

# University cloud and virtualization



**Tomáš Sapák**

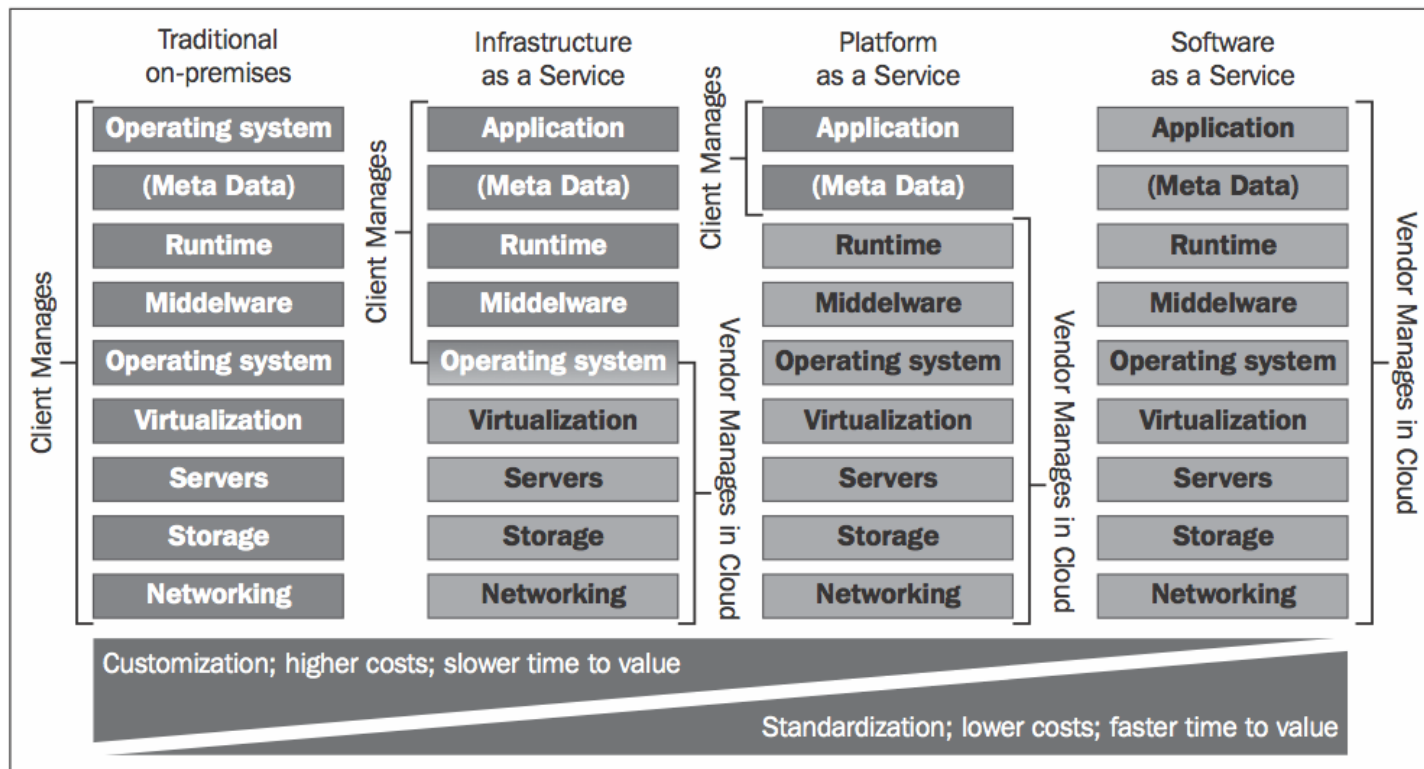
**sapakt@ics.muni.cz**

# ■ Cloud in theory

## What is cloud?

- **Private vs public**
- **Infrastructure as a service (IaaS)** - Amazon Elastic Compute Cloud, Google Compute Engine, Microsoft Azure, Open Stack, Open Nebula
- **Platform as a service (PaaS)** - Google App Engine (PHP, Python, Java, Go), Amazon Elastic Beanstalk (Ruby, PHP, Python, .NET, JAVA, JavaScript), Microsoft Azure Websites (PHP, Python, .NET, JavaScript)
- **Software as a service (SaaS)** - mail, social network, on-line office, file services, image libraries, music libraries, video libraries, web sites ...

# ■ NIST service model





## ■ **IaaS at Masaryk university**

- **Infrastructure for production applications**
  - **stable, highly available environment for production level applications**
- **CERIT-SC cloud**
  - **scientific computations with intensive CPU workload**
  - **covered in session by Aleš Křenek and Tomáš Rebok**

# Infrastructure for production applications

- virtualization (and consolidation) of physical servers since 2006 on MU (VMware)
- most systems are nowadays virtualized
- Pros:
  - easier management
  - space and power consumption efficiency
  - easy migration to different hw, location
  - easier disaster recovery
- Cons:
  - hw failures have worse consequences (Many eggs in one basket)

## ■ Production infrastructure

- approximately 360 VMs
- critical systems (DNS, DHCP, E-mail, file servers, databases, webservers ...)
- downtimes are not tolerable
- quality and redundant infrastructure is essential

## ■ Components

- Physical servers
- Network
- SAN and storage
- Datacenters
- Virtualization layer

## ■ Physical servers

- core of virtual infrastructure
- diskless intel servers - virtual machines are stored on disk arrays
- 9 hosts in production, 2 for testing, 2 in secondary location, 2 in tertiary location
- evolution in form of number of cores (4 per CPU in 2009 vs 10 per CPU in 2014) and amount of RAM (48 GB per host in 2009 vs 384 GB per host in 2014)



## ■ Network

- 2x 10G Ethernet card per host
- each card is connected to different router
- active/active - all hosts can use both interfaces for better performance (single VM can use only one)
- 46 VLANs used by VMs
- Primary and secondary datacenter part of same LAN

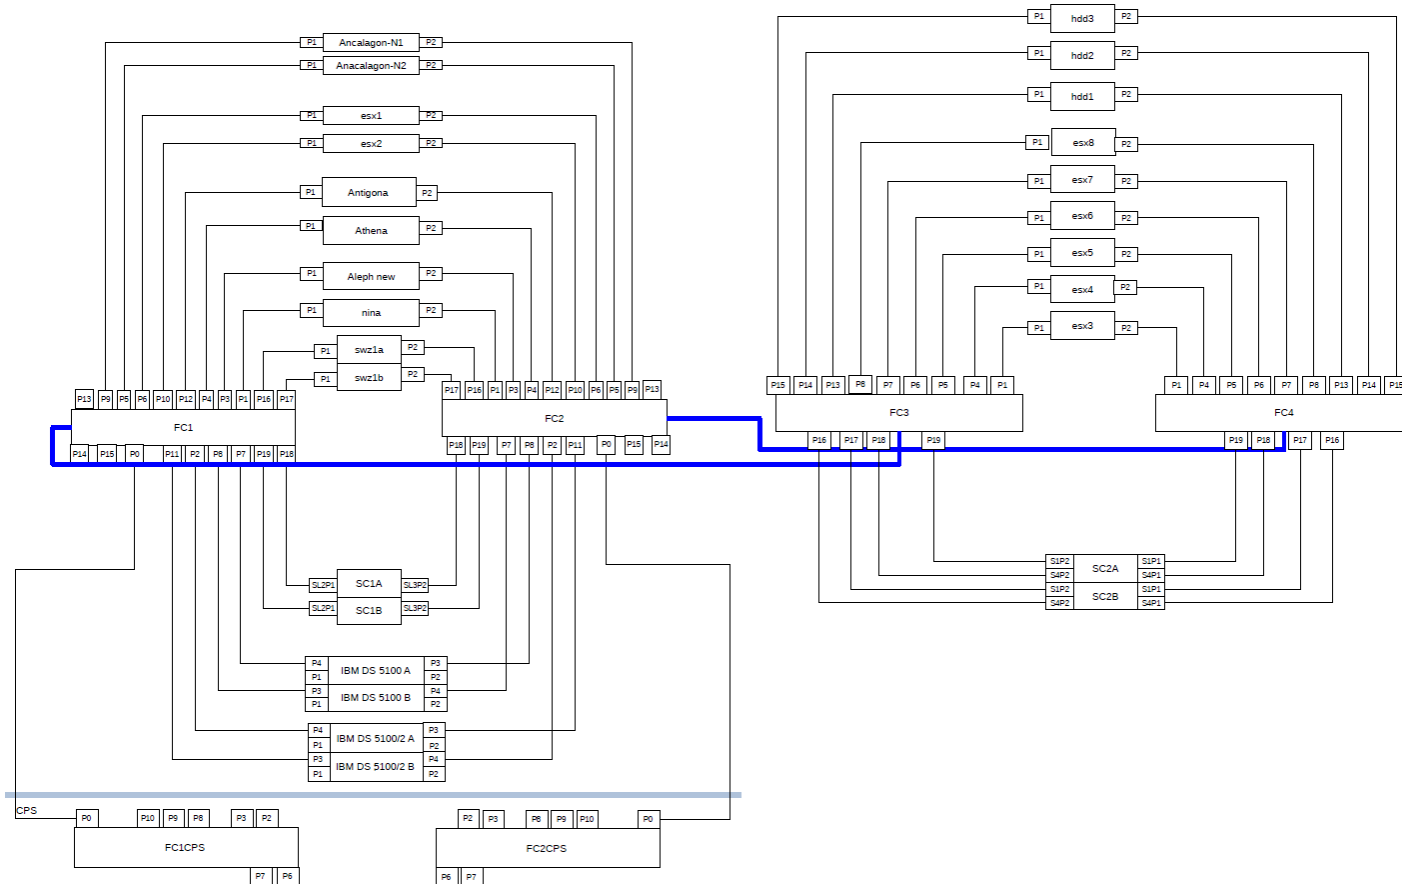
## ■ SAN

- storage area network
- dedicated network for storage traffic (between storage arrays and hosts)
- 8G fibre channel
- fully redundant
- 2x fc switch, 2x hba port per host, per storage array controller
- Connects primary and secondary datacenter

# SAN



Botanická



## ■ Storage

- all VMs are stored on disk arrays
- 3 mid-range arrays for VMs, several low-ends for large data and backups
- currently there are 2 tiered arrays, one of them equipped with SSDs

# ■ Automated Storage Tiering

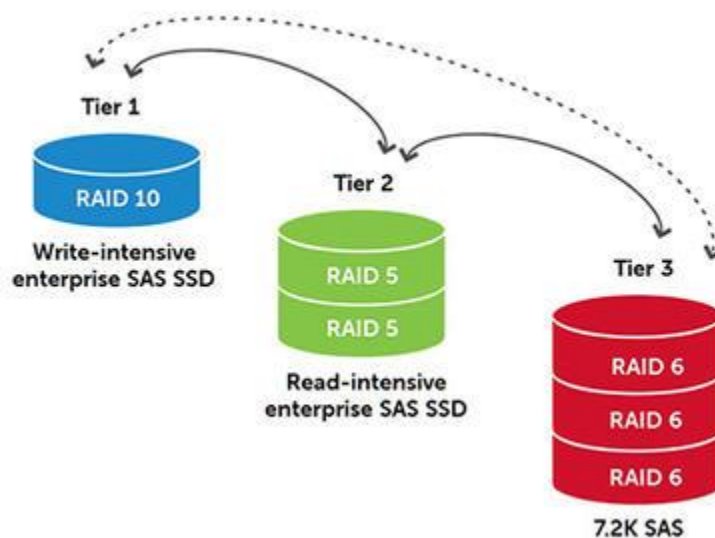
**Traditional storage arrays problems:**

- **Expensive and low size of fast drives (SSDs, FC, SAS)**
- **Low performance and long rebuild times with high capacity SATA (NL-SAS) drives**

# ■ Automated Storage Tiering

- **Storage array consists of two storage layers (tiers) with different types of disk drives (slow and large, fast and small)**
- **Written data are partitioned and stored in the fast tier**
- **For a few days usage of data is evaluated and data with little usage are sent to lower tier**
- **For even better performance and space utilization some vendors use different RAID types for reading and writing**

# ■ Automated Storage Tiering




# ■ Example





# ■ Storage of tomorrow

- All flash array – online compression and deduplication
- Virtual SAN





## ■ **Datacenters**

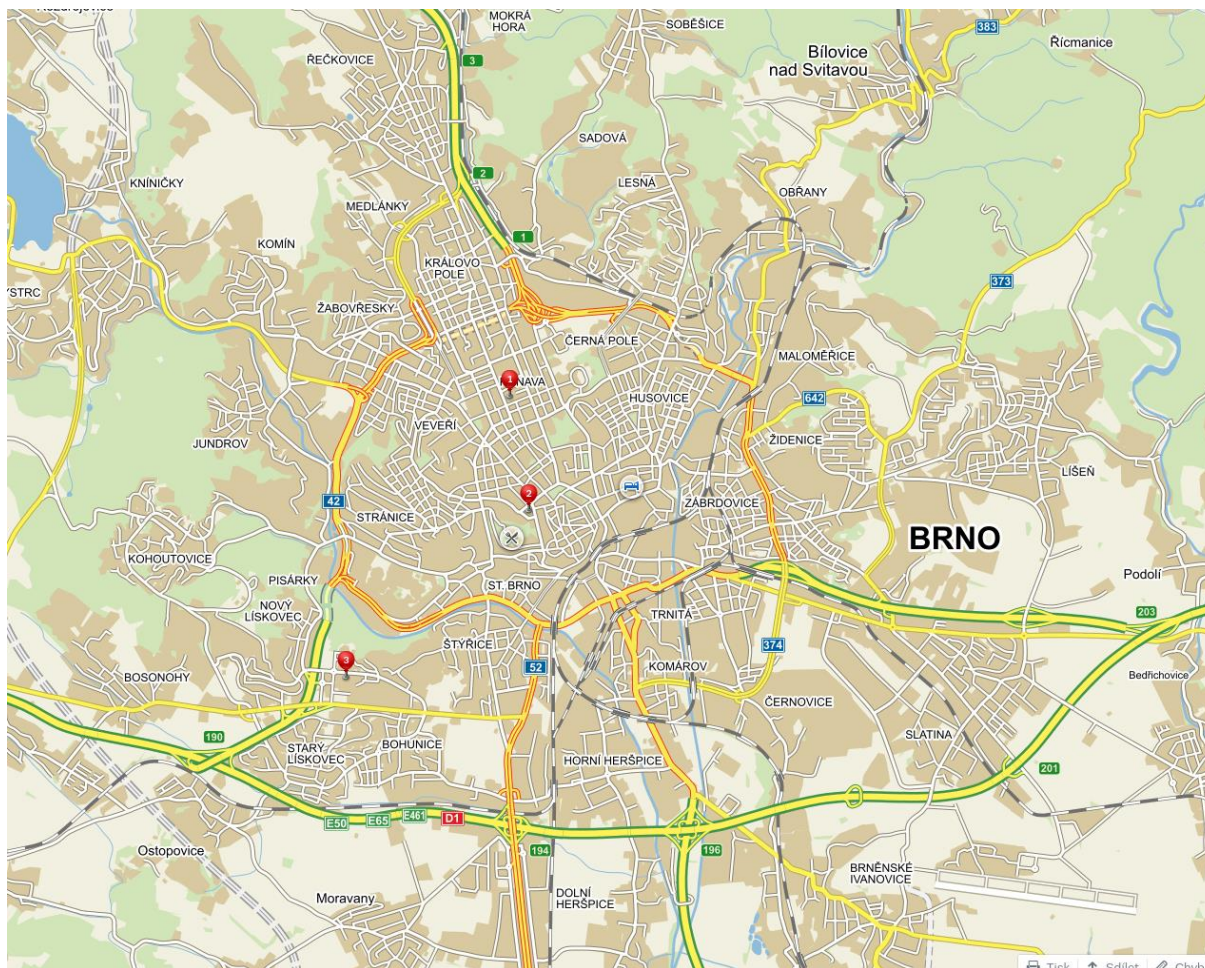
**All datacenters are fully equipped with cooling, UPS. Primary and tertiary locations have electric generator for longer blackouts.**

**Primary location – ICS/FI building, Botanická.**

**Secondary location – University Computer Center, Komenského, part of same LAN and SAN as primary location**

**Tertiary location – University Campus Bohunice, completely separated from primary and secondary locations**

# ■ Datacenters



## ■ Virtualization layer

- VMware vSphere 5.5
- VMware virtual infrastructure consists of two fundamental elements – vCenter and ESXi
- ESXi is hypervisor – software that creates and runs virtual machines
- vCenter is central point for management of hypervisors, virtual machines and users

## ■ **Virtualization layer**

- **Primary and secondary datacenters managed by single vcenter**
- **13 hypervisors form 4 clusters:**
  - **Production cluster**
  - **Production cluster with special demands**
  - **Cluster for testing**
  - **Cluster on secondary location**

## ■ Features

### **Distributed resource scheduler (DRS):**

- **Automatically chooses less utilized host in the cluster for newly started VM**
- **Automatically migrates VMs among nodes in the cluster to balance load evenly**
- **Automatically migrates off VMs from hypervisor for maintenance**

### **High availability (HA):**

- **In case of failure of hypervisor automatically powers on VMs on healthy hosts in the cluster**

# ■ Features

## Data Protection:

- Image-level virtual machine backup (agentless)
- A few click restore
- Deduplication
- Changed Block Tracking Backup/Restore

## ■ Features

### Replication:

- image-level replication
- **Minimum 15 minutes RPO (Recovery Point Objective)**
- **Not suitable for every application**
- **Eliminates single point of failure of storage array (all other components are redundant)**



# ■ Example

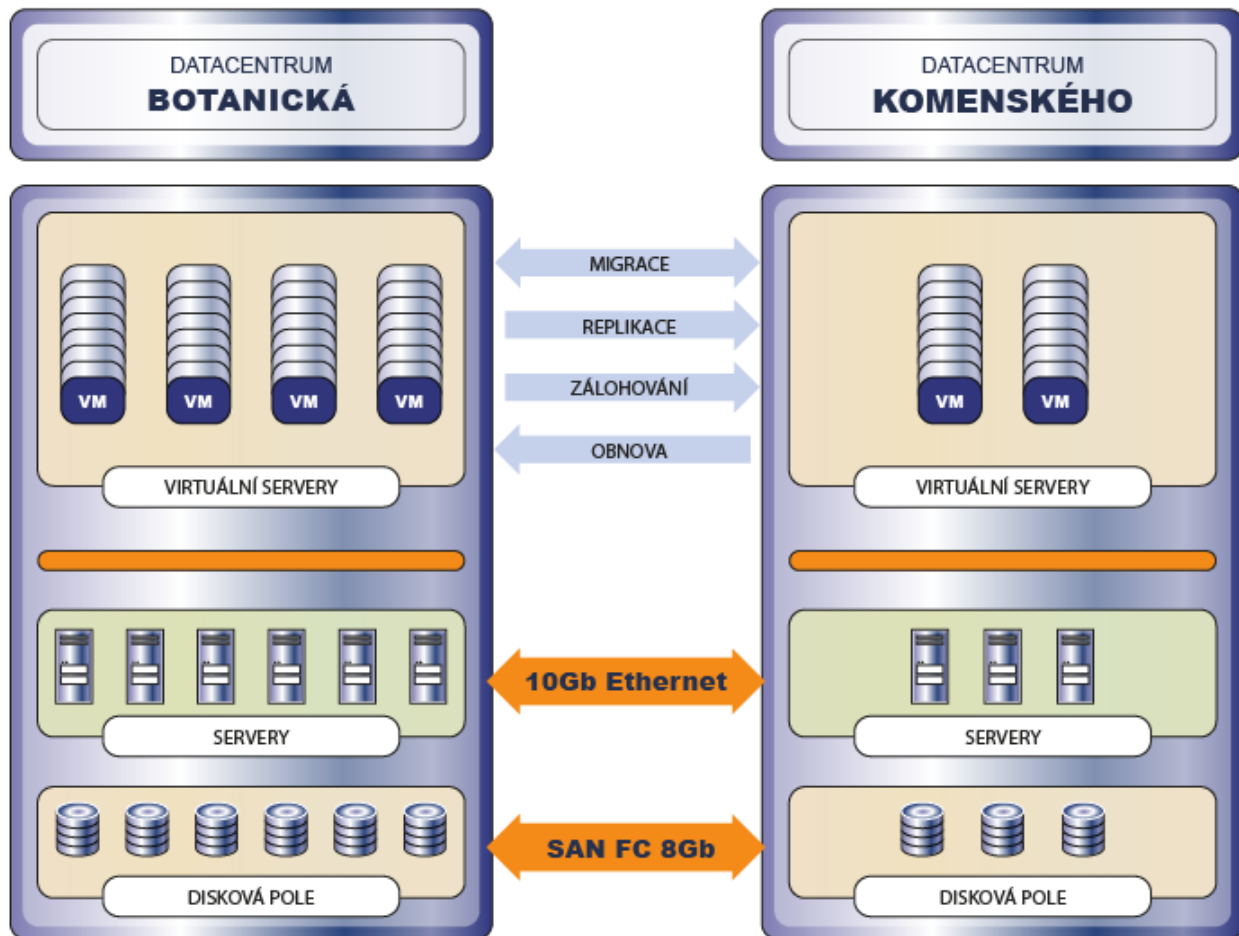


## ■ Features


### Fault Tolerance:

- VM runs on two different hosts in the same time
- In the case of host failure, failover is automatically triggered – zero downtime

# Overview



**Thank you for your attention**

A decorative graphic at the bottom of the slide consists of several wavy lines. A prominent thick red wave is at the bottom, with a grey wave above it. Above these are several thin grey wavy lines. Three light blue dots are connected by thin lines, forming a path across the upper part of the graphic.

**Tomáš Sapák**  
**sapakt@ics.muni.cz**