OFFLINE-FIRST STRATEGIES - A NOVEL CONCEPT FOR THE MIGRATION OF WORKLOADS USING VIRTUAL MACHINES OMITING LIMITATIONS OF TRADITIONAL SERVICE DEPLOYMENT CONCEPTS

21st International Conference on Applied Computing 2024
26th - 28th October 2024
Zagreb, Croatia

Henry-Norbert Cocos
cocos@fb2.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/`

Use Case SKILL/VL

- Students access the SKILL/VL environment over network
- Students use the SKILL/VL environment for lectures
- Students use client hardware to interact with SKILL/VL infrastructure

Problem

The access and 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!

Thesis project 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

Dimension of Scaling

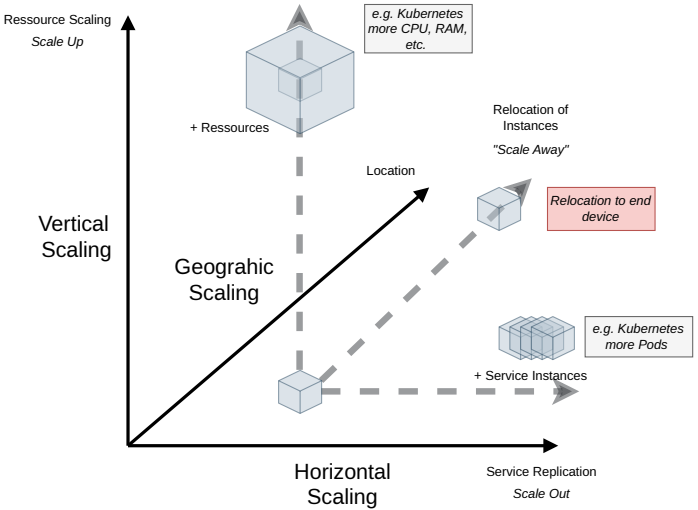

- **Horizontal scaling**
  $\Longrightarrow$ ***Scale out***
  - Adding additional service instances
- **Vertical scaling**
  $\Longrightarrow$ ***Scale up***
  - Adding additional resources to the service instance
- **Geographic scaling**
  $\Longrightarrow$ ***Scale away***
  - Relocating the service instance geographically closer to end user

Research questions

1. **How can resources on clients be used to save cloud resources and consequently bring applications closer to the end user?**

   Using virtualization technologies for compute, storage and networking and extending its base of operation is crucial!

2. **When does the migration to clients make sense and how does it contribute in reducing the load on the core service?**

   Whenever the network bandwidth limits the interaction with the service or the service needs offline operation capability!

3. **How does outsourcing services to the client affect the quality of services?**

   It affects the operation of the service depending on the resource consumption of the service!

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Field experiment – SKILL/VL

- 24 individual workstations for students.
- The resources of the lab are limited!
- Interesting application scenario: **Expansion of physical resources with virtual, logical resources!**

Lab 1-237

Computer networks laboratory in room 1-237 with 24 workstation set up for experiments. Each machine has multiple NICs (Network Interface Cards). Designed for practical application scenarios in computer networks.

Field experiment – SKILL/VL

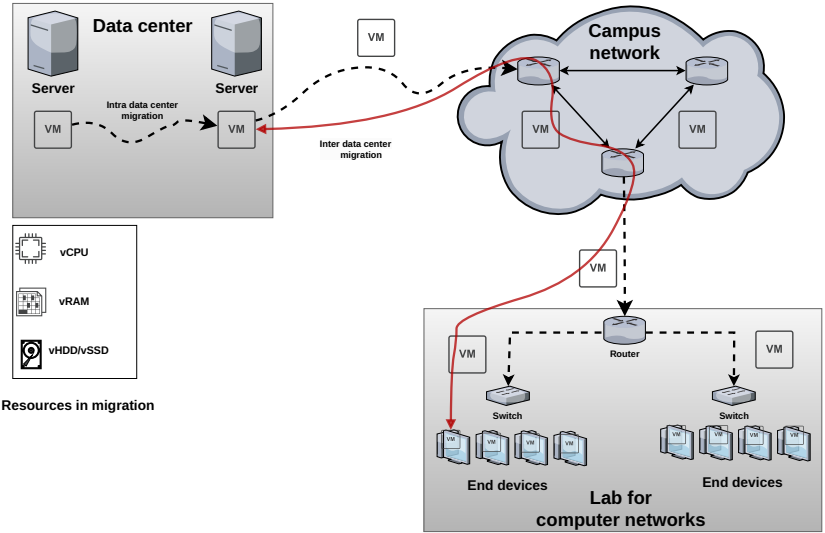

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Conducting Lab Experiments

Experiments shall investigate the following characteristics:

- Network performance over WAN
- Individual performance of applications
- Overall performance of service
- Applicability of Offline-First Strategy

Applicability of Offline-First Strategy

The main question of the service migration over WAN is the applicability. When does it make sense to migrate services over WAN?

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Lab experiment setup – VM migration

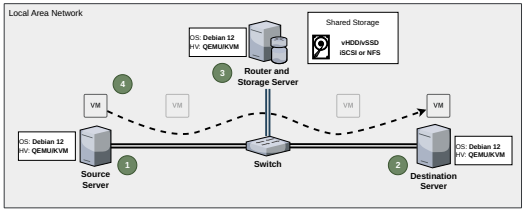

Figure: Lab experiment VM migration over LAN

- 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 Debian 12.

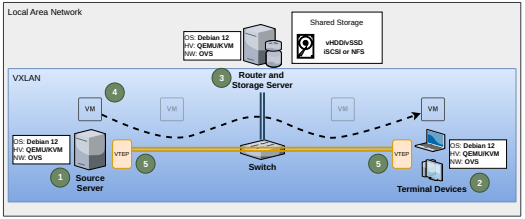

Figure: Lab experiment VM migration over WAN

- Lab experiments to compare VXLAN to LAN setup.
- Well defined environment for the experiments.
- Environment set up with QEMU/KVM, Open vSwitch and Debian 12.

Variables and key data of the experiments (1/3)

Table: Independent variable host[1]

Independent variable	Magnitude
Hardware type	x86-Prozessor Intel Core i7-10700K 8 Cores
CPU frequency	3.80GHz
RAM size	64 GB
HDD capacity	1 TB SSD
OS	Debian 12 Kernelversion 6.1.0-18-amd64
Hypervisor	KVM/QEMU Version 7.2.9
Storage	iSCSI or NFS Version 2.6.4
NIC	Ethernet – 1 GBit/s

[1] The values for the hosts are taken from laboratory 1-237.

Variables and key data of the experiments (2/3)

Table: Independent variables virtual environment

Independent variable virtual machine	**Magnitude**
VM-OS	Debian 12
vCPU	x86-Processor
vRAM	2GB
vHDD	10GB
vNIC	Intel E1000
Application	Testapp[a]

[a] The test applications are presented in the following slide!

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Variables and key data of the experiments (3/3)

Table: Test applications

Application	Implementation	Unit
Downtime Measurement	Python Script	Milliseconds
Monte Carlo Pi Estimation	Python Script	Milliseconds
CPU Load generation	Shell Script (`stress-ng`)	% CPU is busy
RAM Load generation	Shell Script (`stress-ng`)	% memory utilization

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Experiments

The following experiments were conduced:

- Generation of CPU load (0%, 25%, 50%, 80%, 100%)
- Generation of RAM load (0%, 25%, 50%, 80%)
- Simulation of a client-server application (π estimation)
- The WAN was simulated using `netem`.

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Experimental results migration time – 1 Gbit/s (1/5)

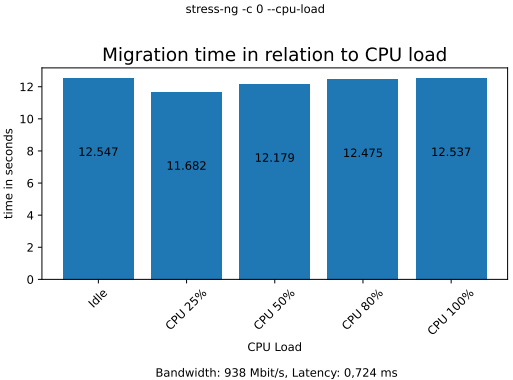

(a) Variation of CPU load (1 Gbit/s)

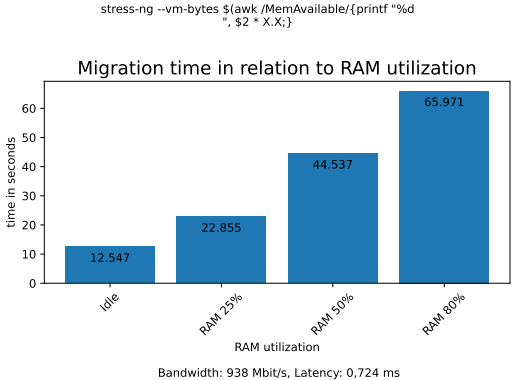

(b) Variation of RAM utilization (1 Gbit/s)

Figure: Migration time bandwidth 1 Gbit/s

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Experimental results migration time – 37 Mbit/s (2/5)

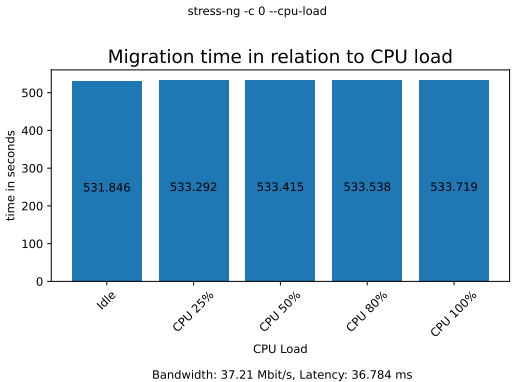

(a) Variation of CPU load (37 Mbit/s)

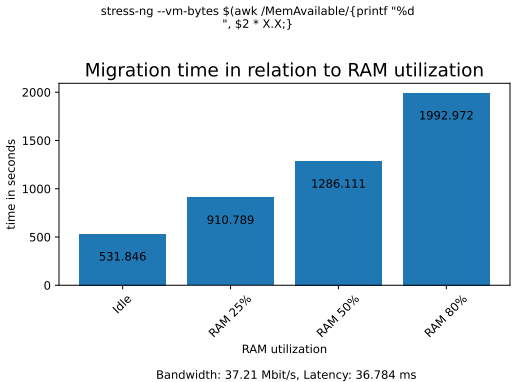

(b) Variation of RAM utilization (37 Mbit/s)

Figure: Migration time bandwidth 37Mbit/s

Experimental results response time – π (1 Gbit/s) (3/5)

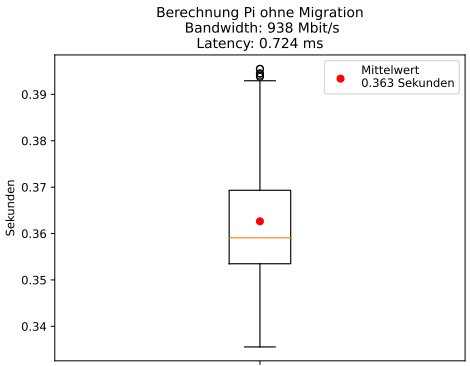

(a) Estimation of π (1 Gbit/s) – **No Migration**

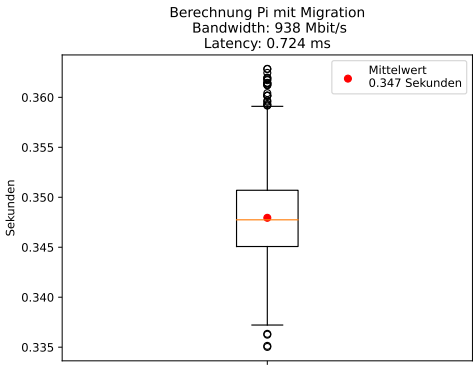

(b) Estimation of π (1 Gbit/s) – **With Migration**

Figure: Response times bandwidth 1 Gbit/s

Experimental results response time – π (37 Mbit/s) (4/5)

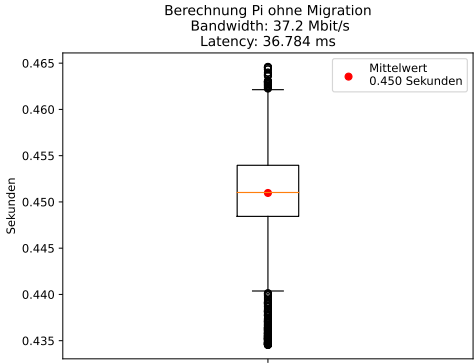

(a) Estimation of π (37 Mbit/s) – **No Migration**

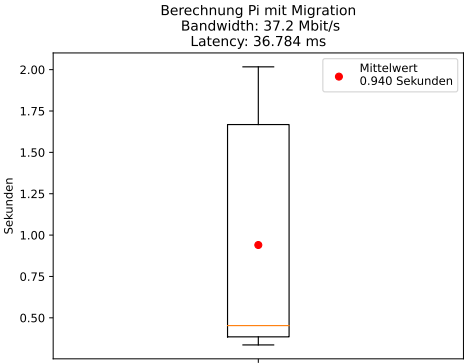

(b) Estimation of π (37 Mbit/s) – **With Migration**

Figure: Response times bandwidth 37Mbit/s

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Experimental results (5/5)

The results indicate that...

- ...CPU-intensive tasks have a minor impact on the migration time! (Plots 3a and 4a)
- ...Memory-intensive tasks yield a huge impact on the migration time! (Plots 3b and 4b)
- ...the response times for applications are heavily impacted by migration and bandwidth! (Slides 18 and 19)
- ...the downtime of the VMs are constantly around 2-4 milliseconds! (*no plot provided*)
- ...the bandwidth has a huge impact on the migration time! (that is common sense!)

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Conclusion

- Proposition of a Offline-First-Strategy for distributed and virtualized resources.
- Experiments show an impact of memory not CPU intensive workload.
- Live migration of VMs need the same hypervisor (in version as well!) on the server and client.
- Live migration of VMs between different CPU-architectures (e.g. Intel and an AMD) is not possible!
- A shared storage for the VMs virtual hard disks is needed at all times!

FRANKFURT UNIVERSITY OF APPLIED SCIENCES

Next Steps

- Measurements in VXLAN setup and comparison
- Setup of field experiment from slide 9
- Analysis of result in field experiment and comparison to lab experiments
- Interpretation of results and impact on setup

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