universität wien

IT4S IT for Science FORUM

11
Oct 22

IT for Science Forum

The speakers

Ronald Maier

Ronald Maier has been Vice-Rector for Digitalisation and Knowledge Transfer at the University of Vienna since 1 October 2019. He is Professor of Information Systems at the University of Innsbruck.

Ulf Busch

Ulf Busch has been CIO of the University of Vienna since 2010. He is the head of Zentraler Informatikdienst.

Toma Susi

Toma Susi is Associate Professor at the Faculty of Physics in the research unit Physics of Nanostructured Materials.

The speakers

Raman Ganguly

Raman Ganguly is Head of IT Support for Research at Zentraler Informatikdienst. The department works in the areas of content and data management, especially in the fields of research data management and open science.

Christian Kracher

Christian Kracher is Head of Coordination Digital Transformation at Zentraler Informatikdienst. Among other, the staff unit is responsible for the topics digitalisation and cloud computing at the University of Vienna.

Fabian Jusufi

Fabian Jusufi is a staff member of Coordination Digital Transformation at Zentraler Informatikdienst. He is cloud administrator of the public cloud platforms Microsoft Azure and AWS at the University of Vienna.

The speakers

Enes Bajrović

Enes Bajrović is a member of the research group Scientific Computing at the Faculty of Computer Science.

Eva Karall

Eva Karall is Head of E-Learning at Zentraler Informatikdienst. The staff unit is responsible for the technical support and further development of the e-learning platform Moodle and tools for digital teaching (BigBlueButton, Zoom) as well as the u:stream services at the University of Vienna.

Michaela Bociurko

Michaela Bociurko is Head of IT Communications and Marketing at Zentraler Informatikdienst. The staff unit bridges the gap between technical experts and users by providing user guides and documentation, organising information events and advising on product design and usability.

The topics

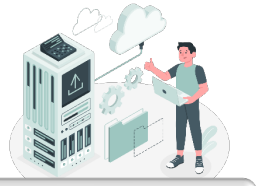

Opening

- Digitalisation for Researchers at the University of Vienna, **Ronald Maier**
- The Role of the ZID in Supporting IT for Researchers, **Ulf Busch**

IT for Open Science

- Open Science Landscape, **Toma Susi**
- Research Data Management and European Open Science Cloud: FAIR Data, **Raman Ganguly**
- Jupyter Notebook – Prototype in Computational Chemistry, **Raman Ganguly**

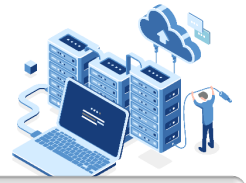

Cloud

- Cloud Strategy of the University of Vienna, **Christian Kracher**
- Cloud Services for Research and Teaching
  - Public Cloud Services: Microsoft Azure and AWS, **Fabian Jusufi**
  - Cloud Computing for Teaching and Research, **Enes Bajrović**
  - Overleaf – Collaborative Online LaTeX Editor, **Eva Karall**

Summary and Outlook

Opening

- Digitalisation for Researchers at the University of Vienna
- The Role of the ZID in Supporting IT for Researchers

universität wien

IT4S IT for Science FORUM

Digitalisation for Researchers at the University of Vienna

Presentation by Ronald Maier

Dimensions of Digital Transformations

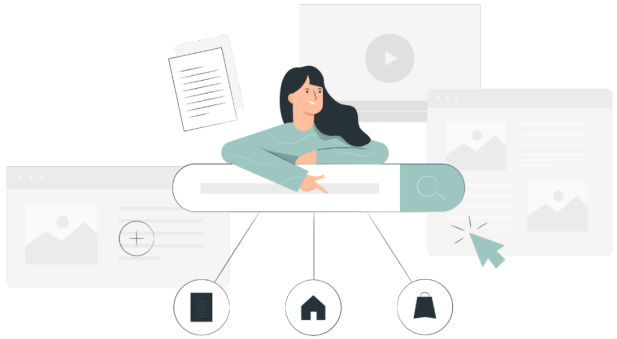

© Image by storyset on Freepik

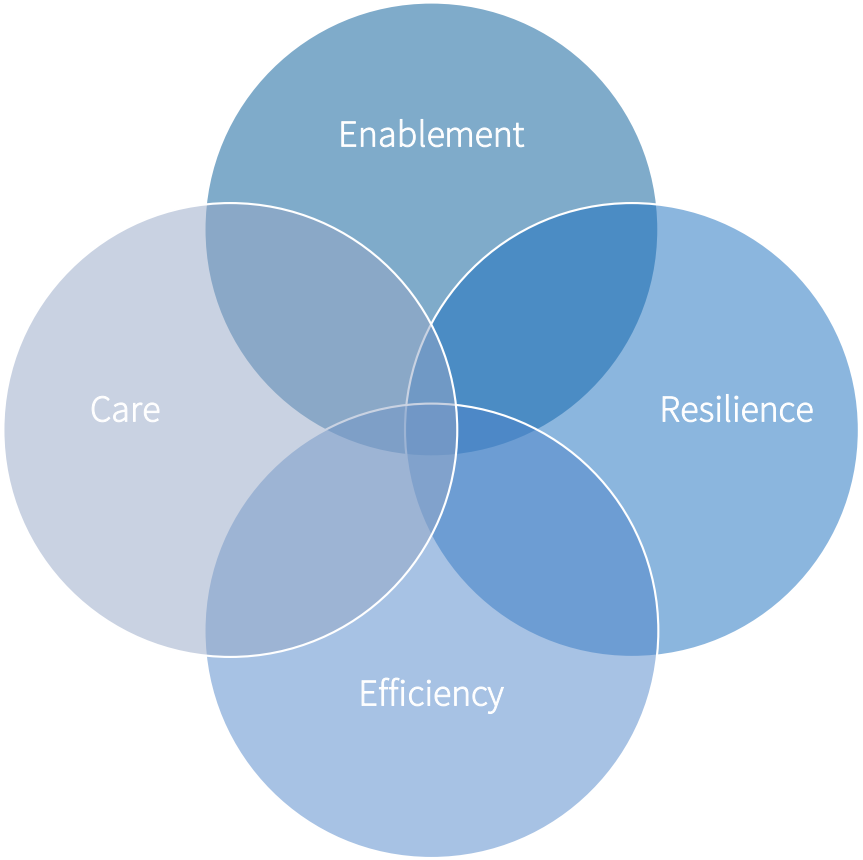

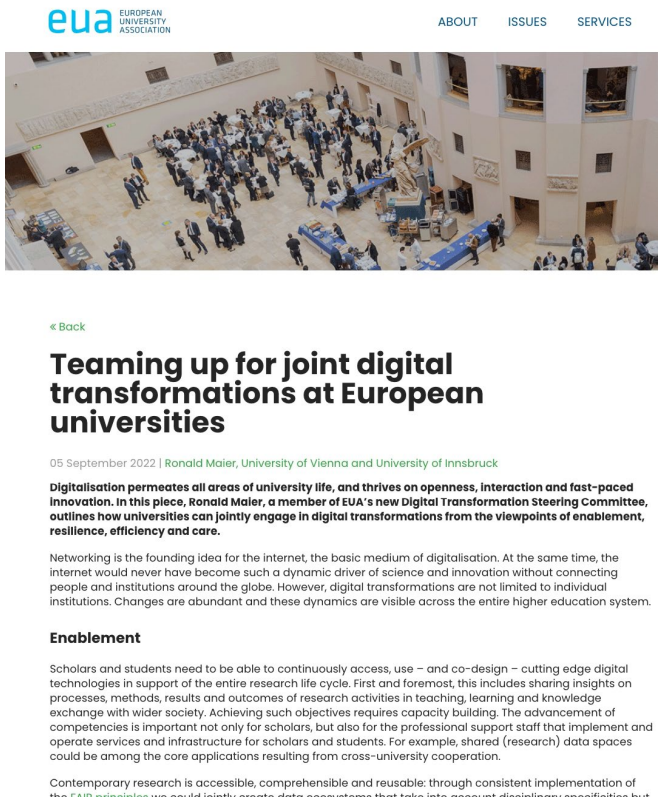

eua EUROPEAN UNIVERSITY ASSOCIATION

ABOUT ISSUES SERVICES

« Back

Teaming up for joint digital transformations at European universities

05 September 2022 | Ronald Maier, University of Vienna and University of Innsbruck

Digitalisation permeates all areas of university life, and thrives on openness, interaction and fast-paced innovation. In this piece, Ronald Maier, a member of EUA's new Digital Transformation Steering Committee, outlines how universities can jointly engage in digital transformations from the viewpoints of enablement, resilience, efficiency and care.

Networking is the founding idea for the internet, the basic medium of digitalisation. At the same time, the internet would never have become such a dynamic driver of science and innovation without connecting people and institutions around the globe. However, digital transformations are not limited to individual institutions. Changes are abundant and these dynamics are visible across the entire higher education system.

Enablement

Scholars and students need to be able to continuously access, use – and co-design – cutting edge digital technologies in support of the entire research life cycle. First and foremost, this includes sharing insights on processes, methods, results and outcomes of research activities in teaching, learning and knowledge exchange with wider society. Achieving such objectives requires capacity building. The advancement of competencies is important not only for scholars, but also for the professional support staff that implement and operate services and infrastructure for scholars and students. For example, shared (research) data spaces could be among the core applications resulting from cross-university cooperation.

Contemporary research is accessible, comprehensible and reusable: through consistent implementation of

Research and Digitalisation - Initiatives

© Image by storyset on Freepik

- Coordination Digital Transformation

  including Resonance Board

- BMBWF Projects on Digital and Social Transformation

  – Cluster Research Data

- Research Data Management Group:

  Research Services, Teaching Affairs and Student Services,

  University Library, and Vienna University Computer Center

Research and Digitalisation - Initiatives

- Data Stewardship Pilot
- 7 Lead Project Proposals submitted to BMBWF Call on (Digital) Research Infrastructures
- Open Science Initiatives
- Digital Accessibility
- And new: IT for Science Forum

Why IT for Science Forum?

- Creating a platform to
  - **present researchers with IT-solutions**, trend-setting developments and projects in the field of IT
  - **foster sharing of good practice solutions** among researchers

© Image by storyset on Freepik

The Role of the ZID in Supporting IT for Researchers

Presentation bei Ulf Busch

Developing strong IT for researchers

- Support of complex IT landscapes requires competent experts and cooperation between all areas of expertise
- Permanent secure storage of digital assets: **PHAIDRA**
- High Performance Computer: **VSC-5**
- September 2021: New policy of the University on **Research Data Management**
- New department subsequently created at the ZID in September: **IT Support for Research**
  - Sign of increased importance at the University
  - Close cooperation with researchers and Vienna University Library
    → aim: development of IT infrastructure and applications

© Image by storyset on Freepik

universität wien

IT4S IT for Science FORUM

IT for Open Science

- Open Science Landscape
- Research Data Management and European Open Science Cloud: FAIR Data
- Jupyter Notebook – Prototype in Computational Chemistry

Open Science Landscape

Presentation by Toma Susi

Open science landscape

- International landscape for Open Science is rapidly evolving
- Initial work by the European Commission set the stage
  - Open Science Policy Platform (–2018)
  - The EU's Open Science Policy (2019)
  - The European Open Science Cloud (2018–)
- *UNESCO Recommendation on Open Science* broadened the scope (2019–2021)
  - Adopted by 193 member states: first global agreement!
- Most recently, funder mandates have provided a boost for adoption
  - Plan S (2018), Horizon Europe (2019), White House OSTP mandates (2013 and 2022)

https://op.europa.eu/s/w6su

https://op.europa.eu/s/w6su

Open Science Policy Platform (OSPP)

Expert group of key stakeholder representatives advising the European Commission

- Recommendation report issued in 2018

> *"For Open Science to be successful, it must become **embedded at every level** and in every aspect of the scientific endeavour and not be perceived as separate from (or even in competition with) current practice.*
>
> *… To enable this, **all stakeholders** in research and its communication need to **take responsibility for supporting** Open Science activities.[1]"*

[1] https://op.Europa.eu/s/w6su

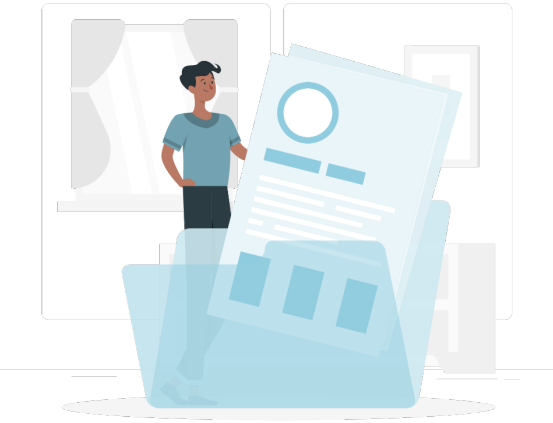

© Image by storyset on Freepik

EU Open Science Policy

Based on OSPP report, the Commission set an Open Science Policy for the EU in 2019

- Basis for grant conditions in Horizon Europe

> *"Open Science is a policy priority for the European Commission and the **standard method of working** under its research and innovation funding programmes as it improves the quality, efficiency and responsiveness of research.*
>
> *When researchers share knowledge and data as early as possible in the research process with all relevant actors it helps diffuse the latest knowledge.* [2] *"*

© Image by storyset on Freepik

[2] https://ec.europa.eu/info/research-and-innovation/strategy/goals-research-and-innovation-policy/open-science_en

UNESCO Recommendation on Open Science

Global agreement on recommendations for Open Science

- Adopted by the General Conference of UNESCO at its 41st session, in November 2021
- Recommendations are intended to influence the development of national laws and practices

> *"The Recommendation defines **shared values and principles** for Open Science, and identifies **concrete measures** on Open Access and Open Data, with proposals to bring citizens closer to science and commitments to facilitate the production and dissemination of scientific knowledge around the world. The Recommendation was developed through a regionally balanced, multistakeholder, inclusive and transparent consultation process.[4]"*

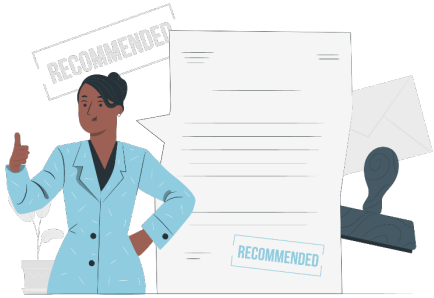

[4] https://unesdoc.unesco.org/ark:/48223/pf0000379949.locale=en

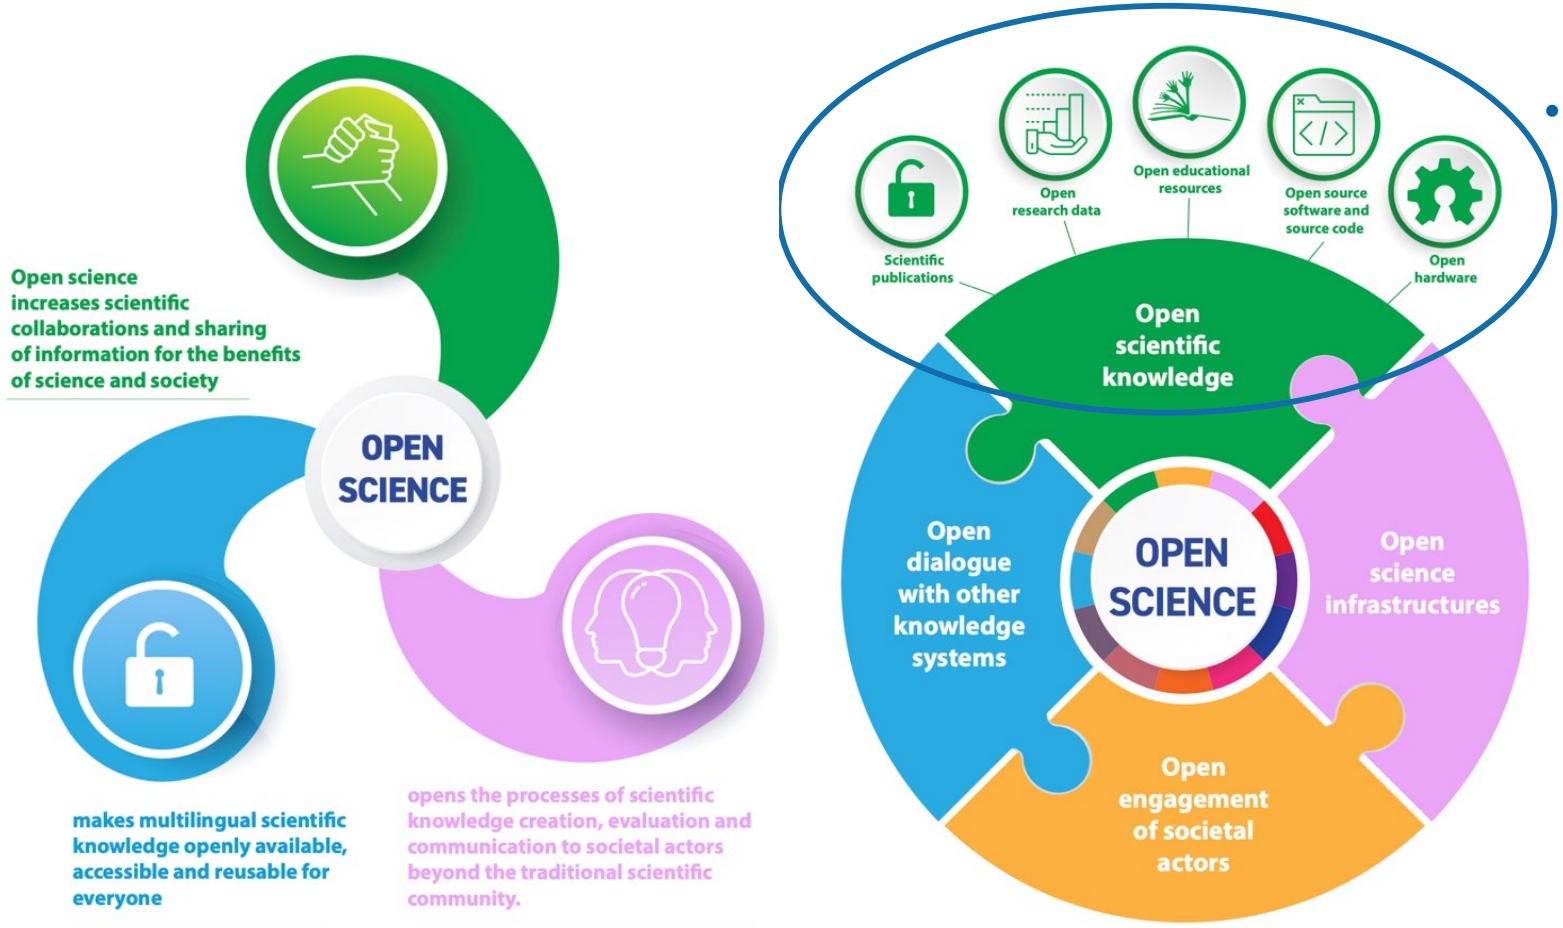

- Diversity of practices, not only publications!

https://unesdoc.unesco.org/ark:/48223/pf0000379949.locale=en

Funder mandates (Europe)

- Plan S open access mandate (in effect from 2021)
  - Public and private funders across Europe and beyond agreed on a common policy

> *"All scholarly publications on the results from research funded by public or private grants provided by national, regional and international research councils and funding bodies, must be published in Open Access Journals, on Open Access Platforms, or made* ***immediately available*** *through Open Access Repositories without embargo."*

- Horizon Europe Open Science Policies (starting in 2020)
  - Open Access to publications (*immediate, mandatory, CC-BY license*)
  - Open Access to related research data (*as soon as possible, mandatory*)
  - Data management plans (*mandatory, FAIR)*

© Image by storyset on Freepik

Funder mandates (USA)

- Office of Science and Technology Policy (OSTP) mandate in 2013 (“Holdren memo”)
  - Public access to much of federally funded research with a 12-month embargo
  - Only the biggest federal agencies included, only recommendations for Open Data

- OSTP mandate in 2022 (“Nelson memo”)
  - Applies to all federal agencies
  - No publication embargo allowed, immediate access
  - Related data must also be available immediately
  - Coming into effect by 2025 at the latest

www.whitehouse.gov/ostp/news-updates/2022/08/25/ostp-issues-guidance-to-make-federally-funded-research-freely-available-without-delay/

Reforming research assessment

- Increasing agreement that current ways are not fit for purpose
  - Transparent, reproducible and robust science is not properly rewarded
  - Encourages (toxic) competition instead of collaboration and sharing
  - Holds back Open Science and innovative publishing models
- Coalition for Advancing Research Assessment launched in September 2022
  - Co-creation of 350+ organisations from 40+ countries

> *"It is now time to go beyond existing declarations and define clearly what we want for the future of research assessment. Universities, researchers and all stakeholders* ***need to choose*** *how they want to be assessed and need to choose* ***now.***[4] *"*

[4] https://coara.eu

© Image by storyset on Freepik

Questions about Open Science Landscape?

- Please ask your questions as a chat comment.

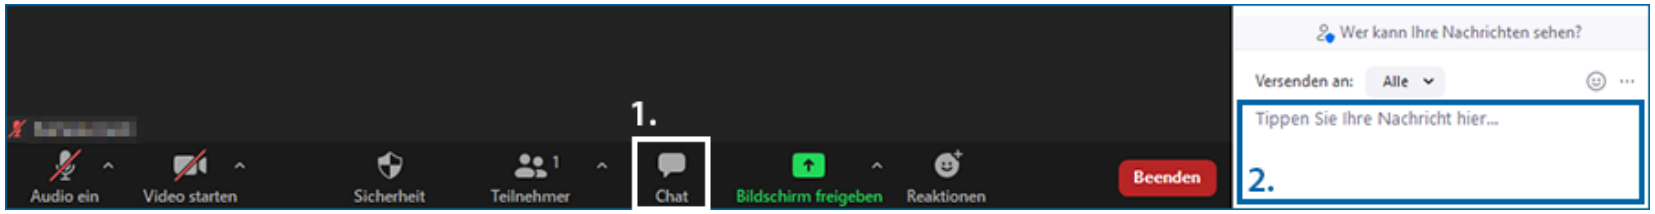

Research Data Management and European Open Science Cloud: FAIR Data

Presentation by Raman Ganguly

Model for data management

- Pre-ingest: generation of data
- Ingest: handover to the data management
- Management: preserve data over time
- Re-use: use the data delivered from the data management

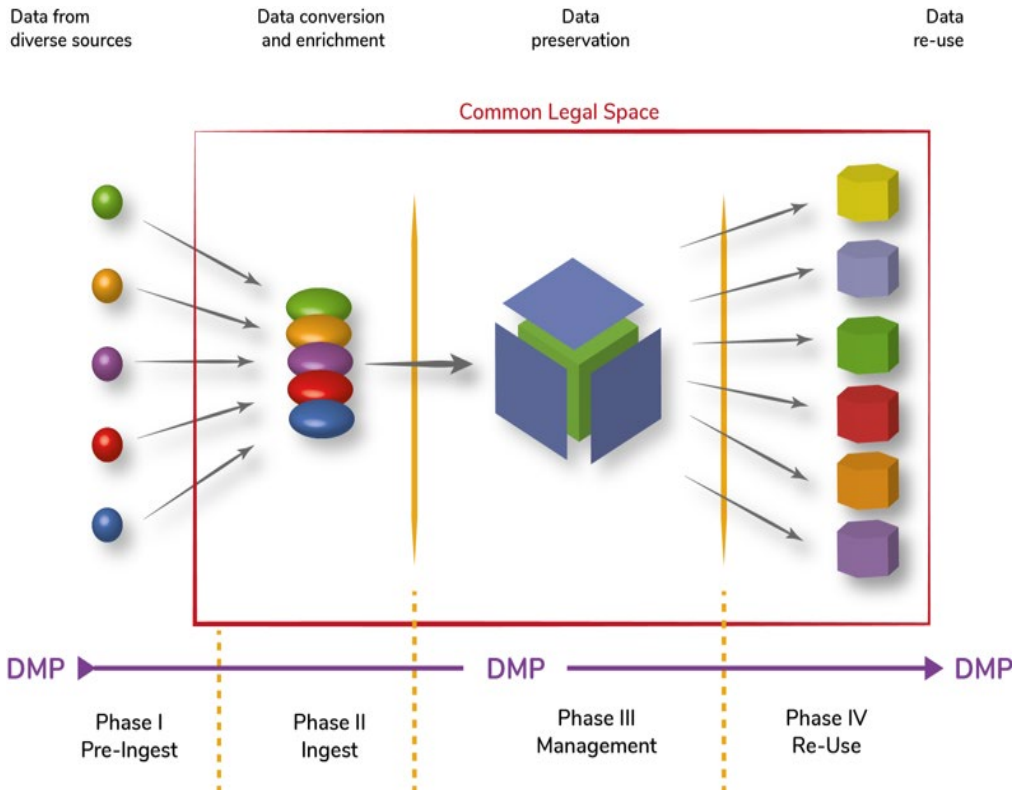

FAIR data principles

- **F**indable, **A**ccessible, **I**nteroperable, **R**eusable
- Data principles for data preservation
- Using open standards to ensure a long-term reusability
- The principles are for data and their metadata

Research data

- “research data” means all data that is created in the course of scientific research and artistic creation processes (e.g. text, tables, video, audio, graphics, etc.) and/or on the basis of which your research results and/or works of art are based – e.g. through experiments, source research, measurements, surveys, digital copies or drafts.

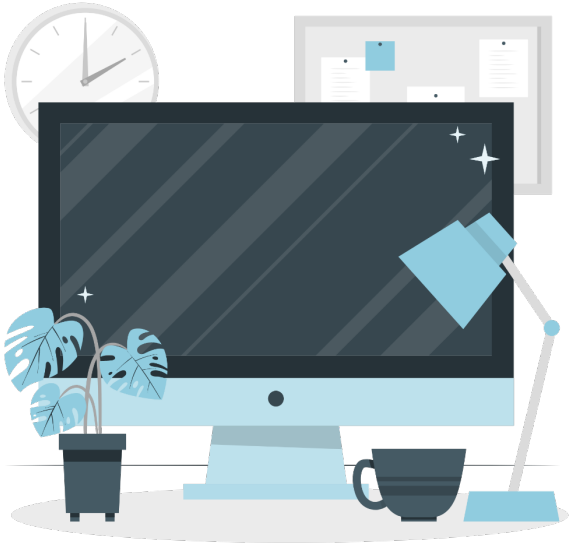

© Image by storyset on Freepik

Enabling data management

- Organising data during data creation
- Support during the ingest
- Enrichment of data with metadata
- Infrastructure for preserving different kind of data
- Providing data catalogues with metadata

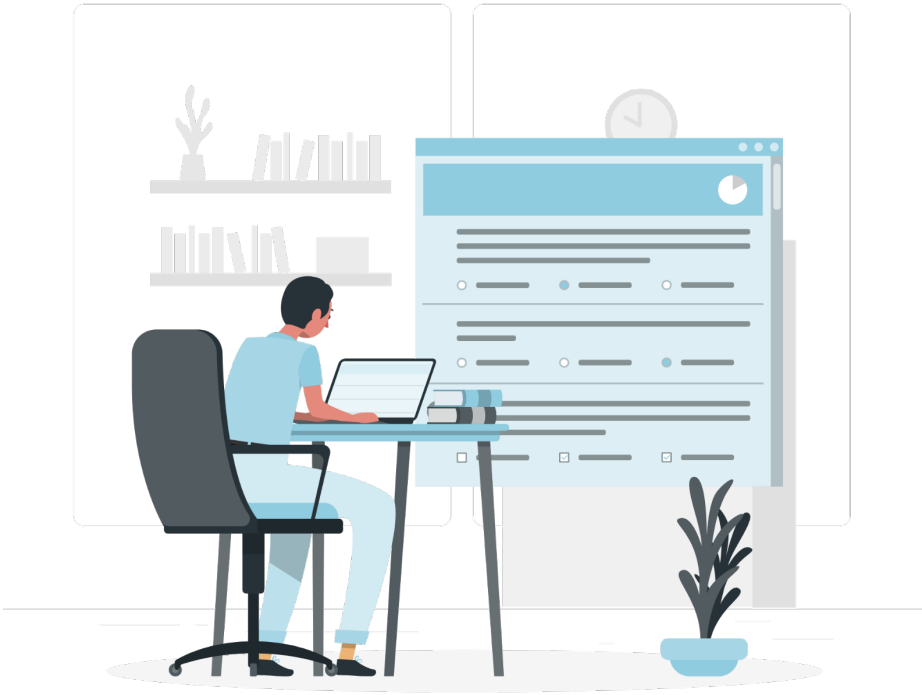

European Open Science Cloud

- Foundation are existing infrastructures for research
- Re-use of data is main driver, FAIR data principles are base for re-use
- European data market for research data
- Infrastructure for computation and preservation of data
- Participation on different levels
- EOSC Office and FAIR Data Office in Austria

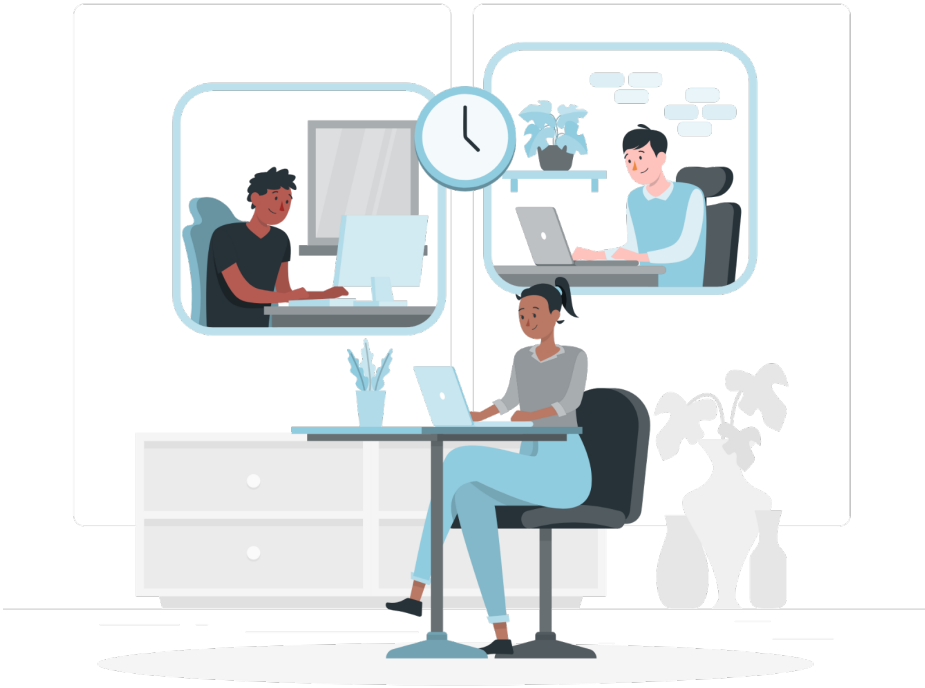

Questions about Research Data Management and European Open Science Cloud: FAIR Data?

- Please ask your questions as a chat comment.

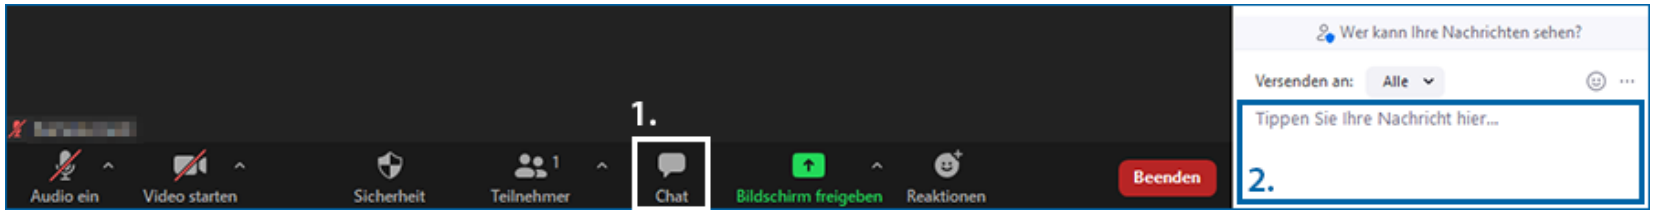

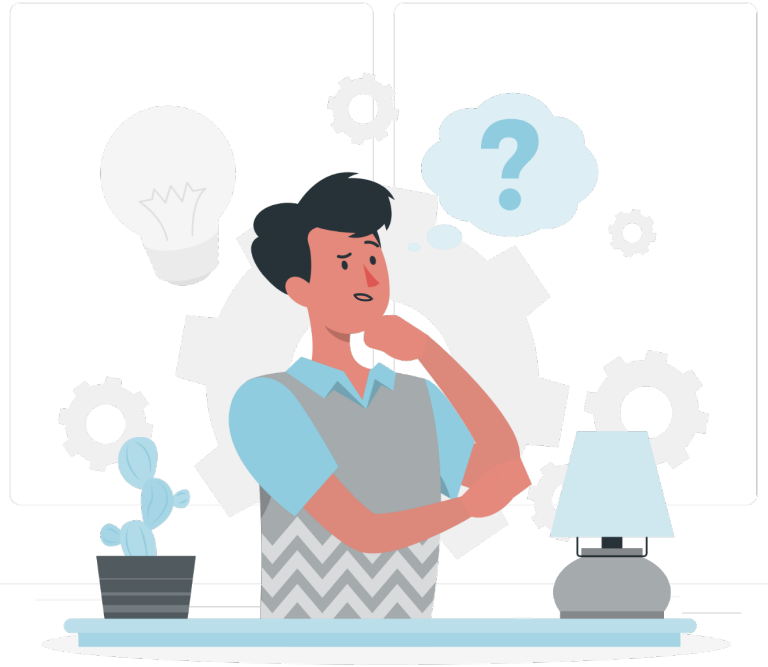

© Image by storyset on Freepik

universität wien

IT4S IT for Science FORUM

Jupyter Notebook – Prototype in Computational Chemistry

Presentation by Raman Ganguly

Initial situation

- Use of Jupyter Notebooks in teaching
- Test in winter semester 2021/22 with department for Computational Biological Chemistry
- Cooperation with e-learning team, integration in Moodle
- Institute already had experience in teaching with Jupyter Notebooks
- Request: cooperation with ZID for better support
- Pilot project financed by the Rectorate via “Active Studying” project

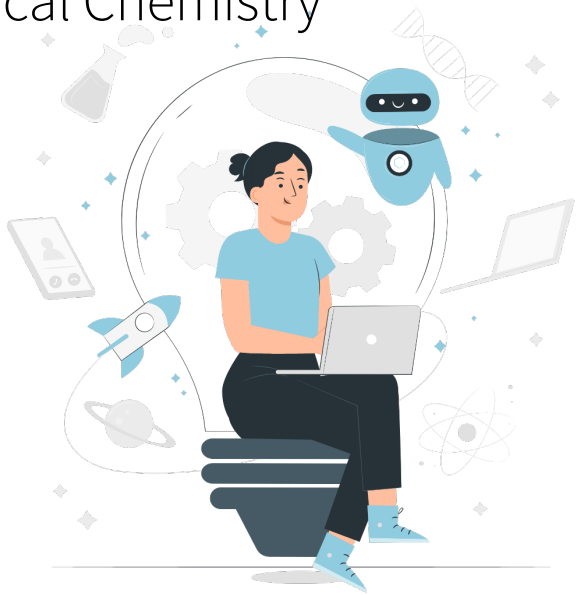

Implementation

- Technically very large planned, as required resources not foreseeable
- Infrastructure of Jupyter Hub developed together with colleagues of the institute
- Integration into Moodle realised with e-learning team

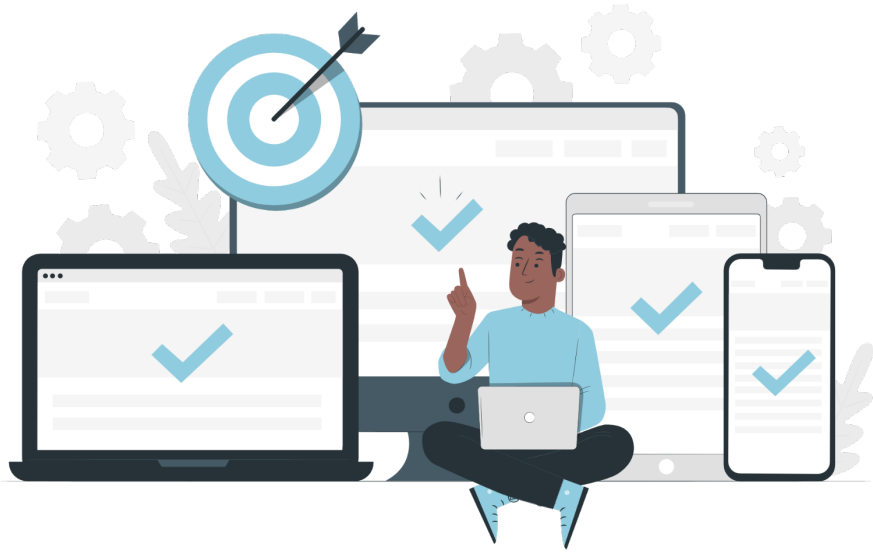

© Image by storyset on Freepik

Conclusion and next steps

- Significantly fewer computing resources were needed than planned
- Service not ready for people yet with lower experience
- Service not ready to scale up with current setup
- Aiming for continuation in current semester
- Extension of pilot is planned, as well as development of a service

Questions about Jupyter Notebook – Prototype in Computational Chemistry?

- Please ask your questions as a chat comment.

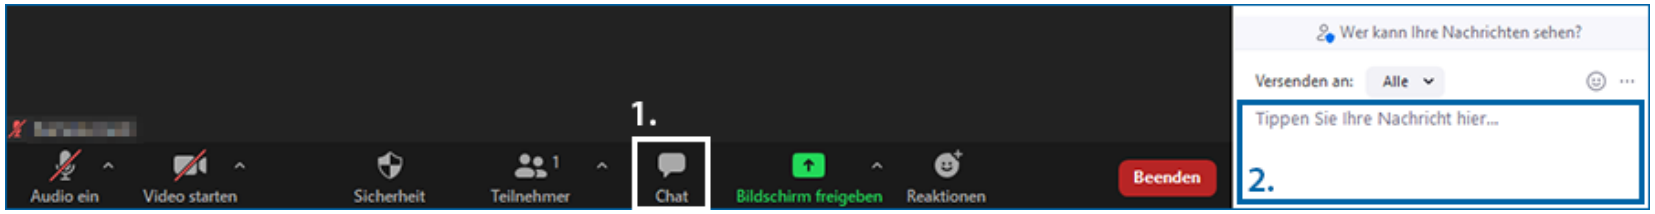

© Image by storyset on Freepik

Cloud

- Cloud Strategy of the University of Vienna
- Cloud Services for Research and Teaching

universität wien

IT4S IT for Science FORUM

Cloud Strategy of the University of Vienna

Presentation by Christian Kracher

What happened until now

- Cloud Policy of the University of Vienna according to the IT Master Plan
- Recommendation for Cloud Strategy has been developed
- Cloud Strategy has been published

© Image by fullvector on Freepik

Cloud Policy according to the IT Masterplan (v. 1.0, 2015)

"The operation of services in the public cloud is not desired due to unresolved questions regarding data security, data protection and long-term perspectives. The development of services within the framework of a private cloud, which is operated exclusively by the University of Vienna or within the framework of an Austrian university network, is being pursued. Should a service be available exclusively as a SAAS solution, explicit coordination with the Rectorate must take place. It must be ensured that the questions regarding data security, data protection and long-term perspective, which the operation of services in the public cloud/hybrid cloud entails, have been clarified.[5]"

Der Betrieb von Services in der Public Cloud ist auf Grund von ungeklärten Fragen zu Datensicherheit, Datenschutz und Langzeitperspektive nicht gewünscht. Der Aufbau von Services im Rahmen einer Private Cloud, die ausschließlich von der Universität Wien bzw. im Rahmen eines österreichischen Universitätenverbunds betrieben wird, wird verfolgt. Sollte ein Service ausschließlich als SAAS-Lösung verfügbar sein, so hat eine explizite Abstimmung mit dem Rektorat zu erfolgen. Es ist sicherzustellen, dass die Fragen zu Datensicherheit, Datenschutz und Langzeitperspektive, die der Betrieb von Services in der Public Cloud/Hybrid Cloud mit sich bringt, geklärt wurden.

[5] IT Masterplan, p. 14

Recommended actions for implementing the Cloud Strategy

- Creating awareness among members of the University of Vienna in the processing of their data
- Creating and approving organisational guidelines
- Accompanying change processes resulting from cloud use in the organisation
- Considering characteristics of the cloud when introducing and developing IT services (e.g. self-service)
- Supporting research and teaching with services

©Joseph Krpelan (Cover), Cloud Strategy of the University of Vienna

Cloud Strategy

- Strategy for members of the University of Vienna
- Perspective for users and operators
- Explanation for cloud computing
- High Level Mission Vision Statement
- Positioning of the University of Vienna on important topics in the context of cloud usage
- Strategic goals when using cloud resources
- Responsibility of a service owner
- zid.univie.ac.at/en/it-worlds/cloud-computing/

Questions about Cloud Strategy of the University of Vienna?

- Please ask your questions as a chat comment.

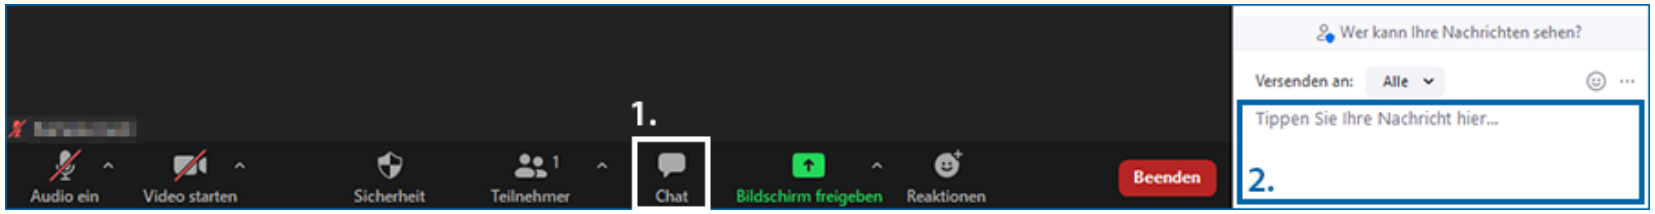

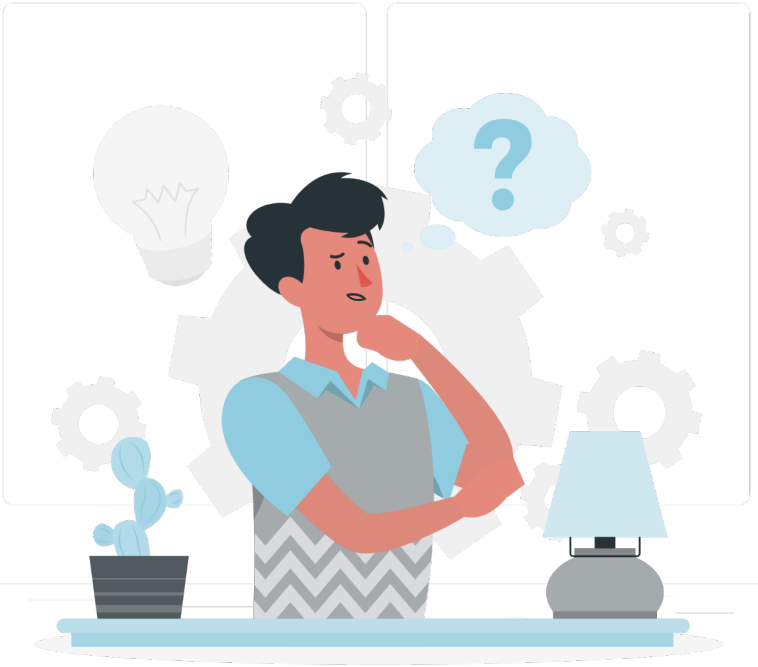

© Image by storyset on Freepik

universität wien

IT4S IT for Science FORUM

Cloud Services for Research and Teaching

- Public Cloud Services: Microsoft Azure and AWS
- Cloud Computing for Teaching and Research
- Overleaf – Collaborative Online LaTeX Editor

universität wien

IT4S IT for Science FORUM

Public Cloud Services: Microsoft Azure and AWS

Presentation by Fabian Jusufi

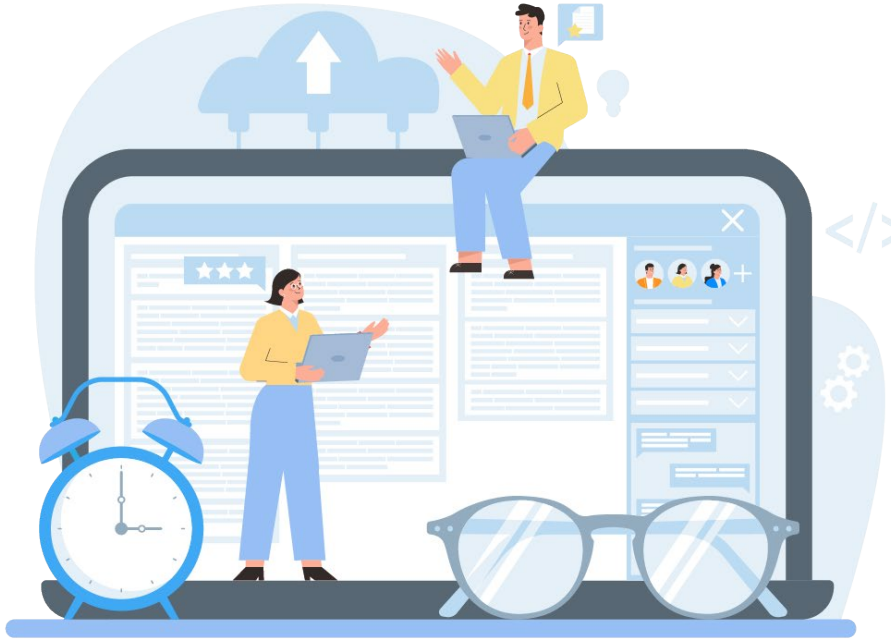

Agenda

- Public cloud for researchers
  - Definition
  - Benefits and risks
- Cloud testing and offering
  - Microsoft Azure
  - AWS (Amazon Web Services)

© Image by storyset on Freepik

Definition of cloud computing

- Five essential characteristics of cloud computing according to NIST (National Institute of Standards and Technology):

Source: cic.gsa.gov/basics/cloud-basics

Benefits of public cloud infrastructure (1/2)

- On demand self-service and broad network access:
  - Agile environment: deployment of resources anytime via the internet
- Rapid elasticity:
  - Fast, on-demand scalability
- Resource pooling and measured services:
  - Dynamic and efficient environment with cost saving potential
  - Live consumption
  - No upfront costs, pay-as-you-go

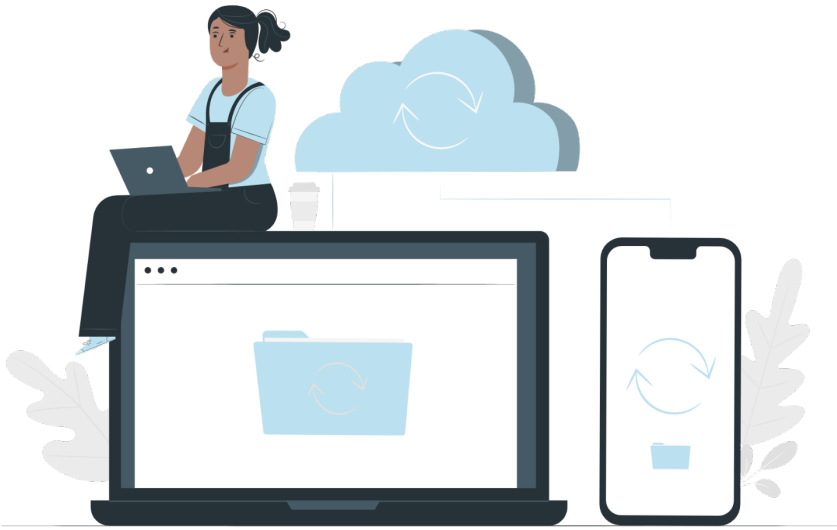

© Image by storyset on Freepik

Benefits of public cloud infrastructure (2/2)

- Always up-to-date:
  - Applications and hardware maintained by the CSP
  - Researchers can focus fully on research goals
- No waiting time for the latest hardware:
  - Nowadays long lead times when ordering hardware
- Easy collaboration with researchers of other institutions
- Services that are not available on-premises:
  - i.e. many AI and machine learning services

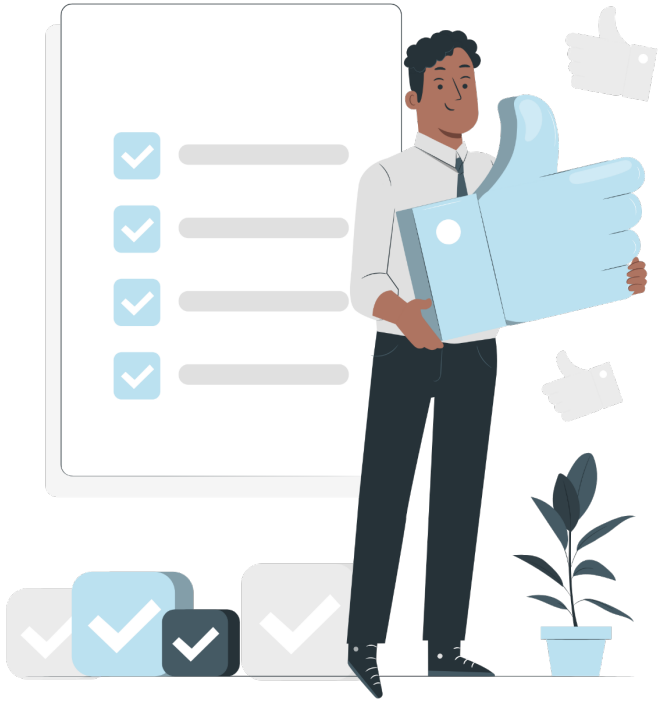

© Image by storyset on Freepik

Risks of using the public cloud

- Data and information leave own organisation
  - Need to trust a third party (CSP)
  - Fear of loss of control over own data

- Possible compliance and security issues for researchers
  - Compliance with policies of the University of Vienna
  - GDPR conformity

© Image by storyset on Freepik

Benefitting from public cloud while minimising risks

- We understand the need to use cloud services
  - However, risks must be addressed and mitigated

- ZID as central point of contact for public cloud services
  - Enabler of services and support
  - GDPR conform access
  - Organisational and technical measurements
    - Collaboration with staff unit IT Security and Data Protection Officer of the University of Vienna

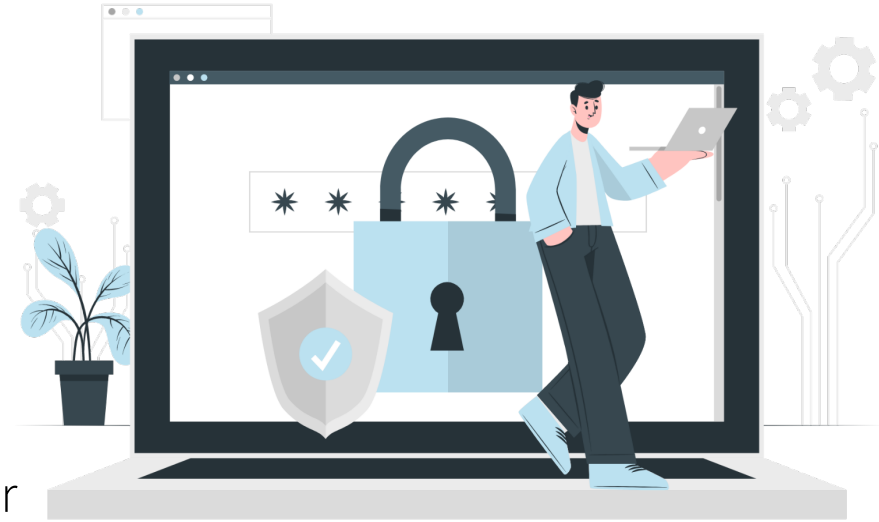

© Image by storyset on Freepik

Microsoft Azure

- Public cloud platform of Microsoft
- More than 200 services available in 58 regions worldwide (region Austria by 2023)

Infrastructure services

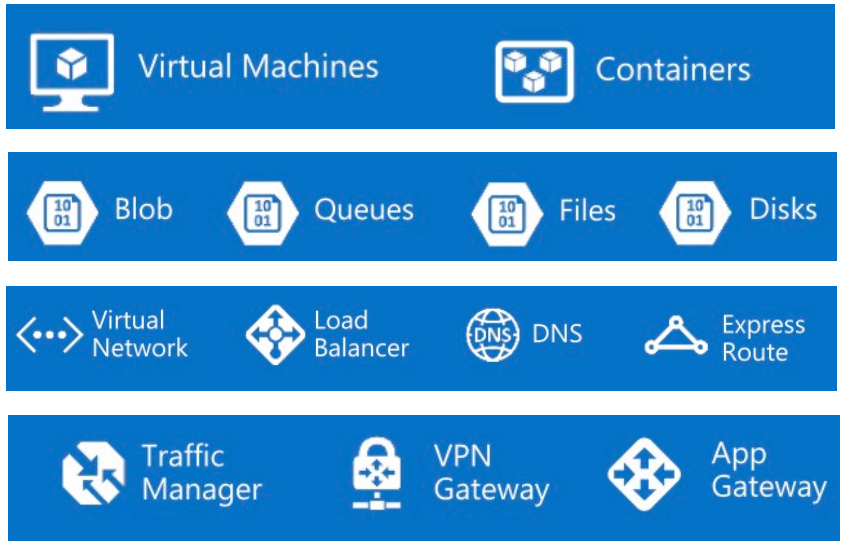

Platform services

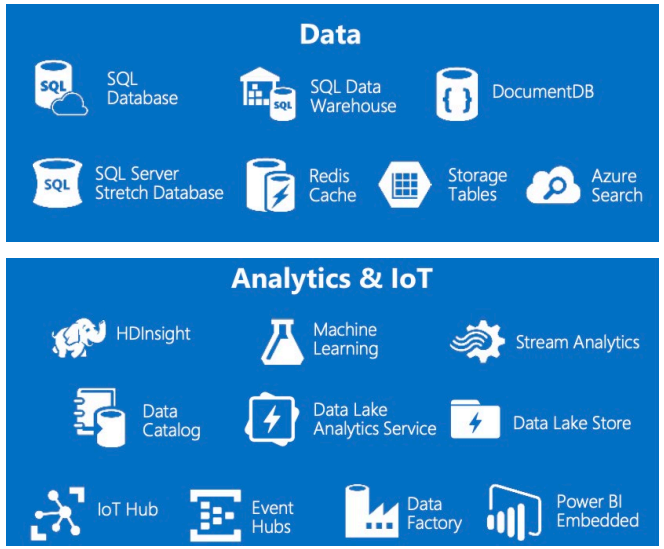

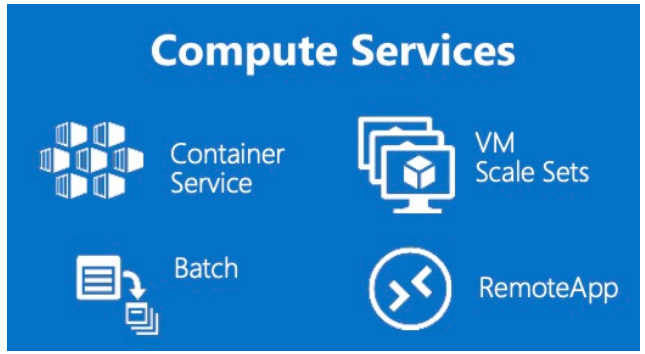

Microsoft Azure – our organisation

- Testing since 2018
  - Collaboration with Faculty of Computer Science since 2019

- Goals:
  - First: learning about the platform
  - Later: trying to set up an organisation in the cloud and making it an added value for the University of Vienna

- Current numbers: 1 tenant (azure.univie.ac.at)
  43 users
  48 running applications

© Image by storyset on Freepik

Azure – collaboration with Faculty of Computer Science

- Strong collaboration with Univ.-Prof. Dr. Wolfgang Klas
  - Research group Multimedia Information Systems
  - Dedicated student projects

- Mainly used services:
  - Virtual machines and networks
  - Azure container instances
  - App services
  - Blockchain services
  - Storage accounts
  - Kubernetes service
  - Web app for container
  - SQL databases

© Image by storyset on Freepik

- Positive feedback with high interest to expand the collaboration

Azure as official service of the ZID (1/2)

- Ongoing project for the introduction of Microsoft Azure as official service of the ZID

- Key benefits:
  - Log-In with u:account credentials
  - 15 % discount as part of the OCRE framework
  - GDPR compliant environment
  - Implemented measures for data security
  - Support from the ZID and Microsoft support plan included

© Image by storyset on Freepik

Azure as official service of the ZID (2/2)

- Current project progress:
  - Most technical requirements completed
  - Documentation of service ready
    - Waiting for final approval

- Nevertheless possibility to use current test environment with special conditions
  - Interested?
    E-mail with concrete research use case

© Image by storyset on Freepik

AWS (Amazon Web Services) – our organisation

- Public cloud platform of Amazon (100 % subsidiary)
- Since 2018: parallel testing to Microsoft Azure
  - Trying to understand benefits of each platform, differences and similarities
  - Currently running test environment with multiple accounts
  - ZID-internal tests but also access for some researchers
- Since Autumn 2021: new contract with Rackspace (OCRE framework) with many benefits:
  - 15 % discounts
  - Egress data free in most scenarios

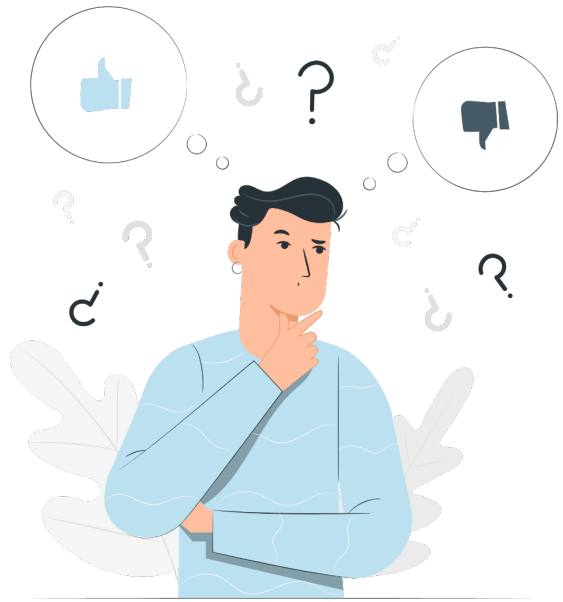

© Image by storyset on Freepik

AWS – collaboration with researchers

- Department of Government
  - AWS Transcribe: AI service that converts automatically speech to text

- Faculty of Computer Science
  - Prof. Dr. Siegfried Benkner
  - Research Group Scientific Computing
  - EC2 (virtual machines) and high performance computing

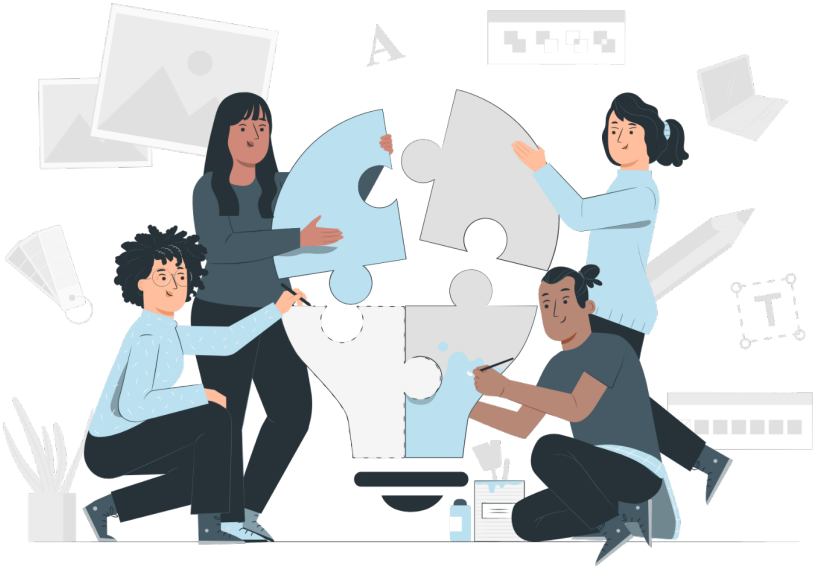

© Image by storyset on Freepik

Our offer to researchers

- Private accounts with personal credit cards are not the right way of doing research.
- If you need cloud services to reach your research goals, contact us.
- We can support you achieving them while protecting your data and saving money.
- More information: zid.univie.ac.at/en/it-worlds/cloud-computing/

- Contact:
  - Fabian Jusufi / Coordination Digital Transformation
  - E-mail: fabian.jusufi@univie.ac.at

© Image by storyset on Freepik

Questions about Public Cloud Services: Microsoft Azure and AWS?

- Please ask your questions as a chat comment.

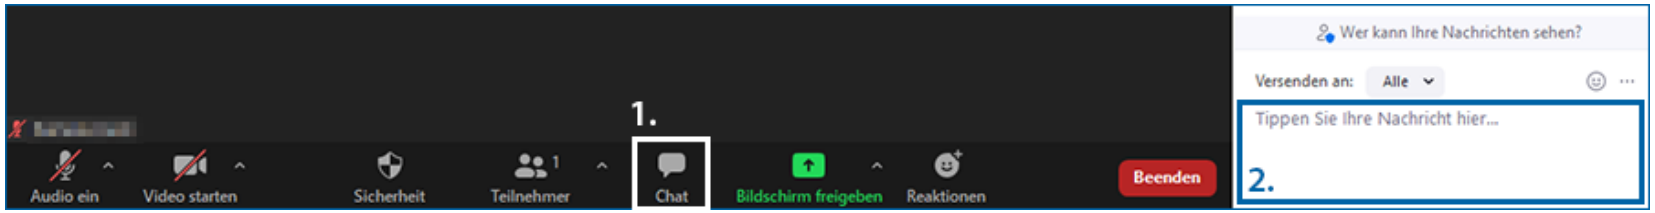

universität wien

IT4S IT for Science FORUM

Cloud Computing for Teaching and Research

Presentation by Enes Bajrović

Agenda

- Introduction
- Cloud computing for research
  - Potential benefits for research in Computer Science
  - Experiences with AWS
- Cloud computing and teaching
  - “Cloud computing” lecture
  - Experiences with cloud computing in our lectures
- Conclusions
  - Summary and outlook

© Image by storyset on Freepik

Research group Scientific Computing (Head: Prof. Siegfried Benkner)

- Parallel and high performance computing, cloud computing, scientific applications
- Programming models and tools to help users in solving increasingly complex compute-, and data-intensive problems in science and engineering
- Participated in many international research projects in HPC and cloud computing
- AutoTune, Peppher, VPH-Share, Aneurist (all EU), FWF Retida, ...

How can we use cloud computing as a research vehicle for our research and teaching activities?

sc.cs.univie.ac.at

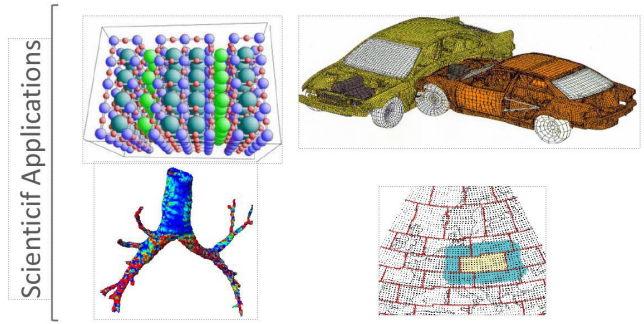

Research challenges (1/2)

- Keeping up with fast developing hardware
  - Early access essential for research
  - Parallel and specialised processors (GPU, FPGAs, NE, …)
  - HPC clusters comprised of different (heterogeneous) nodes
- Scalability
  - Testing on larger clusters essential

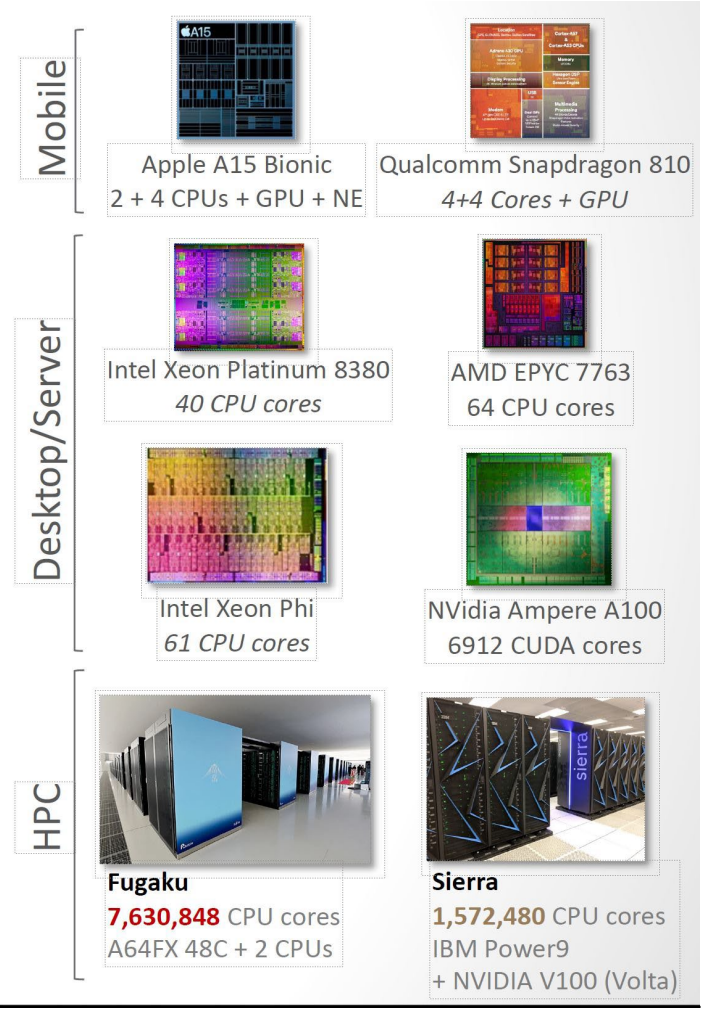

Research challenges (2/2)

- Performance and portability
  - Access to different types of architectures and cluster configurations
- Challenging to keep on with on-premises architecture
  - Using the newest available architectures
  - Keeping up with the hardware innovation cycles
  - Quickly scale and
  - Keeping the costs low

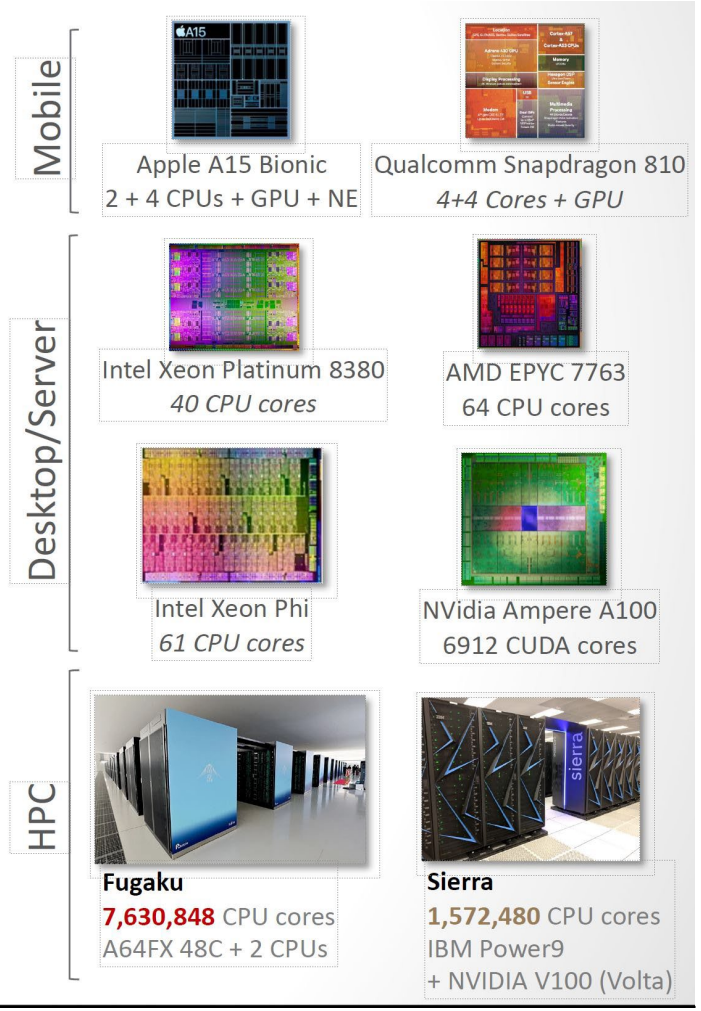

Cloud computing is …

“… on-demand delivery of compute power, database, storage, applications, and other IT resources through a cloud services platform via the internet with pay-as-you-go pricing.” (AWS)

- Service models
  - SaaS (Software as a Service), FaaS (Function as a Service), PaaS (Platform as a Service), IaaS (Infrastructure as a Service)
- Deployment models
  - Public cloud, community cloud, private cloud, hybrid cloud
- Public cloud offers
  - Amazon EC2, Microsoft Azure, Google Cloud Platform, …

© Image by storyset on Freepik

Research: cloud potential advantages

- Early access to state-of-the-art architectures essential for research
  - Rapid innovation cycles of modern HW architectures
  - New CPUs sometimes available in the cloud much earlier than acquirable on the market
  - Global shortages also affect us (e.g., COVID-19)
- “Immediate” availability of additional resources
  - No physical upgrades necessary
  - Rapid scaling
  - Access to different IT services (technology, tools, storage)
- Using cloud computing as a research vehicle

© Image by storyset on Freepik

Research: AWS – HPC in the cloud (1/2)

Using AWS Parallel Cluster for HPC in the cloud

- Build your own virtual cluster
- EC2 instances
  - Different flavors of CPUs (AMD, Intel, ARM, …)
  - Different types of accelerators/machine-learning processors (GPUs, FPGAs, …)
  - General or compute/memory/storage/accelerated-computing optimised
  - From small instances (1 vCPU) to many high performance cores
- Pre-installed operating systems (Linux and Windows)

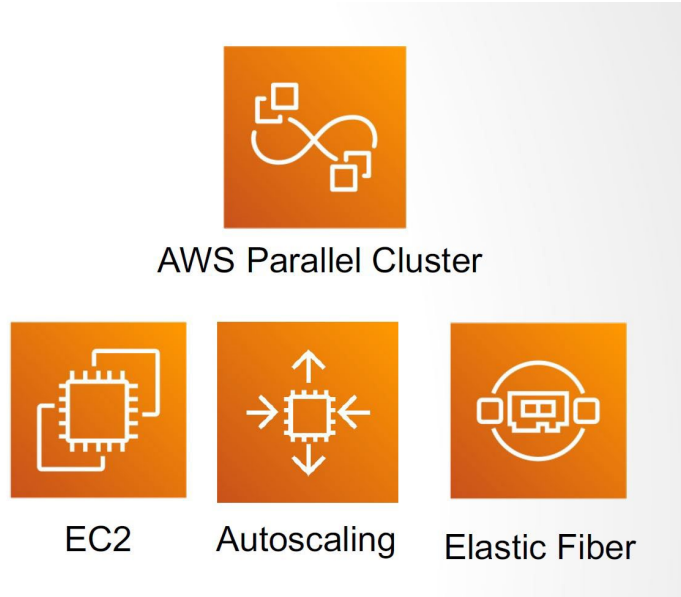

Research: AWS – HPC in the cloud (2/2)

Using AWS Parallel Cluster for HPC in the cloud

- High-performance file system
- Different node communication options available
  - Fast interconnect (EFA – up to 100 Gigabit network)
  - Essential for scaling
- Automatic resource scaling
- Easy cluster management

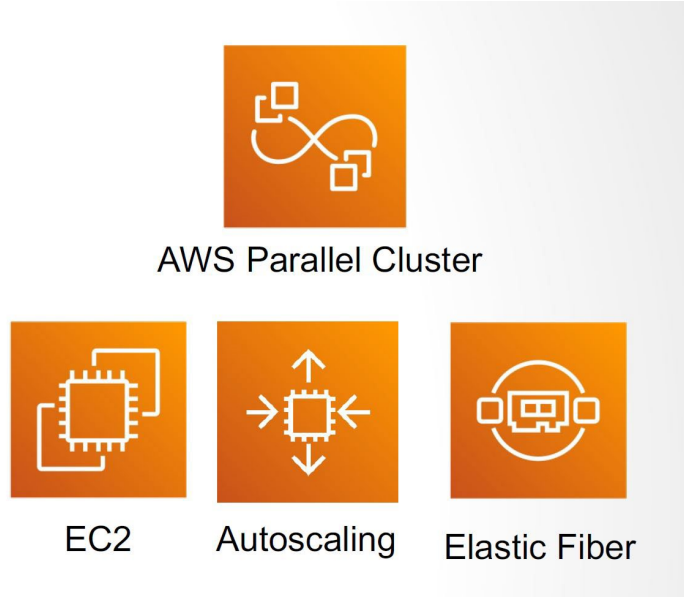

aws.amazon.com/hpc/parallelcluster/

Research: AWS – HPC in the cloud

- Relatively easy to configure
  - With scripts and configuration files
  - Web UI
- Workload management software (Slurm)
- Environment
  - Similar to a regular HPC cluster
  - Working on frontend and submitting parallel jobs (applications) for execution
- Node instances scale as requested (via Slurm)
- Nodes automatically released after 10 minutes of idle time

```
Region: eu-central-1
Image:
  Os: centos7
HeadNode:
  InstanceType: t3.small
...
Scheduling:
  Scheduler: slurm
  SlurmQueues:
  - Name: q1
    ComputeResources:
    - Name: g4dn.nmetal
      InstanceType: g4dn.metal
      MinCount: 0
      MaxCount: 16
      Efa:
        Enabled: true
    Networking:
...
```

Research: AWS – HPC in the cloud (1/2)

- Machine
  - 8 nodes: 384 CPU cores, 64 GPUs
  - Cloud: ~80 USD per hour + data transfers, storage, …
  - On-premises: ~180.000 USD + electricity, maintenance, …
- Experimental usage (Computer Science)
  - Short performance evaluations and scalability studies
  - We typically need roughly 2 hours for our experiments
  - Cost for 10 experiments (8 nodes): ~1.600 USD

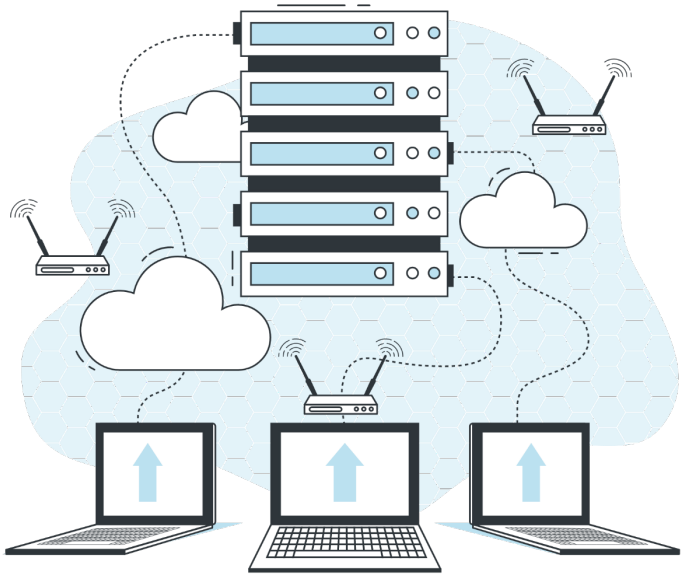

© Image by storyset on Freepik

Research: AWS – HPC in the cloud (2/2)

- Production usage (Domain Science)
  - Long running simulations/applications running days or weeks
  - Cost: ~1.920 USD / day (~1.300 with 1 year commitment)
  - For >100 days it may be cheaper to by the cluster

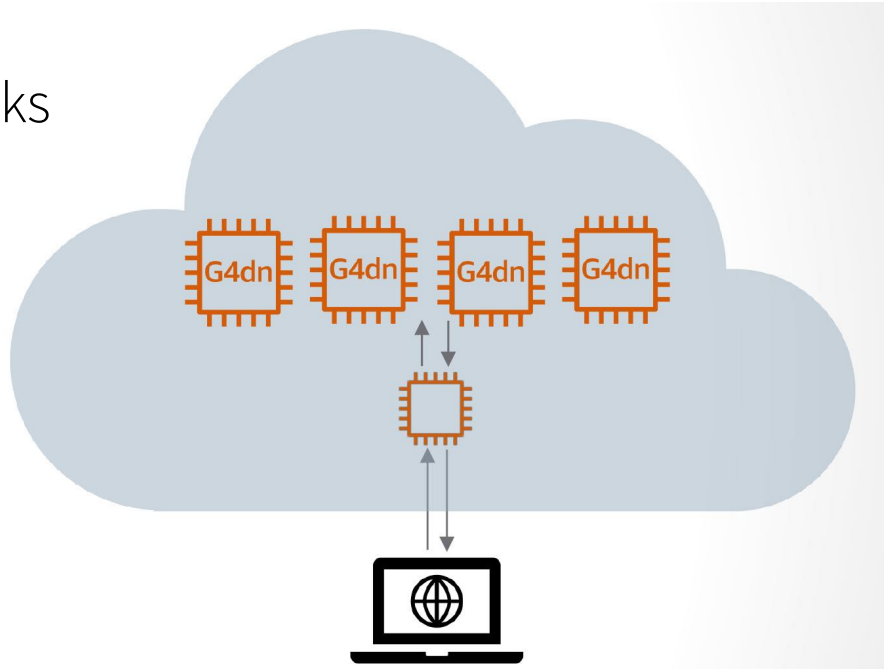

Research: summary

- Cloud computing as necessary driver for research
  - Many potential advantages in our use case:
    - Early access to architectures
    - Rapid scaling
    - Reduced costs
- However, not all use cases are viable in practice
  - Running long simulations on the newest hardware may be expensive
  - Scaling can be limited (depending on the usage scenario)
    - Illusion of unlimited power(?)
    - e.g., acquiring a cluster in a single region (the same data centre/rack) may not always be possible
  - Security, data privacy, …

© Image by storyset on Freepik

Teaching: cloud computing

- Master course “Cloud Computing” (curriculum Computer Science, 70+ students and increasing)
- Core concepts and principles of cloud computing
  - Topics: infrastructure, microservices, Kubernetes, serverless, message brokers, cloud storage, …
- Development of cloud-based applications and infrastructures
  - Provisioning and maintaining virtual machines and other hardware resources
  - Working with public clouds and state-of-the-art cloud services
  - Moderate to advanced use cases as assignments (dealing with microservice application that uses Kubernetes, serverless, machine-learning in cloud, Apache Kafka, distributed storage algorithms)

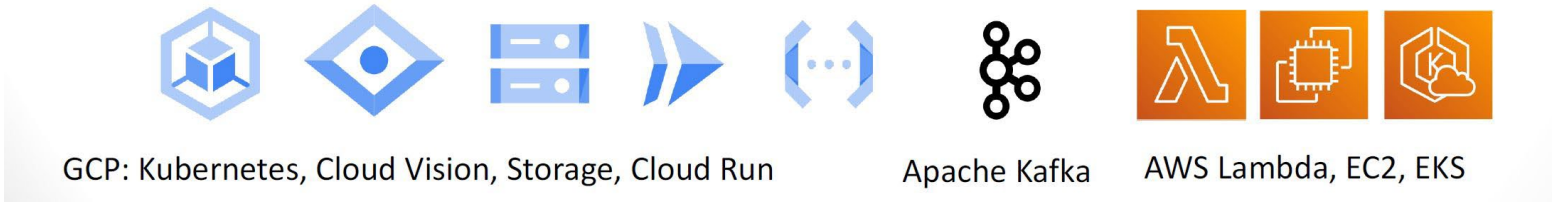

Teaching: cloud computing

- How to use the cloud?
  - Hundreds of services offered by each cloud provider (AWS: 227)
  - How to develop software that uses cloud services or runs fully in the cloud?
  - Cost estimations, vendor considerations, security, data privacy, …

- Understanding the cloud
  - How selected cloud services work behind the scenes
  - What are the software technologies and hardware architectures behind it
  - Virtualisation of IT infrastructure

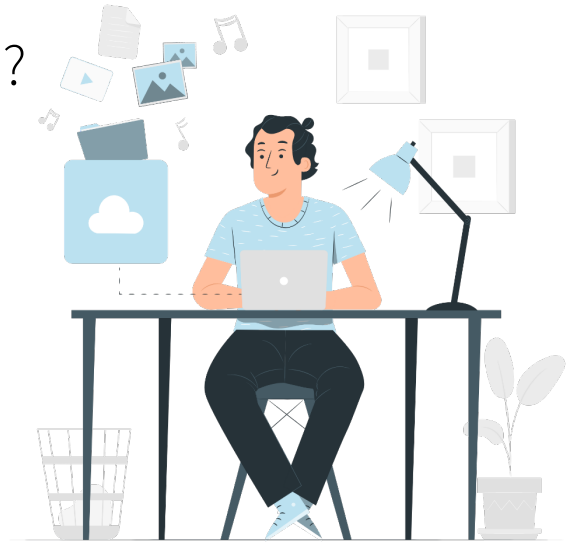

© Image by storyset on Freepik

Teaching: cloud computing (1/2)

Private cloud: In-house infrastructures

- To some extent possible
  - e.g., OpenStack, Cloud Stack, Kubernetes using in-house clusters, …
- Not a perfect fit for teaching
  - Often not scalable
  - Not up to date with latest hardware
  - Not up to date with latest software and cloud solutions

© Image by storyset on Freepik

Teaching: cloud computing (2/2)

Public clouds: Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP)

- Credits for students?
  - Some limited options are available (e.g., GCP)
  - Promising collaboration with ZID for the future
- Many potential advantages for classes
  - For the cloud computing topics
  - For other Computer Science classes (e.g., parallel computing, HPC, parallel architectures and programming models, …)
  - For students: experiences with real-world use scenarios, very valuable for professional careers

© Image by storyset on Freepik

Teaching: cloud computing (1/2)

- Practical exercise: large scale data stream processing in the cloud
  - Resembles many real-world big data applications
  - A virtual cluster on a cloud running Kubernetes engine
  - A microservice application running on Kubernetes
  - Machine learning algorithms (e.g. face detection, recognition)
  - Serverless and containerless
  - Handling streams of data
  - Different architectural styles

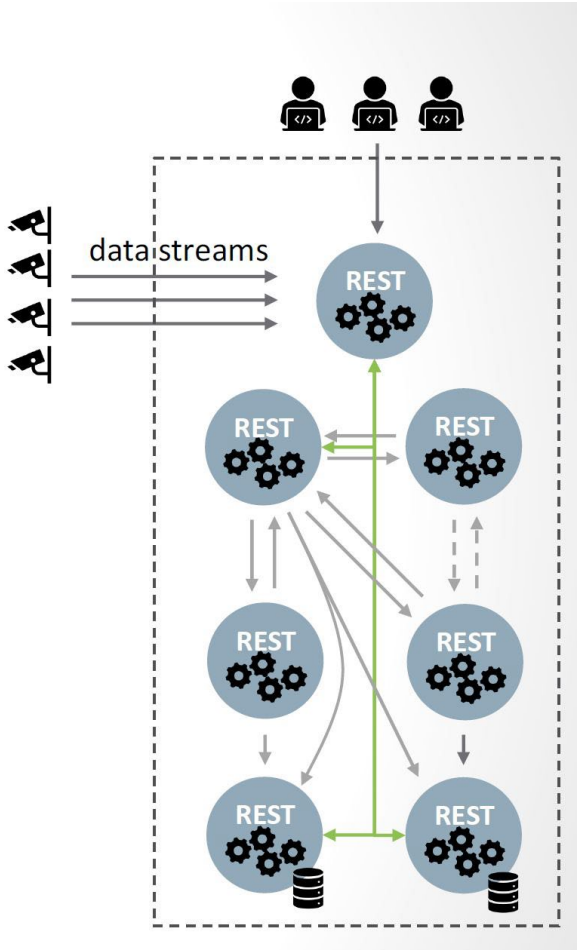

Teaching: cloud computing (2/2)

- Students need to deal with many different issues
  - Networking, service discovery, communication and data flow
  - (Auto) scaling up and scaling down based on demand
  - Costs analysis
  - Different approaches to build software

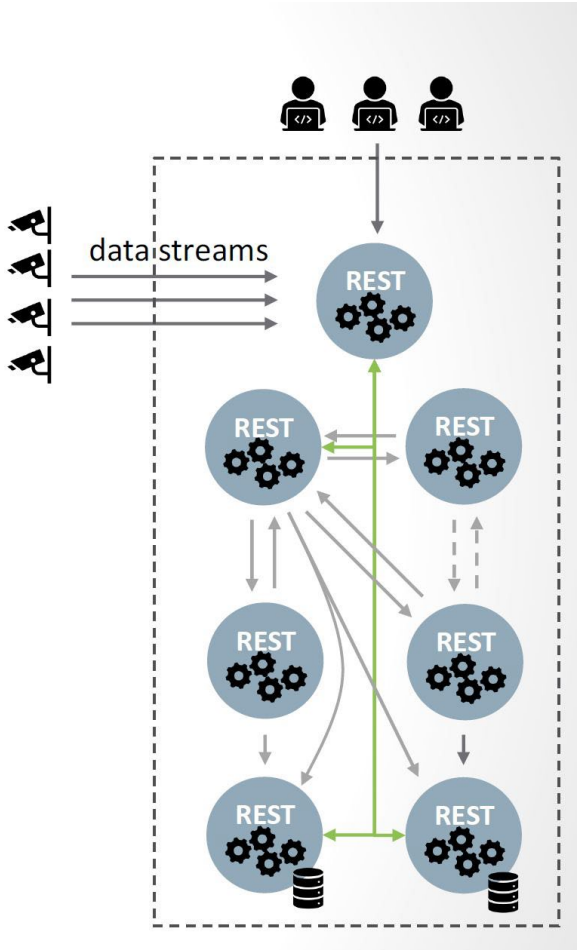

Teaching: cloud computing – challenges (1/2)

- Access to the cloud
  - How can we acquire cloud credits for students?
  - No reliable way to get credits (currently)

- How to control the costs and ensure cost limits?
  - Currently, each student gets a 25 USD voucher
  - Kubernetes: 4 nodes -> 16 vCPUs in total = ~0.6 USD each student per hour
  - ~20 hours to do the exercise and spend 50 % of the credits

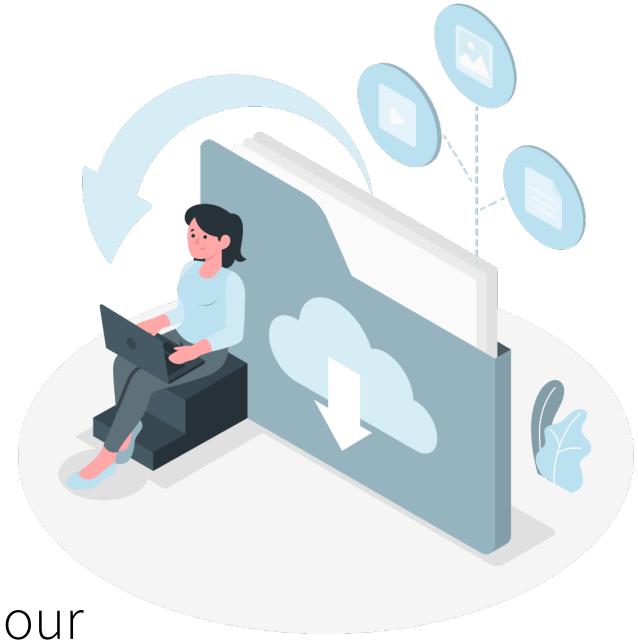

© Image by storyset on Freepik

Teaching: cloud computing – challenges (2/2)

Temporary workarounds

- Development on local machines, testing in the cloud (not ideal)
  - Working with Kubernetes on a laptop can be problematic due to resource constraints
  - Different environments may lead to a series of issues that distract from the problems that happen in the cloud
- Replacing traditional services gradually reduces costs

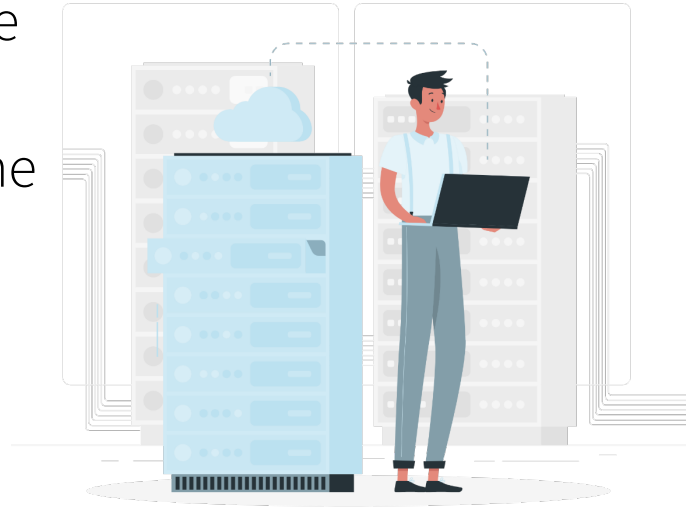

© Image by storyset on Freepik

Teaching: summary

- Cloud computing lecture
  - Using the cloud as a tool for computer science
  - Understanding how the cloud works behind the scenes
  - Students get a hands-on experience in cloud technologies, very valuable for industry

- Challenges
  - Cloud computing is extremely fast developing field
  - Providing access to cloud computing for students is a challenge
  - Cost control is an open issue

© Image by storyset on Freepik

Questions about Cloud Computing for Teaching and Research?

- Please ask your questions as a chat comment.

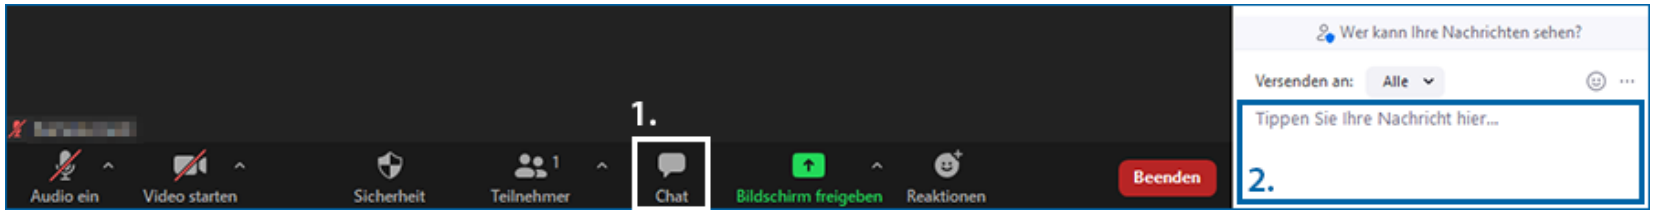

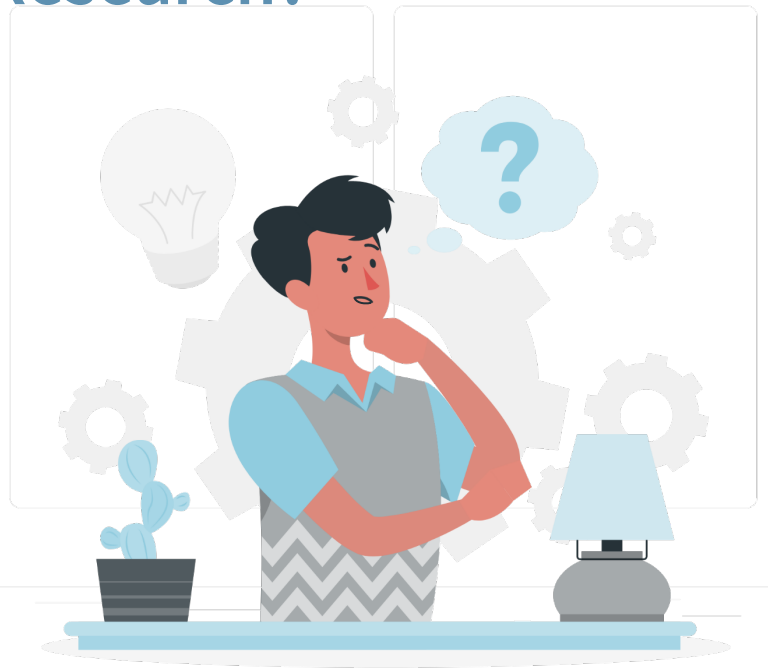

© Image by storyset on Freepik

universität wien

IT4S IT for Science FORUM

Overleaf – Collaborative Online LaTeX Editor

Presentation by Eva Karall

Why LaTeX and how does it work?

- Developed by Leslie Lamport in the 1980s, LaTeX came to its name: **La**mport **TeX**
- A markup language that facilitates the production of well-formatted documents

```
This is a simple math expression \(\sqrt{x^2+1}\) inside text.
```

This is a simple math expression $\sqrt{x^2+1}$ inside text.

What is Overleaf?

- Overleaf is an online LaTeX editing tool that
  - allows you to create LaTeX documents directly in your web browser
  - “compiles” your LaTeX automatically to show you the results

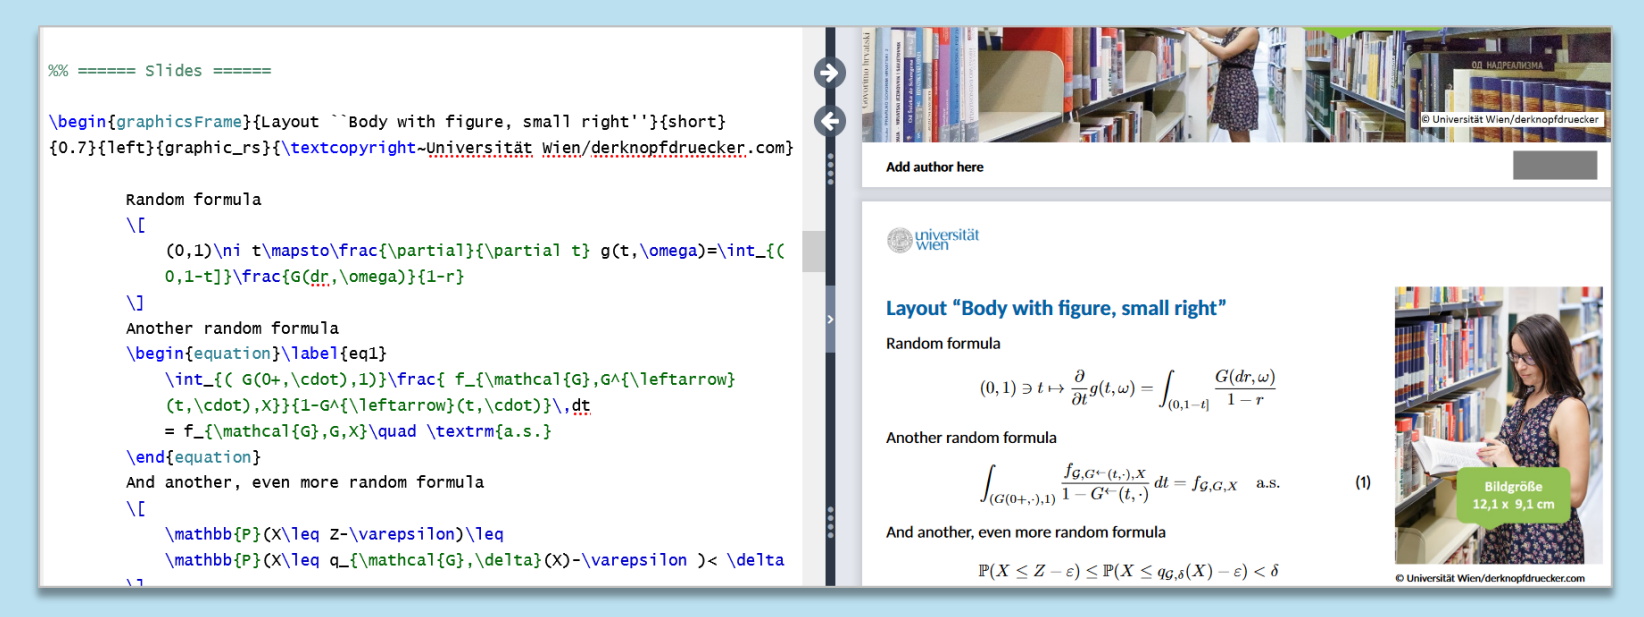

Which license do I get?

- **Students** of the University of Vienna receive an **Overleaf Free license**:
  - Unlimited number of projects
  - Collaborating in real time with another person
  - Using templates

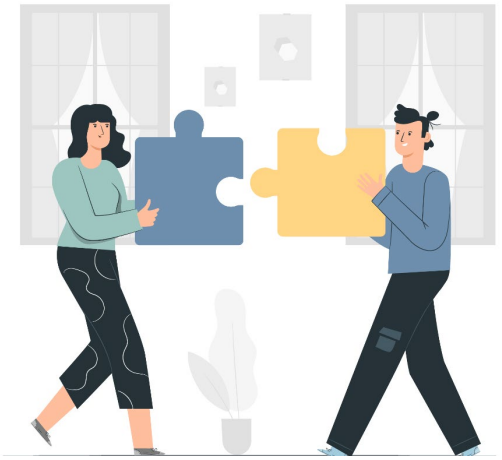

© Image by storyset on Freepik

- **Employees** of the University of Vienna receive an **Overleaf Professional license free of charge:**
  - In addition to the features of the Free license, this enables:
    - tracking changes in real time
    - unlimited number of authors for documents
    - full version history
    - synchronization with e.g. Git/GitHub

University of Vienna on Overleaf

- To start using Overleaf go to www.overleaf.com/edu/univie
- **How to get your license**:
  - Click on the button *Log in through your institution* below
  - Log in via weblogin with your u:account
  - Start new Overleaf project
    - starting from scratch
    - uploading your own files
    - or using one of the many templates available

Templates

- **Gallery** with thousands of templates
- Anyone can submit projects for possible inclusion in the gallery
- Templates available for thesis, CV or submission to **scientific journals** and conference, e.g. the IEEE] and Springer
- Official **University of Vienna templates** available

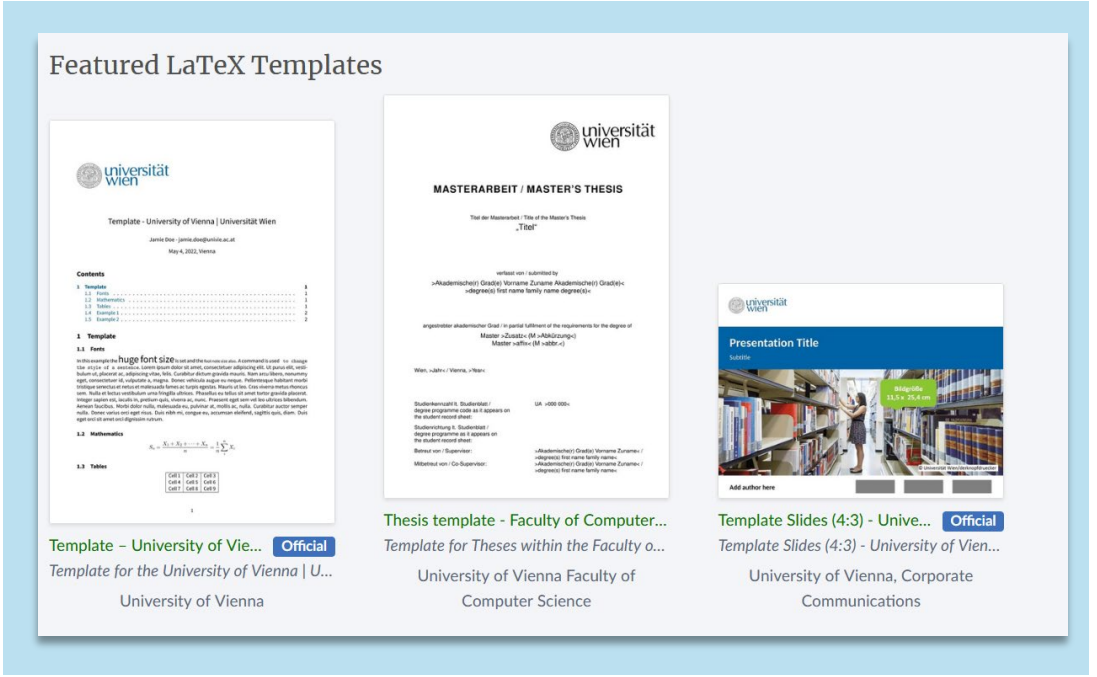

Sharing

- Every document you create on Overleaf is private by default
- Two ways to share your work with collaborators:
  - By private invitation
    - Private invitations allow you to invite selected, named collaborators to access your projects
    - the number of named collaborators you can invite on each project depends on your plan
  - By link-sharing
    - Link-sharing allows you to share your projects via secret links for viewing, commenting and editing

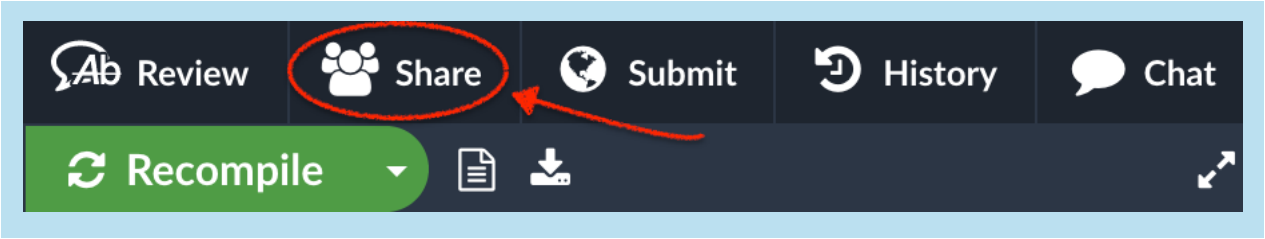

Screenshot: www.overleaf.com/learn/how-to/Sharing_a_project

Collaboration

- Single master version of each document
- Simultaneous collaborative editing

- Discussing within Overleaf
- Tracking all changes
- Keeping one to-do list
- Accepting and rejecting changes

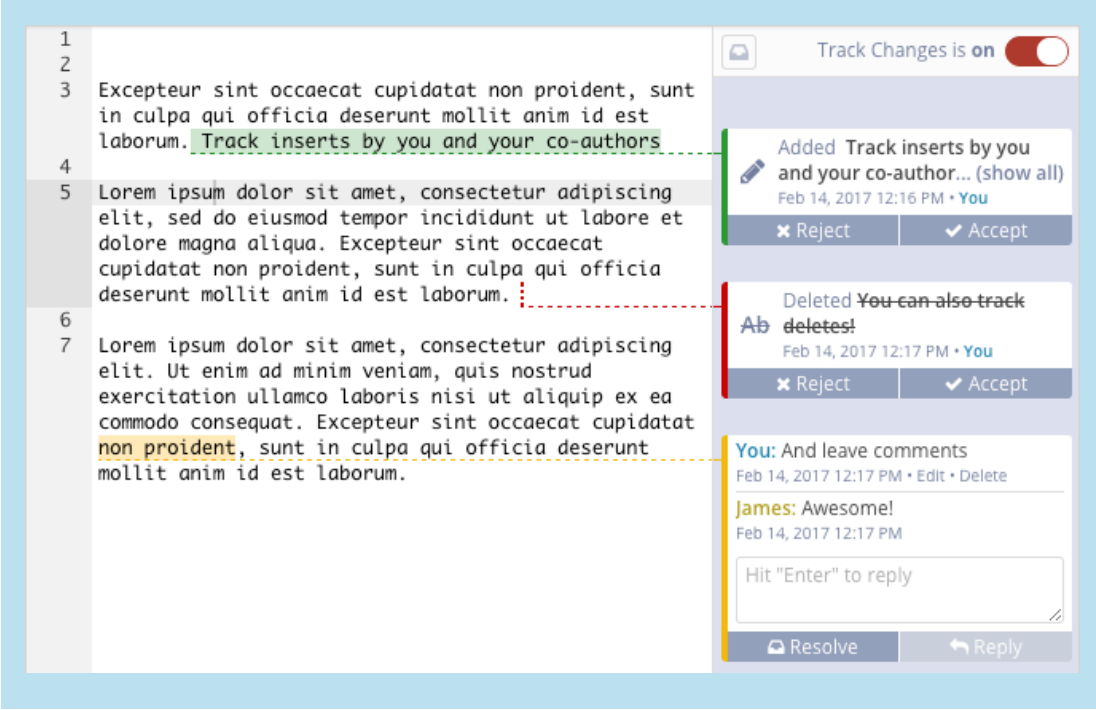

Screenshot: https://www.overleaf.com/blog/track-changes-and-comments-in-latex-2017-03-09

More information

- Service page – zid.univie.ac.at/en/overleaf/
- University of Vienna on Overleaf – www.overleaf.com/edu/univie

- **Instructions and support via Overleaf**
  - Documentation of Overleaf
  - Helpful how-to guides
  - Technical support via contact form or e-mail to support@overleaf.com

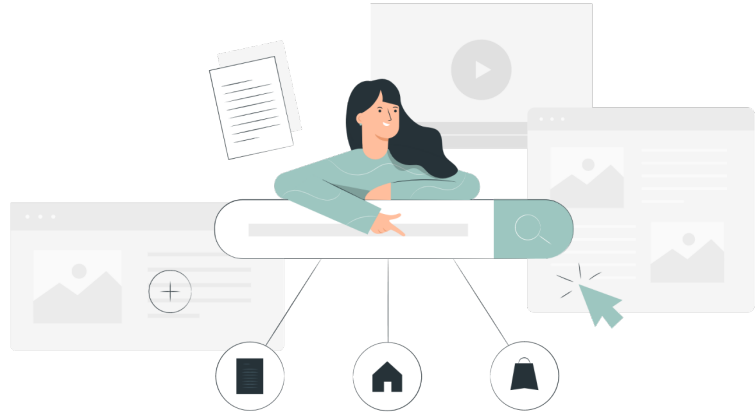

Questions about Overleaf – Collaborative Online LaTeX Editor?

- Please ask your questions as a chat comment.

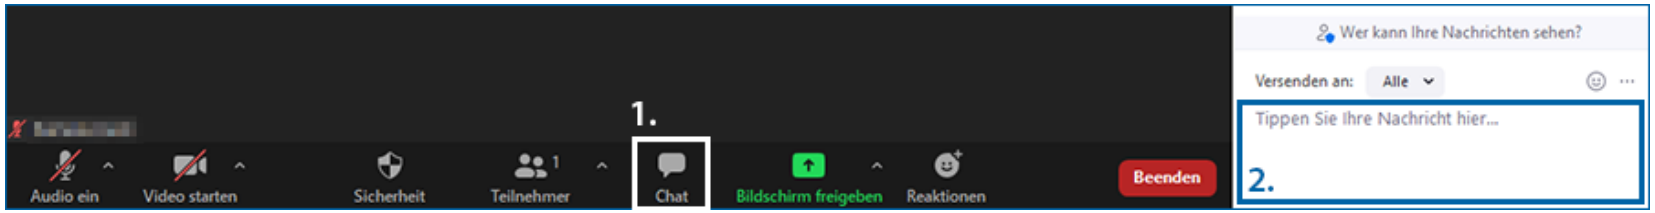

Outlook

universität wien

IT4S IT for Science FORUM

Thank you for your attention!

© Image by storyset on Freepik