METRO|NOM

OpenStack at large retail enterprise - Boon or bane?

14.11.2018 – OpenStack Summit Berlin

About

About me

Thomas Lunkwitz (0x29 y)

- Product Owner for „Compute Cloud“ at METRONOM
  - Working at METRONOM and its precursor for 16 years
  - Career started as network specialist for WAN and LAN
  - Started with OpenStack as project in 2014
  - Product owner for OpenStack since 2017

- Volunteer at German Federal Agency of Technical Relief (THW)

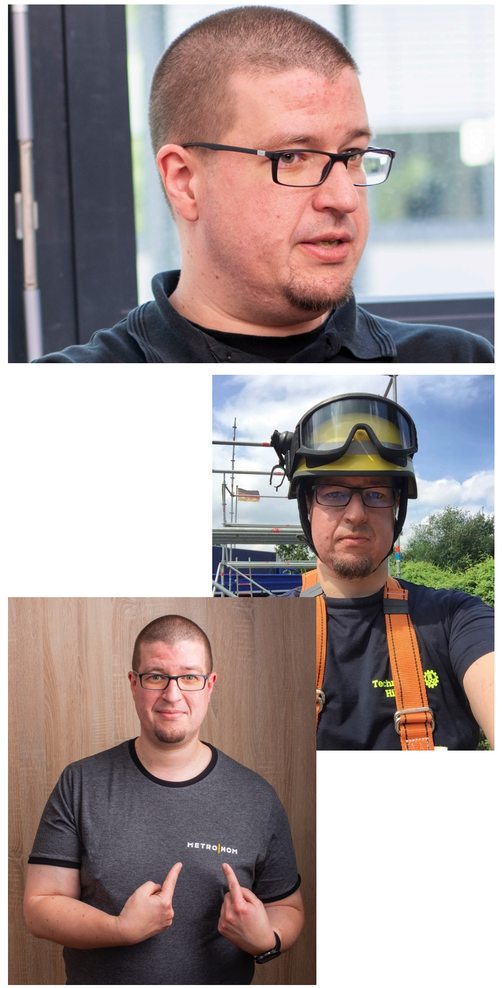

“I am a METRONOMIAN”

METRONOM.com

About METRONOM

We set the pace in food and technology.

- METRONOM is the tech unit of METRO, a leading international wholesale and food specialist company
- Around 2.000 people working for METRONOM mainly in Berlin, Duesseldorf, Hannover, Brasov and Bucharest
- Our Vision:

Our culture, services and digital solutions are unmatched in the global wholesale business. We revolutionize the entire industry – working with us is a privilege!

METRONOM.com

Boon and Bane – But why?

“Bane” because …

- Previous structure of METRO IT was totally vendor driven
  - Open source was used only for niche products
- We were not contributing to or working in communities
- We had a lot of very good teams in different departments
  - These departments were more like Silos
- We were not really acting as one big team
- The world and especially the IT is changing very fast
- IT is becoming consumable like electricity

- **We are changing our culture and mindset**

“Boon” because …

- Cloud is becoming more and more important in the IT market
- In future we will see many more features and software which are only available in the cloud
- Hybrid cloud approach will be the future
  - Using a public cloud environment brings flexibility and additional features
  - Not all data can or should be stored in the public cloud
  - Using an internal cloud based IaaS provides:
    - Safety
    - Alternative to public cloud providers
- OpenStack and cloud forces and supports us to change our culture and mindset

- **OpenStack is open and provides the open infrastructure we need**

METRO|NOM.com

Advantages of private cloud

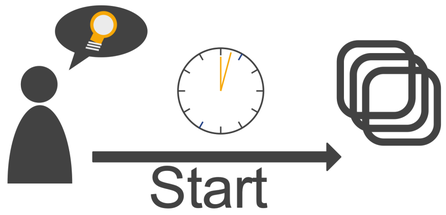

- Everything is consumable as self-service – easy to start

- Fully integrated into corporate network

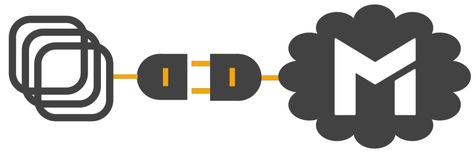

- Data protection and data privacy are guaranteed (on-premise)

- Cost effective (non-profit cost level) – pay-per-use

Consume infrastructure

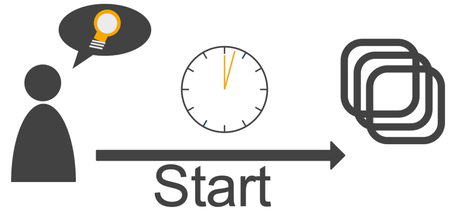

- Everything is consumable as self-service – easy to start
  - As of today a tenant has to be requested by a webform
  - We will create the project and the users in OpenStack
  - Once the project and users are created, everything else is consumable as self-service
- We will work on a shop system, so that even a project can be created automatically
  - Idea is to implement a kind of web shop for infrastructure products
  - The shop system will be based on open source software
  - PoC is not started yet

- **From an idea to start in less then five minutes**

Fully integrated

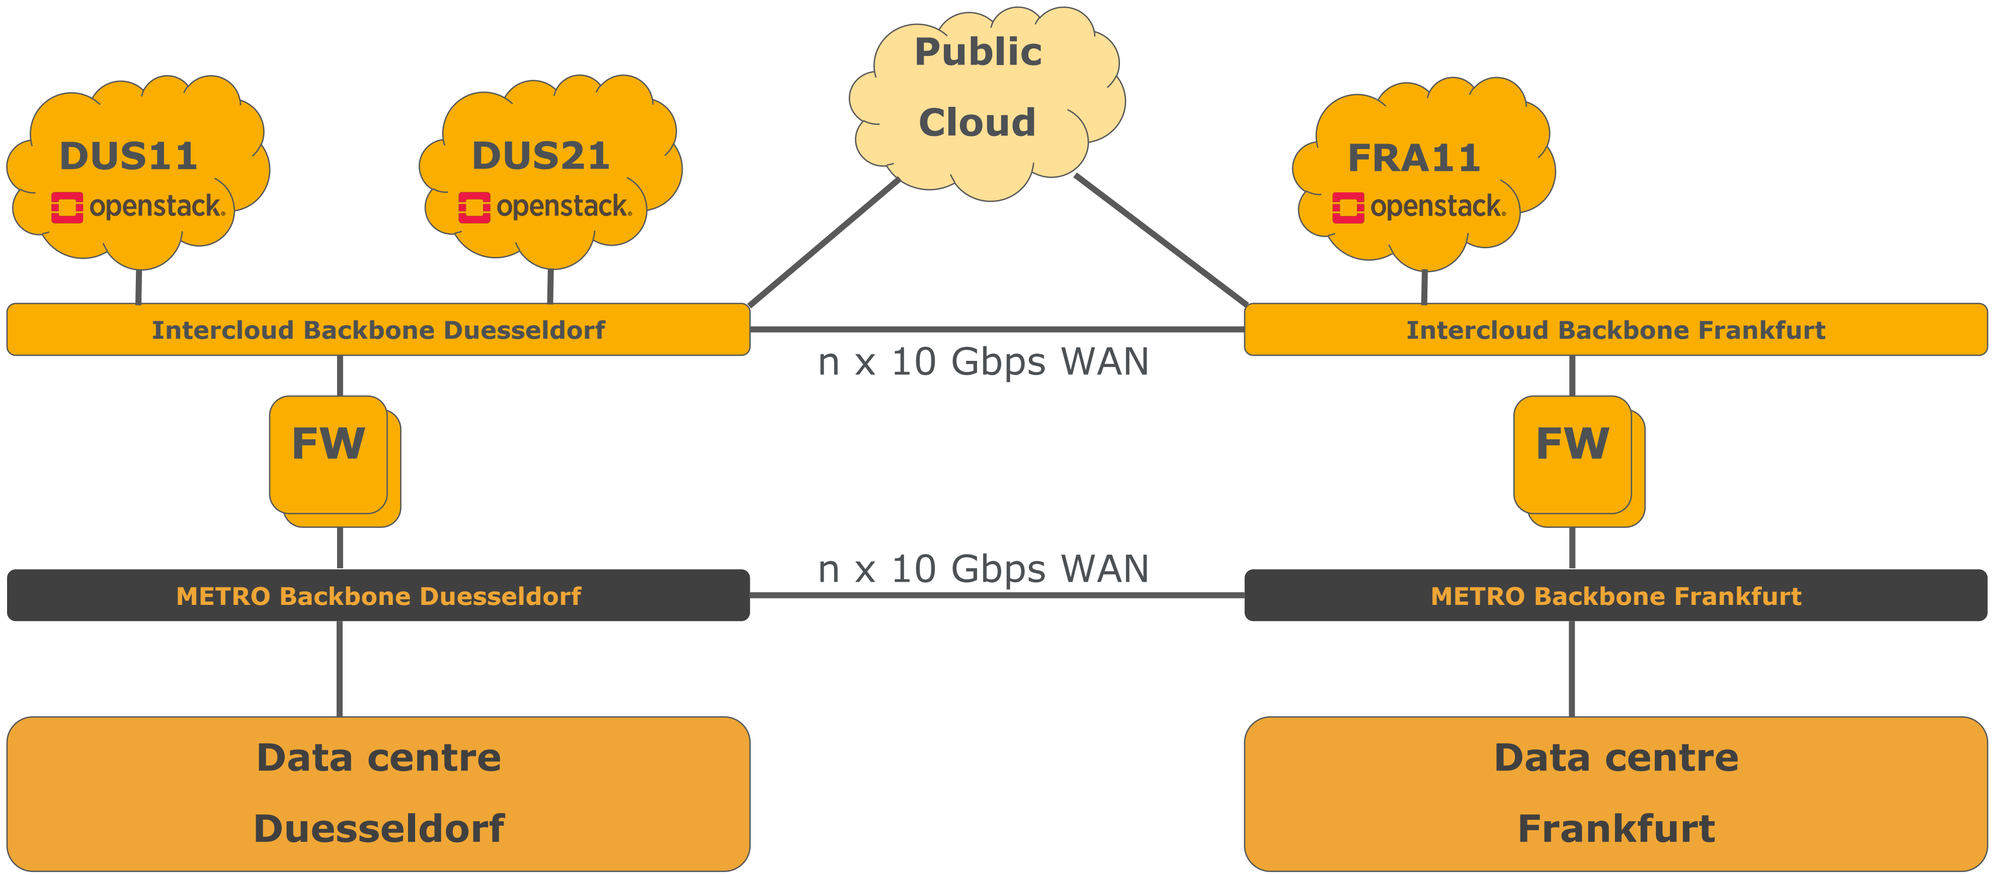

- Fully redundant and highly scalable infrastructure

Data protection and privacy

- Private cloud for METRO only
- All hardware is located in METRONOM or METRONOM managed data centres
  - Full physical access control by METRONOM
  - No issues in terms of GDPR and external service providers
- Data encryption is strongly recommended to our customers
- We provide block and object storage
  - Everything is provided in the context of the project
  - We do not provide any kind of shared storage
  - Local storage volumes will be overwritten when deleted

- **In future we will encrypt all data by default on hardware level**

Cost effectiveness

- METRONOM is a non-profit Tech of METRO
  - We charge only for our costs
- We reduce hardware costs by standardising our server flavours
- We charge on a pay-per-use model
  - Smallest unit is one hour
- We charge for
  - Instances
  - Storage
  - Public floating IPs

- **Full cost transparency for our customers by monthly reports**

METRONOM.com

Our OpenStack Deployments

Three data centres in Germany and China

Our deployments

- We built 6 OpenStack fully independent environments
  - 2 based on Suse OpenStack Cloud 5 (Juno)
    - Those were the first productive environments from 2015
    - The will be decommissioned by the end of 2018
  - 4 based on OpenStack Ansible (Newton)
    - 2 in Duesseldorf
    - 1 in Frankfurt
    - 1 in Shanghai
  - Options for two more
    - 1 in Moscow
    - 1 in Frankfurt

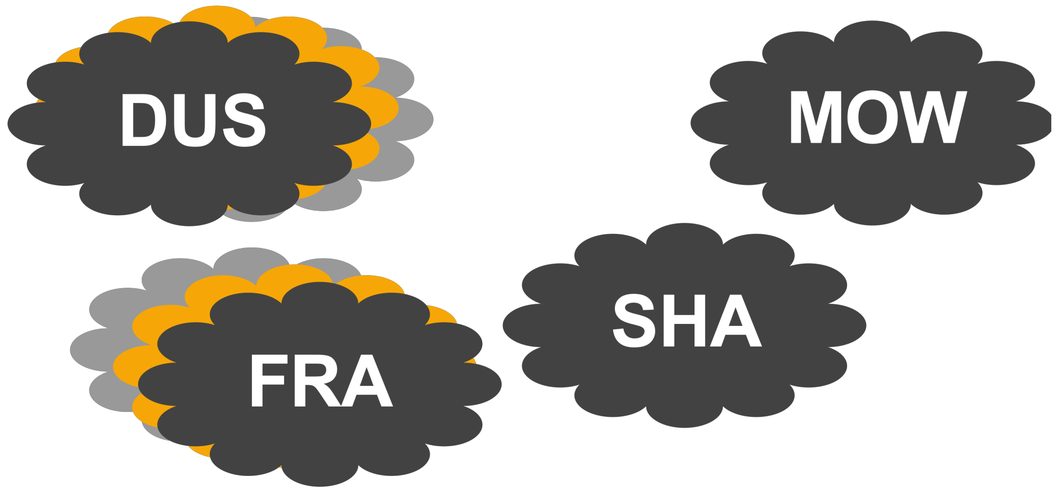

Design your application landscape

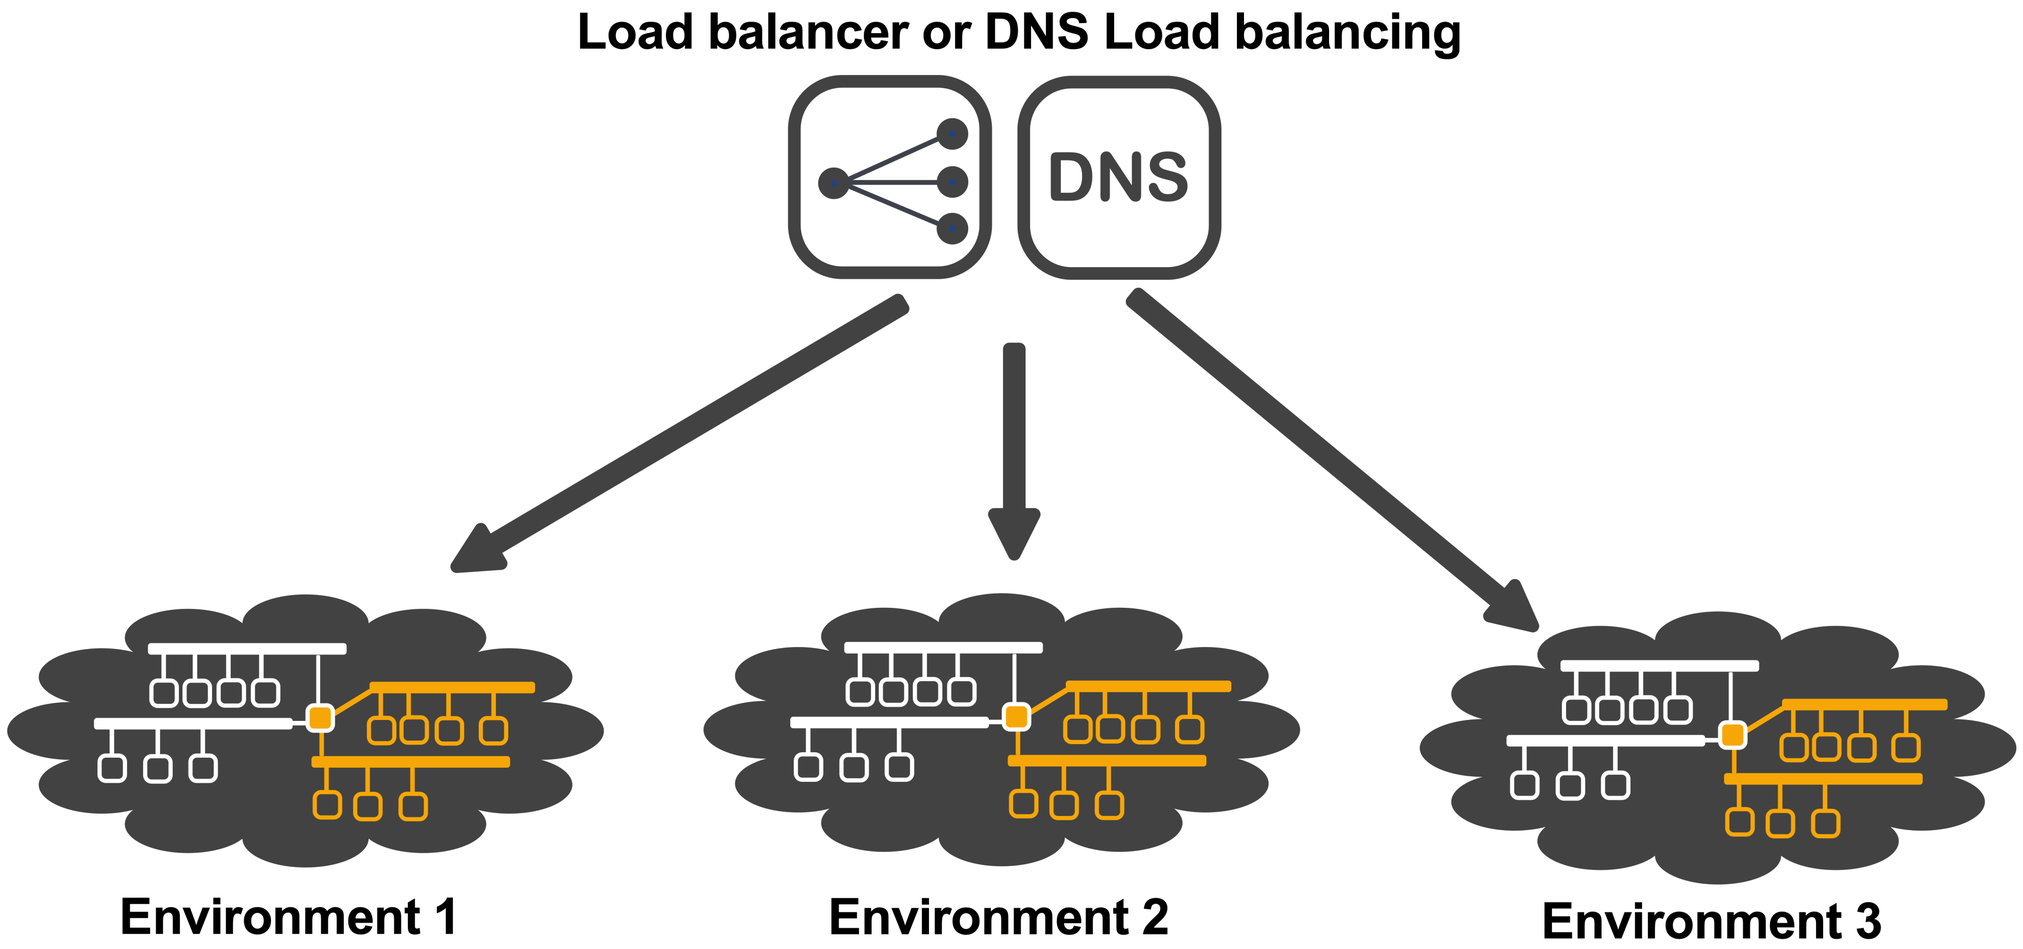

Design your application landscape

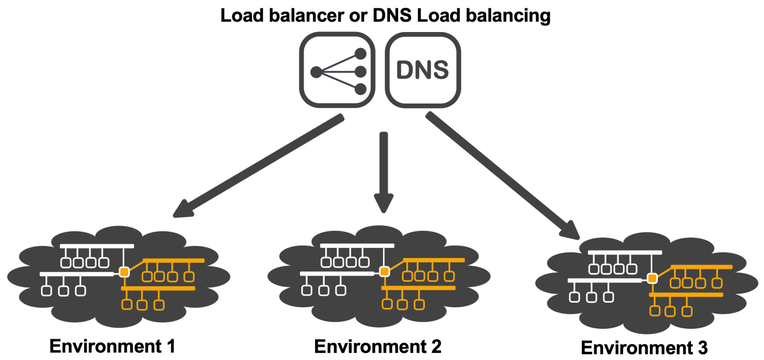

- Prerequisites for our customers:
  - Be prepared for any kind of failure
    - A VM could fail and even be lost at any time
    - Even a compute host could fail and we probably will not repair it
  - It is strongly recommended to spread the application across multiple environments if it is needed in HA mode
- We “educate” our colleagues and customers to handle failures in the infrastructure more efficiently
- We offer support and consultancy **but we do not implement anything for them**

Software we use

- Ubuntu Server 16.04 LTS
- OpenStack Ansible (Newton)
- CEPH-Ansible (Jewel)

- OpenStack modules
  - Keystone
  - Nova
  - Neutron
  - Horizon
  - Cinder
  - Swift
  - Ceilometer (will be replaced)
  - Glance
  - Heat

Our architecture

- Everything is distributed over three racks
  - Per Rack
    - One controller
    - 1 to n compute hosts
    - One monitor node (Rados GW)
    - 1 to n storage hosts
    - Max. 32 nodes in total
    - Free spare RU for TOR Switches
- Actual end-of-the-row switch concept
  - Might change to TOR with Spine-Leaf architecture

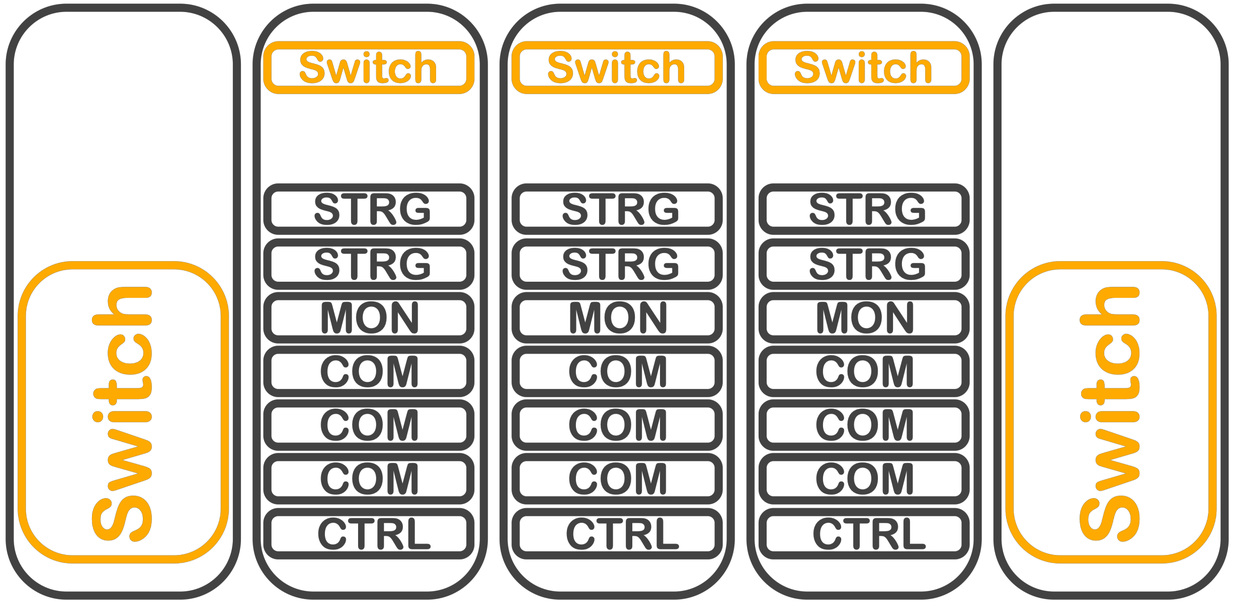

Our network design

- Per controller / monitor / compute host
  - 1 x 1 Gbps Ethernet copper - Hardware management (iLo, CIMC, iBMC)
  - 2 x 1 Gbps Ethernet copper – OpenStack management (active/backup failover mode)
  - 2 x 10 Gbps Ethernet fibre – user traffic (active/active bond)
  - VxLAN with DVR

- Per storage host
  - 1 x 1 Gbps Ethernet copper - Hardware management (iLo, CIMC, iBMC)
  - 2 x 1 Gbps Ethernet copper – OpenStack management (active/backup failover mode)
  - 2 x 10 Gbps Ethernet fibre – storage traffic (active/active bond)
  - 2 x 10 Gbps Ethernet fibre – replication traffic (active/active bond)

Available resources and running VMs

Env.	Comp. Hosts	CPU (1:4)*			RAM (1:1)*	Storage (1:1)*			*VMs*
		Physical	Hyperthr.	vCPU		Local (in compute host)	CEPH HDD	CEPH SSD	
DUS10**	29	550	1.100	4.400	~ 14 TB	~ 60 TB	~ 51 TB	- / -	*266*
DUS11	18	360	720	2.880	~ 9 TB	~ 26 TB	~ 78 TB	~ 15 TB	*179*
DUS21	60	2.220	4.440	17.760	~ 30 TB	~ 160 TB	~ 131 TB	~ 41 TB	*1.109*
FFM10**	14	280	560	2.240	~ 6 TB	~ 30 TB	~ 46 TB	- / -	*176*
FRA11	33	660	1.320	5.280	~ 17 TB	~ 65 TB	~ 47 TB	~ 15 TB	*264*
SHA11	39	1.320	2.640	10.560	~ 20 TB	~ 129 TB	~ 35 TB	~ 27 TB	*523*
Total	**193**	**5.390**	**10.780**	**43.210**	**~ 96 TB**	**~ 470 TB**	**~ 388 TB**	**~ 98 TB**	***2.517***

*): Over commitment *physical:virtual*

**): Discontinued by end of 2018

METRO|NOM.com

Challenges

- Assure installation of an environment
  - Mirror all resources needed for an installation
  - Being able to install even without Internet

- Building a CI/CD pipeline
  - Assure that all environments are on the same version with same settings
  - Optimize monitoring and support process

- Major upgrade to Rocky or newer

It’s all about people!

Culture, people and mindset

"People don't buy what you do, they buy why you do it." *Simon Sinek*

- Success is not only a question of technology, it is rather a question of **passion** and **people**
- We are hiring for **potential** and **personality**
- We work in self-organized teams

We drive innovation, not just follow!

Team, team, team

```
---
# Join the team
# METRONOM GmbH

- hosts: you
  pre_tasks:
  - name: Check Prerequisites
    assert:
      that:
      - NEW_JOB_WANTED == "True"

  - name: Set facts
    set_fact:
      flexible_working_hours: "True”
      employee_benefits: "True”
      teamwork: "True”

  tasks:
  - name: Check applicant soft skills
    soft_skills:
      skill: "{{ item }}”
      state: present
    with_items:
    - Teamplayer
    - Flexibility
    - Humor
    register: soft_skills_match

- name: Check applicant Tech skills
  technical_skills:
    skill: "{{ item }}”
    match_type: regex
    state: present
  with_items:
  - Linux
  - Ansible|Puppet|Chef
  - Python|GO
  - Network|Docker|Kubernetes
  register: technical_skills_match

- name: Prepare and send job application
  uri:
    target: "{{ WWW.METRONOM.COM/APPLY }}”
  when: soft_skills_match == True and
  technical_skills_match == True
```

METRONOM.com

Contact

Thomas Lunkwitz

tlunkw
tlunkw
thomas-lunkwitz
@ thomas.lunkwitz@metronom.com

Drop me a line!
I will respond within 48 hours.

METRONOM.com

Thank you for your attention!

Thomas Lunkwitz