Tampereen yliopisto
Tampere University

Lecture 2
Cloud for Software Engineers
Technical

Tampereen yliopisto
Tampere University

Introduction to ”cloud”

Tampereen yliopisto
Tampere University

Agenda

- History and motivations
- Definitions
- Categories: IaaS, SaaS, PaaS
- Business and SW development perpective
- Technical perspectives in Operating Systems and Cloud Platforms courses

Tampereen yliopisto
Tampere University

Our perception?

Different interpretations

- It is somewhere in the cloud and controlled by a ”big brother”
- We do not need to buy our computing hardware anymore
- Cool technology to play with
- Accessible from anywhere
- Other

Tampereen yliopisto
Tampere University

Where did it start ?

https://www.dataversity.net/brief-history-cloud-computing

- ” In its early stages, the Cloud was used to express the empty space between the end user and the provider. In 1997, Professor Ramnath Chellapa of Emory University defined Cloud Computing as the new

’computing paradigm, where the boundaries of computing will be determined by economic rationale, rather than technical limits alone.’ ”

- In 1999, Salesforce became a popular example of using Cloud Computing successfully. They used it to pioneer the idea of using the Internet to deliver software programs to the end users. The program (or application) could be accessed and downloaded by anyone with Internet access. Businesses could purchase the software in an on-demand, cost-effective manner, without leaving the office.

Tampereen yliopisto
Tampere University

Case salesforce

- From https://www.salesforce.com/:

WHAT IS SALESFORCE?

Salesforce is a suite of web-based CRM applications that help you find, win, and keep customers. Learn more.

Tampereen yliopisto
Tampere University

https://www.computerworld.com/article/3427741/a-brief-history-of-salesforce-com.html

“The way the story goes is that Marc Benioff was floating in the sea just off Big Island in his beloved Hawaii during a sabbatical when he thought:

why can't buying software be as a simple as Amazon is for consumer goods?

This line of thinking eventually led to Benioff and a team of developers pioneering the software-as-a-service (SaaS) model by delivering its customer relationship management (CRM) software

over the internet on a per seat, per month payment plan, instead of deployed on-premise servers under a hefty licensing agreement.”

Tampereen yliopisto
Tampere University

Case Amazon

- Forerunner of web-based retail services.
- Used only 10% of their computing capacity capacity (which was commonplace at the time).
- The Cloud Computing Infrastructure Model gave them the flexibility to use their computer's capacity much more efficiently.
- Soon after, other large organizations followed their example.
- In 2006, Amazon launched Amazon Web Services, which offers online services to other websites, or clients.
- Another of Amazon Web Services' sites is the Elastic Compute Cloud (EC2), allowing individuals to rent virtual computers and use their own programs and applications.
- Now (2020!) selling computing capacity to others (AWS business) brings about half of the operating income (source: https://press.aboutamazon.com/news-releases/news-release-details/amazoncom-announces-first-quarter-sales-17-597-billion)

Tampereen yliopisto
Tampere University

Case Google

- Also in 2006, Google launched the Google Docs services.
- Based on two separate products,
  - Google Spreadsheets (acquired from 2Web Technologies, in 2005) and
  - Writely (Google purchased Writely)
- In 2007, IBM, Google, and several universities joined forces to develop a server farm for research projects needing both fast processors and huge data sets.
- 2007 was also the year when Netflix launched it’s streaming video service, using the Cloud, and provided support for the practice of “binge-watching.”

Tampereen yliopisto
Tampere University

Story continues

- IBM
  - Smart Cloud
  - IBM Cloud (formerly IBM Bluemix and IBM SoftLayer)
- Microsoft
  - Azure
  - Office 360
- Apple
  - ICloud

Tampereen yliopisto
Tampere University

Example SaaS: Adobe Creative Cloud

(http://www.paulpehrson.com/2011/04/11/adobes-new-software-as-a-service-model/)

Product	Full	Upgrade*	SAAS**	Months to justify initial investment***
Design Premium	$1899	$399	$95	20
Web Premium	$1799	$399	$89	20
Production Premium	$1699	$399	$85	20
Master Collection	$2599	$549	$129	20
Photoshop	$699	$199	$35	20
Illustrator	$599	$199	$29	20

Tampereen yliopisto
Tampere University

Towards definitions

Peter Mell; Timothy Grance (September 2011). The NIST Definition of Cloud Computing (Technical report). National Institute of Standards and Technology: U.S. Department of Commerce. doi:10.6028/NIST.SP.800-145. Special publication 800-145.

Tampereen yliopisto
Tampere University

Essential characteristics 1/2

- *On-demand self-service.* A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.
  - Unilaterally?
  - Without human interaction?
- *Broad network access.* Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).
  - What does the heterogeneous platforms mean in practice?
- *Rapid elasticity.* Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.
  - What if scaling is not automatic?

Tampereen yliopisto
Tampere University

Essential characteristics 2/2

- *Resource pooling.* The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.
  - Why is this essential?
- *Measured service.* Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.
  - Why?

Tampereen yliopisto
Tampere University

Service models

SaaS

PaaS

IaaS

Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision
processing, storage, networks, and other fundamental computing resources
where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications.

The consumer
does not manage or control the underlying cloud infrastructure but
has control over operating systems, storage, and deployed applications; and
possibly limited control of select networking components (e.g., host firewalls).

Examples?

Tampereen yliopisto
Tampere University

Service models

SaaS

PaaS

IaaS

Platform as a Service (PaaS). The capability provided to the consumer is to

- deploy onto the cloud infrastructure consumer-created or acquired applications
- created using programming languages, libraries, services, and tools supported by the provider.

The consumer

- does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage,
- but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Examples?

Tampereen yliopisto
Tampere University

Service models

SaaS

PaaS

IaaS

Software as a Service (SaaS). The capability provided to the consumer is to

use the provider's applications running on a cloud infrastructure.

The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface.

The consumer

does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Examples?

Tampereen yliopisto
Tampere University

Deployment models

- *Private cloud.* The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
- *Community cloud.* The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
- *Public cloud.* The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
- *Hybrid cloud*. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

Virtualization – what and why?

Tampereen yliopisto
Tampere University

Use case 1: run ”foreign” software

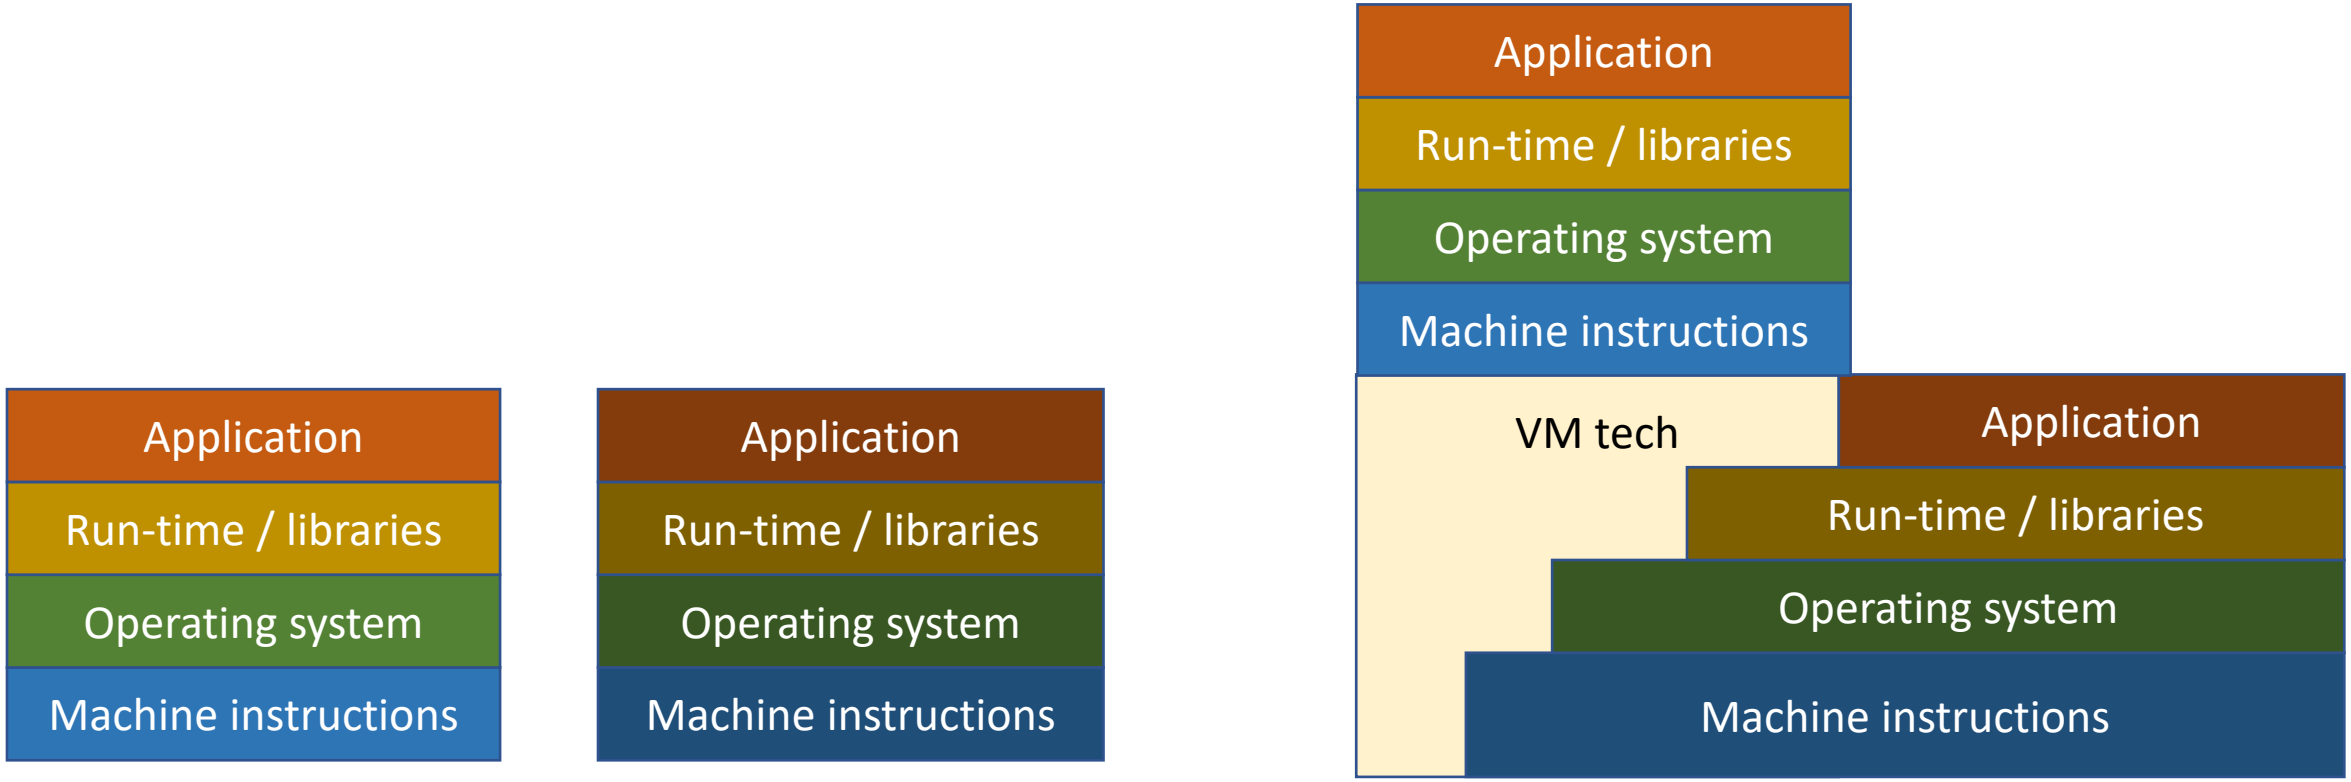

Tampereen yliopisto
Tampere University

Full virtualization

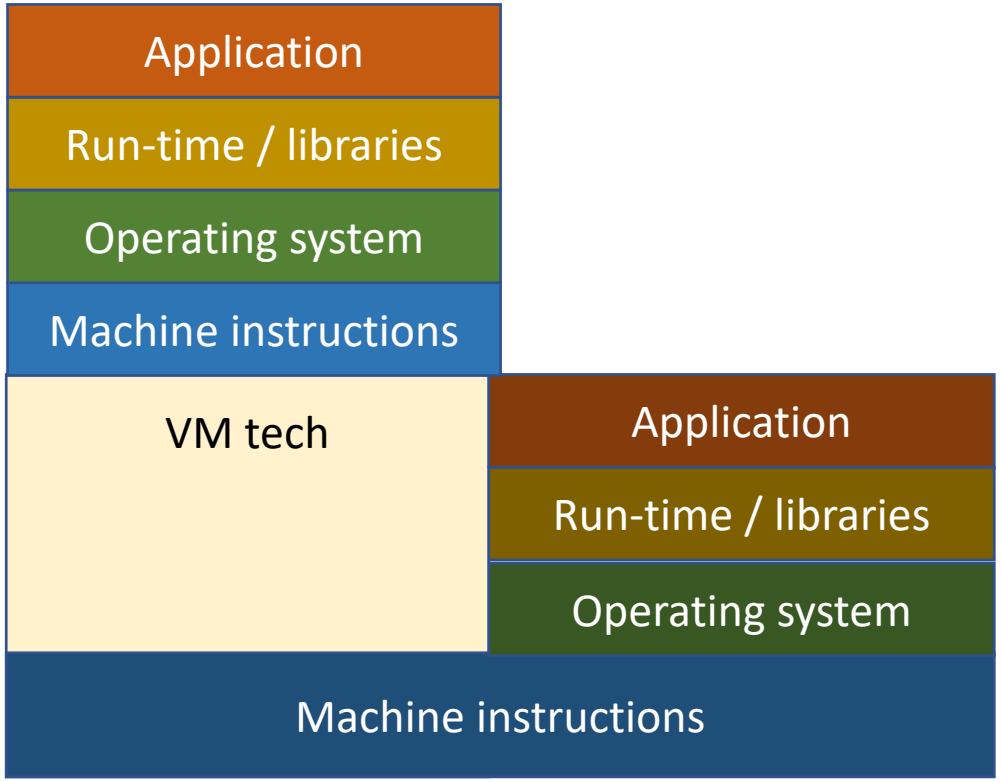

Tampereen yliopisto
Tampere University

OS as virtualization layer

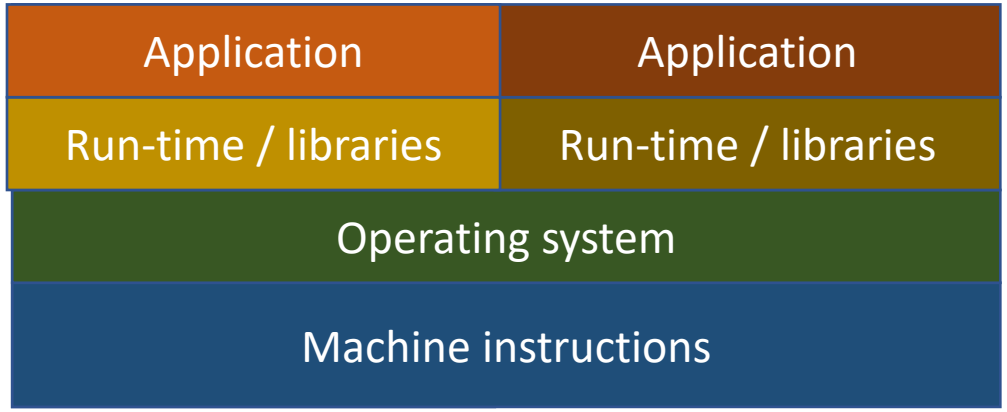

Tampereen yliopisto
Tampere University

Runtime environment (e.g. Java)

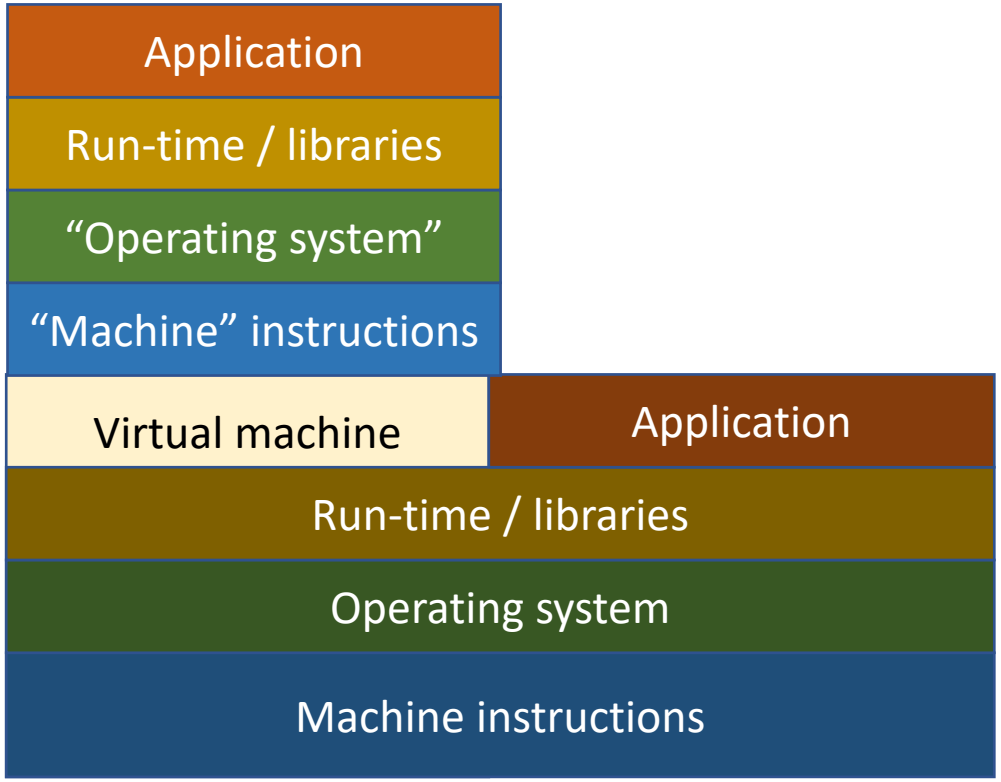

Tampereen yliopisto
Tampere University

Use case 1: run ”foreign” software

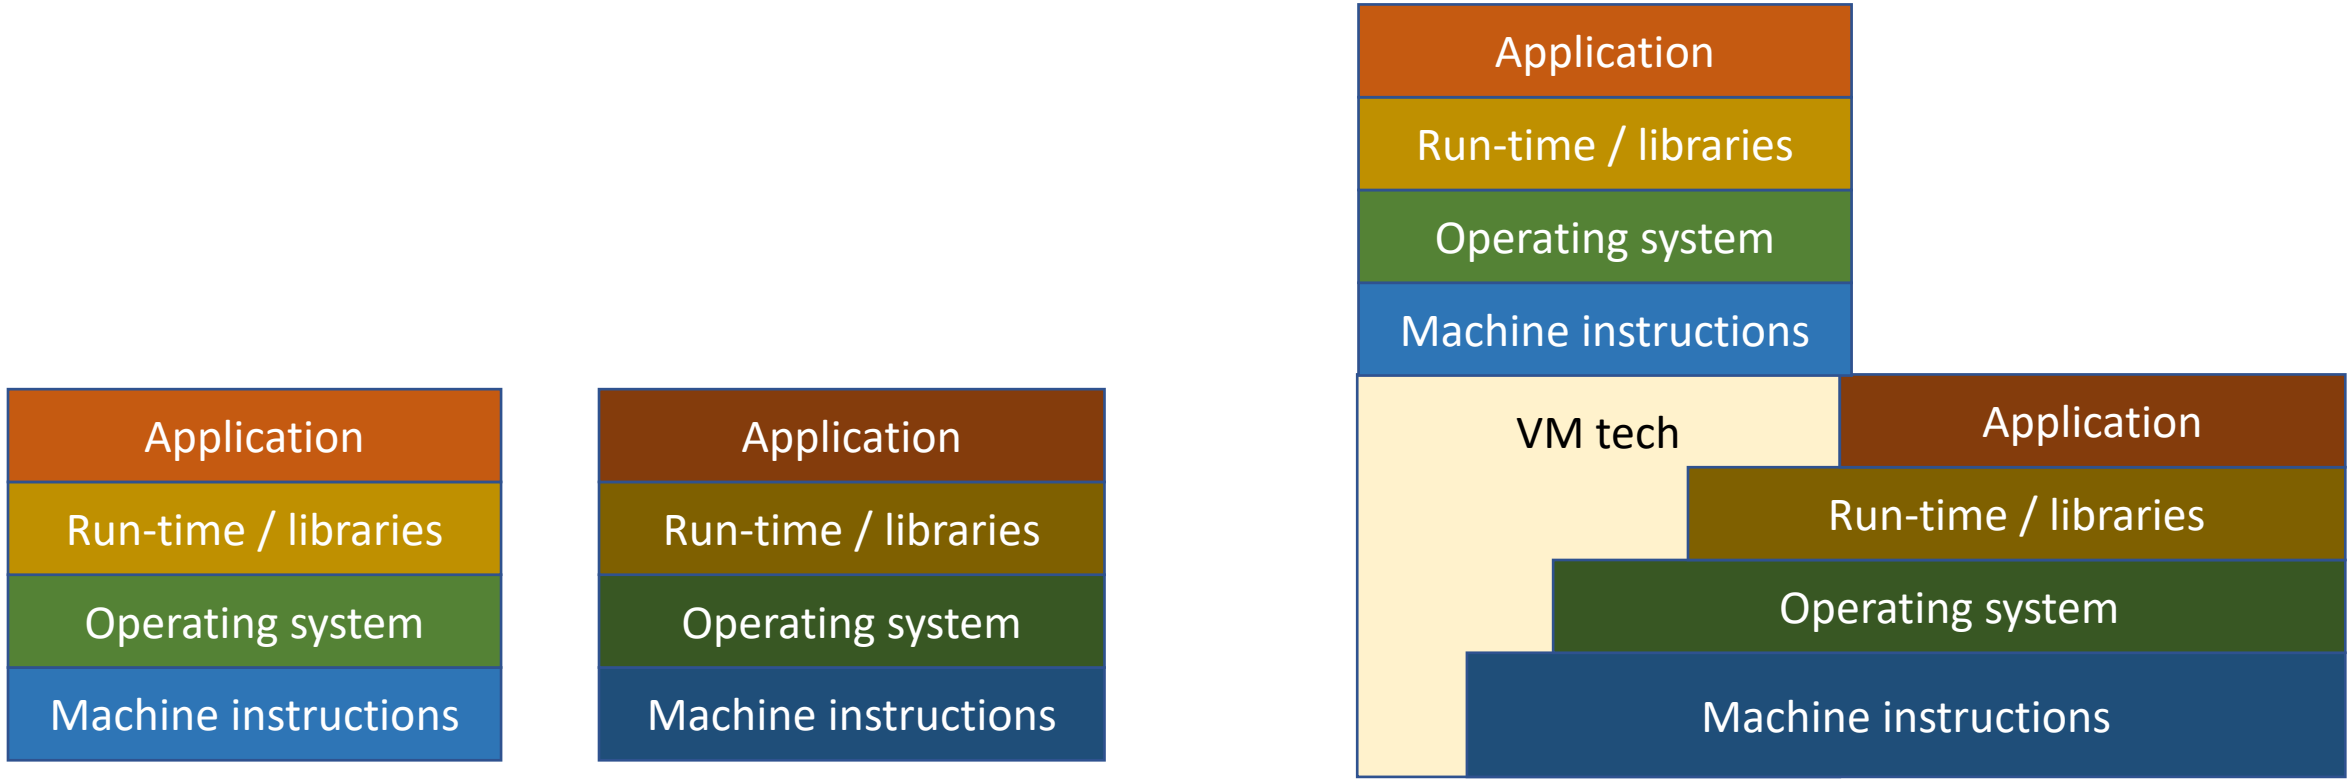

Tampereen yliopisto
Tampere University

Use case 2: isolate

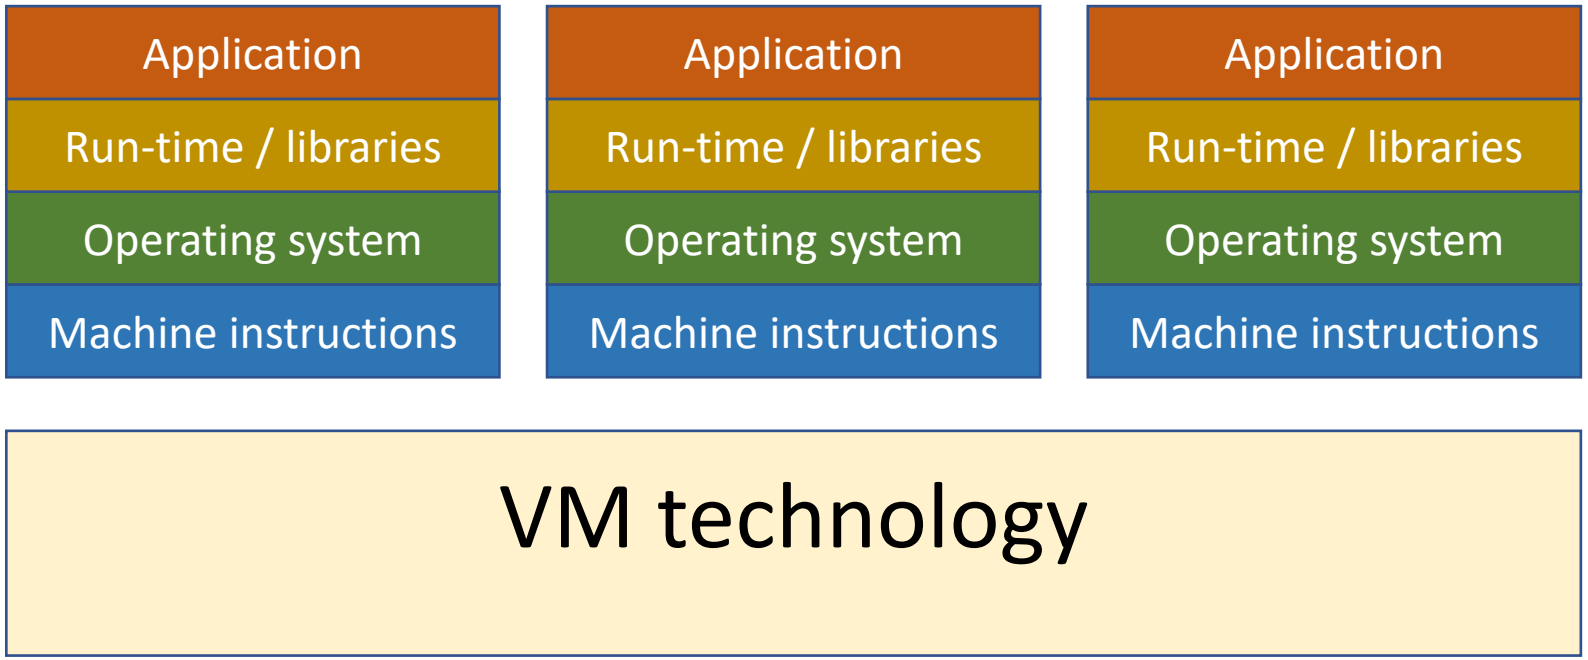

Use case 3: scale

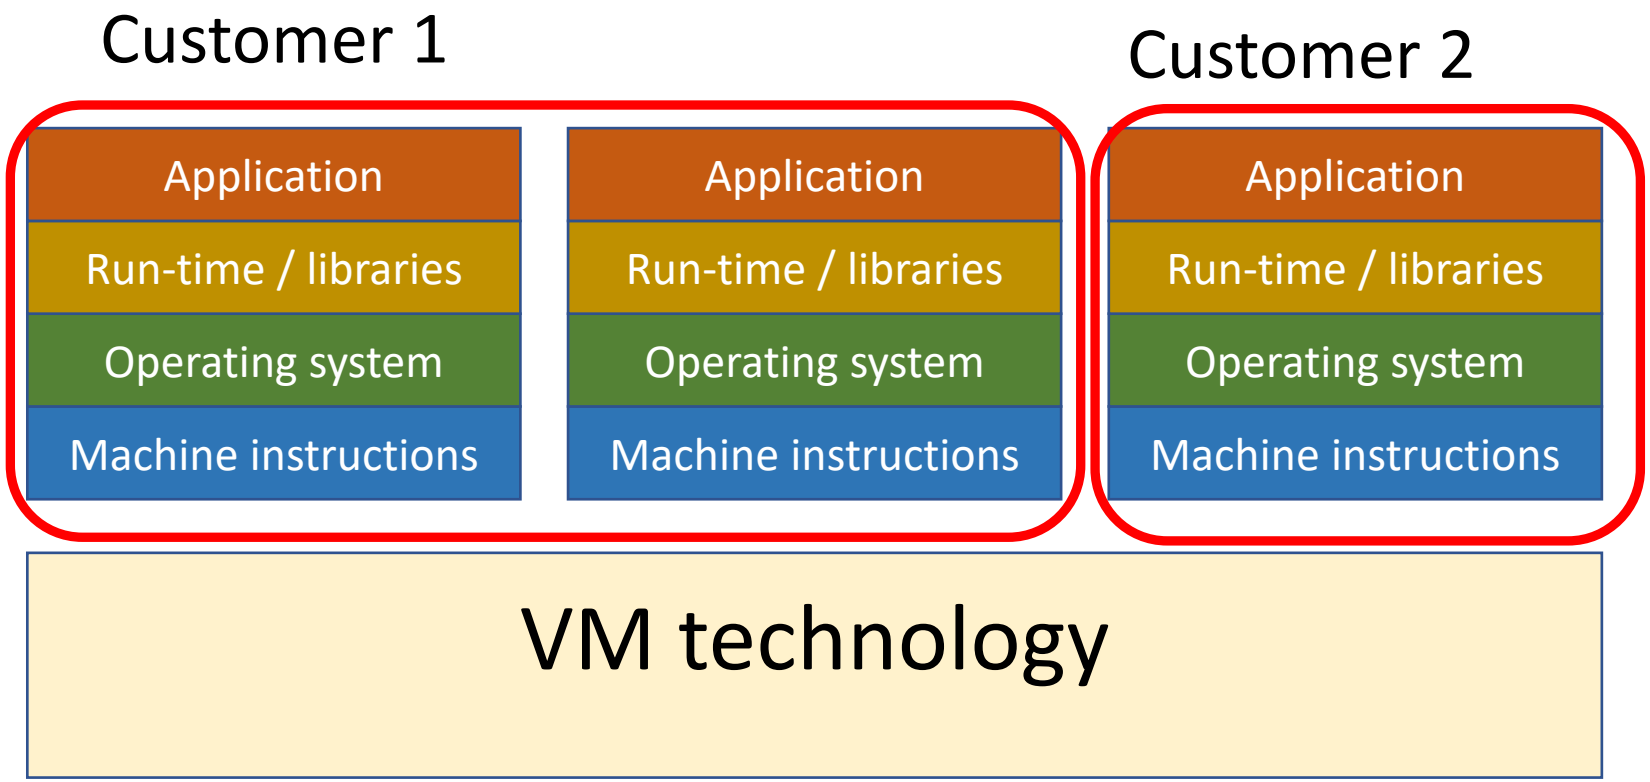

Tampereen yliopisto
Tampere University

Levels of virtualization

- Hardware virtualization
- Operating system virtualization
- Desktop virtualization
- Application virtualization
- Network virtualization

Tampereen yliopisto
Tampere University

Network virtualization

- Network, its HW and protocols, is simulated with software so that it looks like a different network to applications
- Different from OSI layer models

- Is VPN a virtual network?

Application virtualization

- Applications are compiled to machine-independent ”machine ” code
- Applications are run with a virtual machine
- Benefits
  - Same code can be run on different CPUs
  - Increased safety. Why?
- Problems
  - Performance

Tampereen yliopisto
Tampere University

Desktop virtualization

- When you run the ”linux-desktop” (linux-desktop.cc.tut.fi) on your windows machine

Tampereen yliopisto
Tampere University

Guest

Virtual Machine Mngr
(Hypervisor)

Host

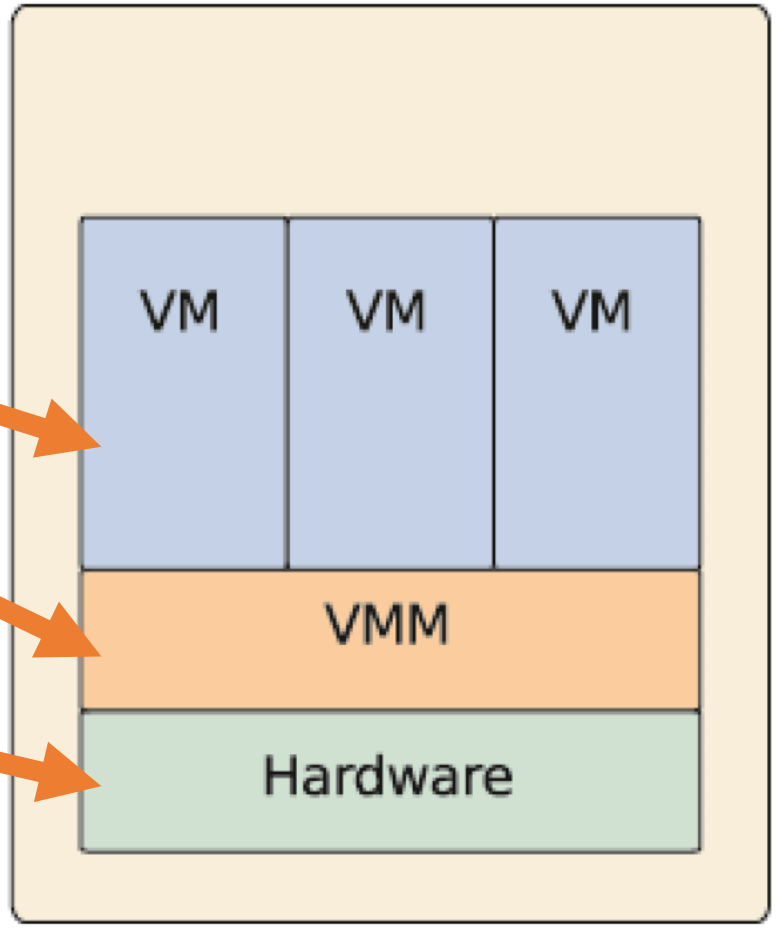

Hypervisor Type I

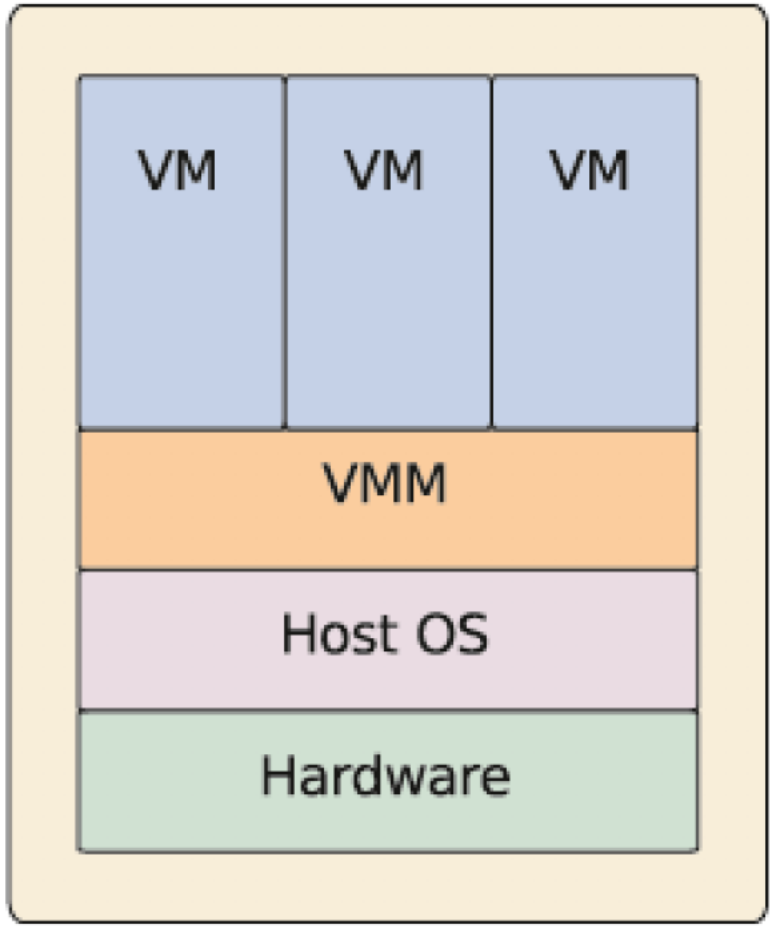

Hypervisor Type II

Mika Kaaretkoski,
Master thesis, 2018

F. Rodrguez-Haro et al., “A summary of virtualization techniques,” Procedia Technology, vol. 3, pp. 267 { 272, 2012, the 2012 Iberoamerican Conference on Electronics Engineering and Computer Science. [Online]. Accessed: 4.8.2018 Available: http://www:sciencedirect:com/science/article/pii/S2212017312002587

Tampereen yliopisto
Tampere University

More info in

- COMP.CS.110 Operating Systems