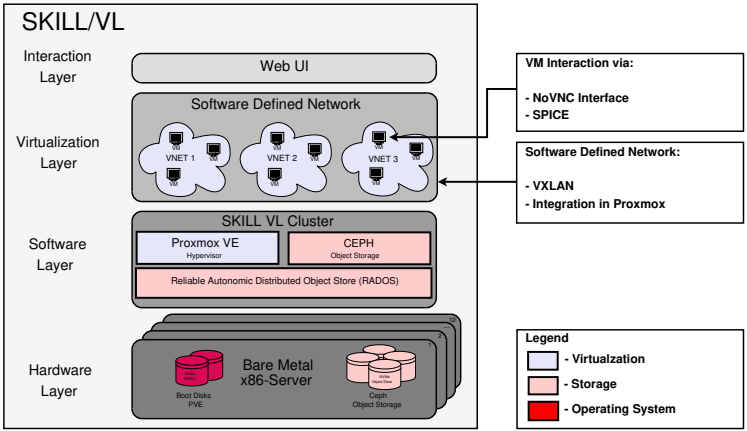

SKILL/VL plattform:

- **Interaction Layer**
  $\Longrightarrow$ React web UI
- **Virtualization Layer**
  $\Longrightarrow$ VMs and SDN network environment
- **Software Layer**
  $\Longrightarrow$ Proxmox VE and Ceph as plattform for virtual resources
- **Hardware Layer**
  $\Longrightarrow$ x86 Server, physical network and storage

Use Case Computer Networks Lab (1-237)

- Students access the SKILL/VL environment with client hardware over network.
- Students use the Workstations as equipment for their exercises.
- Students use lab environment to interact with SKILL/VL infrastructure.

Problem

The computer networks lab has "only" 28 workstations and the use of the SKILL/VL infrastructure generates load on the server infrastructure and network! The network causes latencies in the interaction with the server!

Opportunity

The resources accessed by users (virtual machines, networking, software) are virtualized!

Idea

Use client hardware as an enhancement of the service!

Proposal

- Use client infrastructure to enhance cloud service
- Migrate services from cloud to client whenever possible

Benefits

- Reduced network latency
- Reduced server load
- Increased autonomy and resilience
- Increased geographical scalability

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Research questions

1. What bandwidth and workload conditions allow feasible live VM migration for distributed laboratory environments?
2. How strongly does RAM dirty-page generation influence migration performance compared to CPU load?
3. Can VM migration realistically support offline-first laboratory architectures in WAN environments?

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Field experiment – SKILL/VL

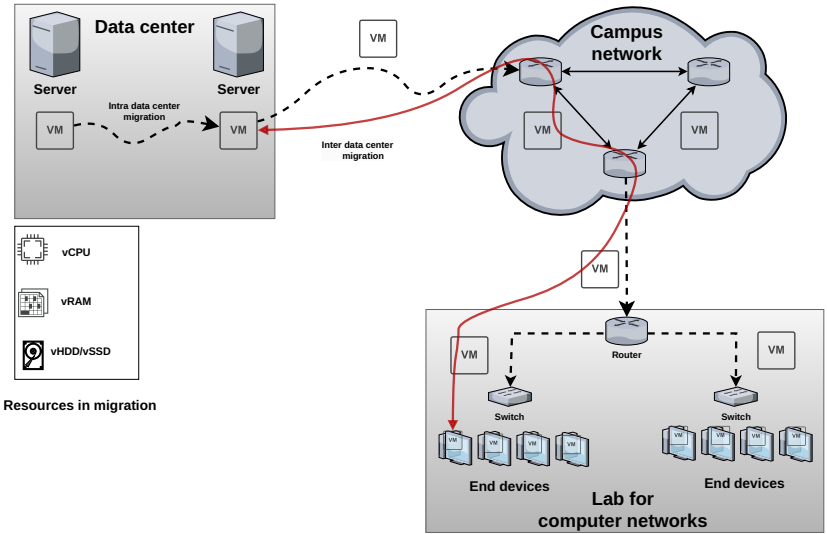

Lab experiment setup – VM migration

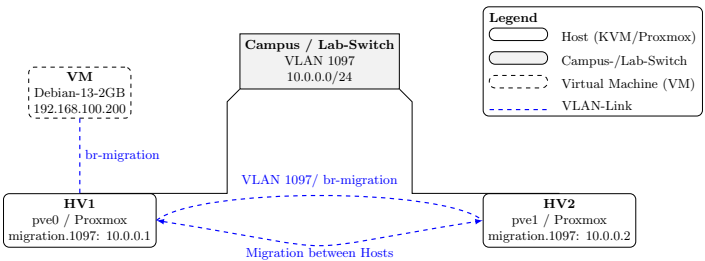

Figure: Lab experiment VM migration over LAN

- Lab experiments to establish a ground truth and a basis for comparison.
- Well defined environment for the experiments.
- Environment set up with QEMU/KVM and Proxmox VE.

Lab experiment setup – VM migration

- 10 migrations over LAN with variable bandwidth and workload parameters.
- **Bandwidths:** 10 Gbit/s, 1 Gbit/s, 500 Mbit/s 250 Mbit/s, 125 Mbit/s, 100 Mbit/s 50 Mbit/s, 25 Mbit/s, 10 Mbit/s.
- **Workloads:** 0% CPU, 25% CPU, 50% CPU, 75% CPU, 100% CPU
  0% RAM, 25% RAM, 50% RAM, 75% RAM, 100% RAM.

VM migration – Downtime for different scenarios

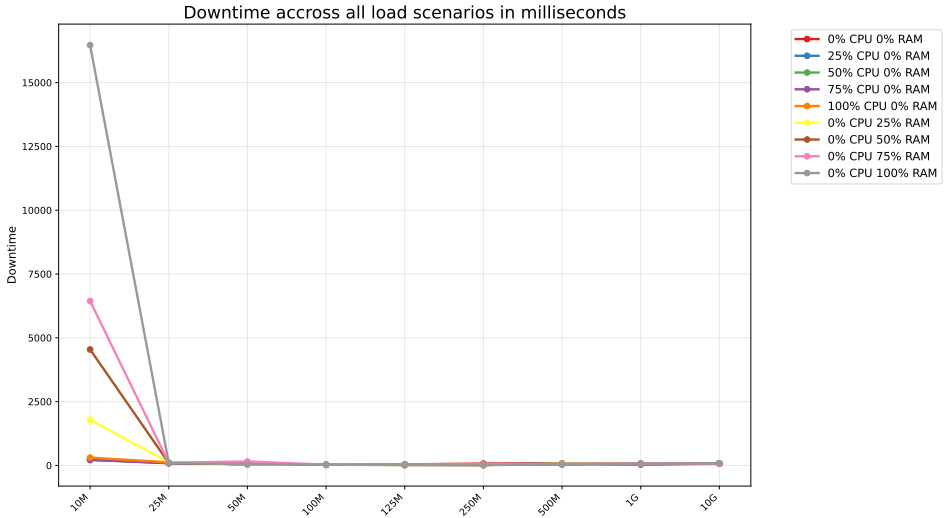

VM migration – Downtime over all scenarios

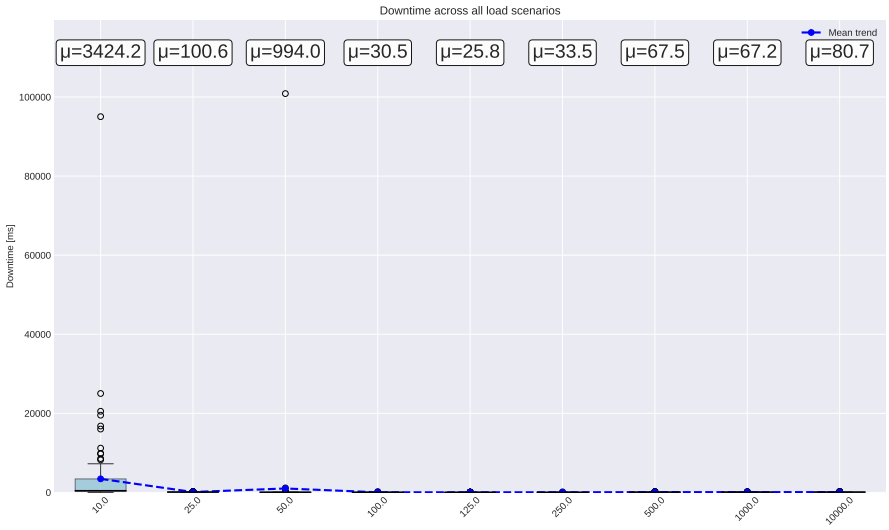

VM migration – Migration time for different scenarios

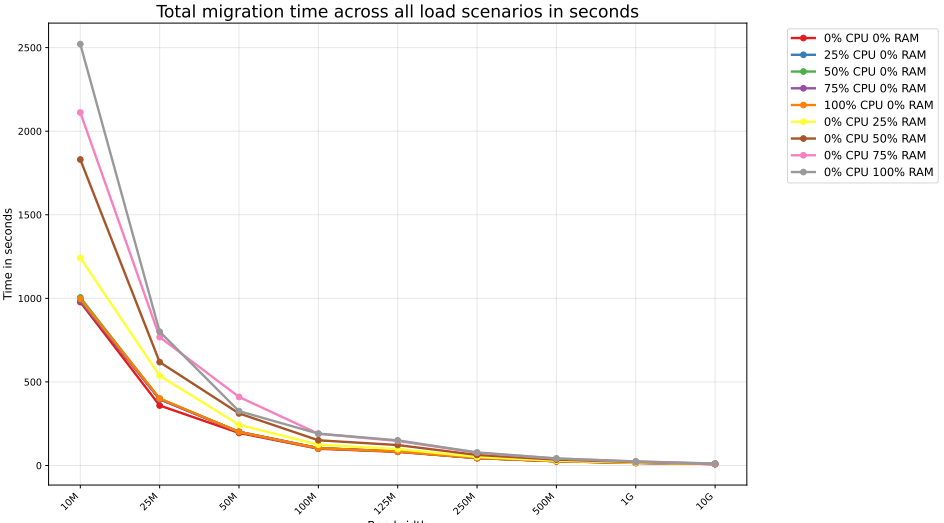

VM migration – Migration time over all scenarios

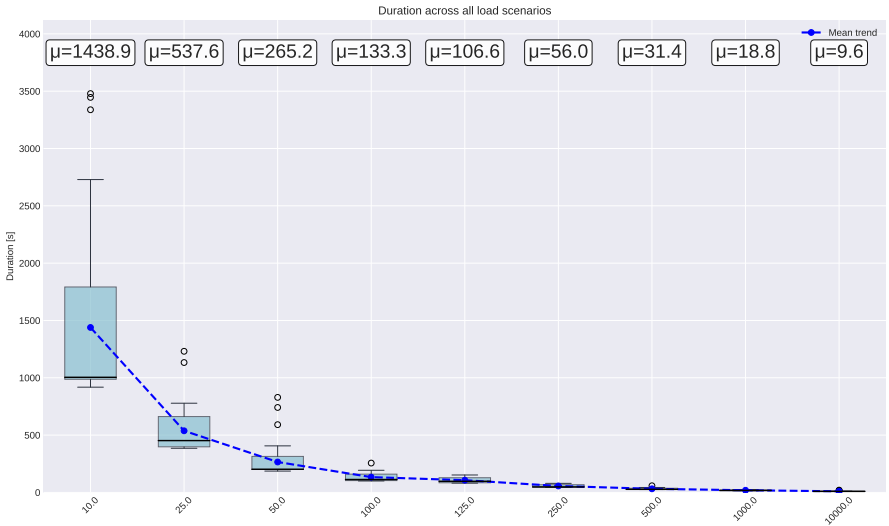

Insights

High bandwidth works well

At 1 Gbit/s and above, downtime remains low and migration time stays in the range of only a few seconds to a few tens of seconds.

Low bandwidths do not perform well

At 10 Mbit/s, migration becomes extremely slow, and downtime moves far beyond what is acceptable for interactive teaching scenarios.

RAM has a far bigger impact than CPU

- **10 Mbit/s - CPU-only** – peak 240.70 ms downtime.
- **10 Mbit/s - 100% RAM** – 16,470 ms peak downtime.
- **29 times higher downtime for RAM-only than CPU-only at 10 Mbit/s.**

Interpretation

- **Dirty-page retransmission** – More memory writes mean more pages must be copied again during pre-copy migration.
- **Throughput erosion** – Effective transfer speed declines because the migration process chases newly dirtied memory.
- **Downtime inflation** – The final switchover becomes more expensive when memory convergence is difficult.

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Limitations

- The WAN setup used 0 ms delay, so it represents a controlled bandwidth baseline rather than a fully realistic Internet scenario.
- High dirty-page rates can dramatically increase migration time or prevent migration entirely.
- At 100% RAM load, migration may become structurally infeasible when page dirtying outpaces network bandwidth.
- Increasing the downtime threshold can enable migration, but it also leads to longer actual downtime.

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Conclusion

The concept is Campus-ready but not suitable for WAN-setups!

What works:

- Intra-campus migration.
- Data-center-to-lab relocation.
- Scenarios with at least 100 Mbit/s, typically 1 Gbit/s or more.

What does not work:

- WAN-connected end devices at 50 Mbit/s and below.
- High-RAM-write workloads over constrained links.
- Interactive teaching scenarios requiring short disruption.

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Future Work

- Scale experiments to 4 and 8 GB VM memory.
- Add delay and loss for more realistic WAN tests.
- Evaluate VXLAN for simpler network abstraction.
- Compare VMs with LXC for lower-overhead migration.

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Thank You For Your Attention!

Henry-Norbert Cocos, M.Sc
Frankfurt University of Applied Sciences
Room 1-230
☎ +49 69 1533-2699
✉ cocos@fb2.fra-uas.de
www.henrycocos.de

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Promotionszentrum Angewandte Informatik
HAW Hessen