



Computing

Clouds

Cloud Computing

AMIs

Workshop AMI

Guy Leonard

Workshop on Genomics
2025

“All Watched Over By Machines Of Loving Grace”

Richard Brautigan



Computing

Clouds

Cloud Computing

AMIs

Workshop AMI

Guy Leonard

Workshop on Genomics
2025

AMIs, the Cloud, and Computing

Fundamentals

The Human Interface to the Digital World

- ⇒ You!
- ⇒ Experience computing through:
 - GUIs
 - Touch and Voice
 - CLIs
- ⇒ Technology exists to extend human capabilities
- ⇒ User Experience (UX), Human-Computer Interface (HCI) and Accessibility



Fundamentals

The Human Interface to the Digital World

- ⇒ How many computers are in this room?
 - ~70 laptops?
 - ~70 phones?
 - Tablets?
 - Smart Watches?
 - Routers
 - Other 'smart' devices?
 - Other microcontrollers/chips?
- ⇒ Ubiquitous!
 - But do you know how they work?



Fundamentals

Bridging User Needs and System Capabilities

- ↪ Directly serves user requirements
- ↪ Productivity Tools
 - MS Office
- ↪ Creative Applications
 - Photoshop
- ↪ Scientific and Research Software
 - The next two weeks...
- ↪ Entertainment and Gaming

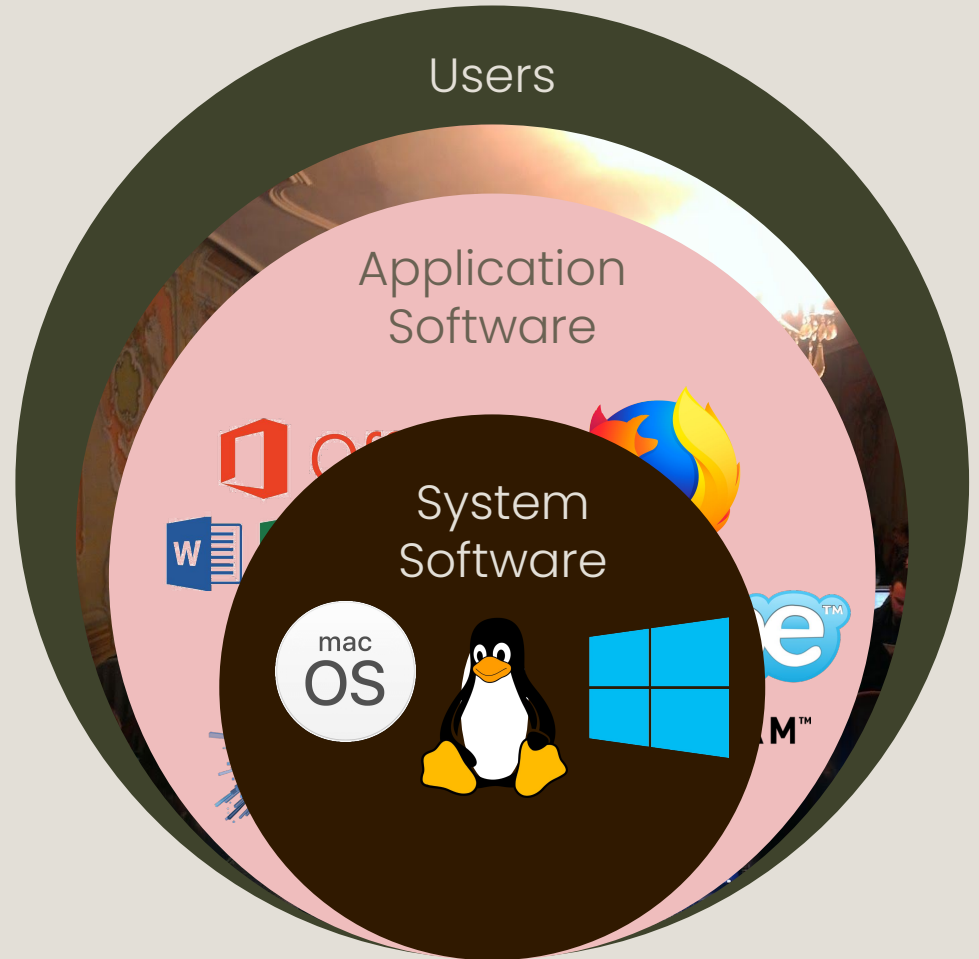
- ↪ Translates user intentions into computational actions



Fundamentals


The Computational Translator

- Acts as an intermediary between hardware and application software
- Key Components:
 - Operating Systems
 - Device Drivers
 - Utility Programs
 - Runtime Environments
- Primary Functions:
 - Resource Management
 - Process Scheduling
 - Memory Allocation
 - Security and Access Control
- Compilers, Linkers and Assemblers



Fundamentals


The Physical Foundation



Fundamentals

The Physical Foundation



- ↪ Key Components:
 - Central Processing Unit (CPU)
 - Random Access Memory (RAM)
 - Storage Devices
 - Input/Output Devices
 - Network Interfaces
- ↪ “fetch-decode-execute” cycle to process program instructions
- ↪ Computer performance depends on cache size, clock speed and the number of cores



Fundamentals

→ CPU


- **Control Unit**
 - “fetch-decode-execute”
 - Transfers data and instructions around the system
- **Arithmetic Logic Unit (ALU)**
 - arithmetic and logical operations
- **Cache**
 - small amount of high-speed random access memory (RAM)
- **Clock**
 - Coordinates all the components, measured in Hertz (Hz)



CPU vs GPU

→ Central Processing Unit

- Core computational operations
- Low compute density
- Complex Control Logic
- Large Cache
- Optimised for serial operations



→ Graphical Processing Unit

- Specialised computational operations
- High compute density
- High computation per memory access
- Built for parallel operations



CPU: Cores vs. Threads


↳ What is a CPU Core?

- A core is an independent processing unit within a CPU.
- Each core can execute its own instructions simultaneously.
- Early CPUs had only one core, but modern CPUs have multiple cores for parallel execution.
- More cores = better performance for multi-threaded applications.


↳ What is a Thread?

- A thread is a sequence of instructions executed by a CPU.
- Each physical core can run one or more threads.
- Single-threaded applications: Run on one core at a time.
- Multi-threaded applications: Can run across multiple cores for improved performance.

Single Core CPU





Dual Core CPU



Mainframes & Supercomputers

- In 1978, the Cray-1 supercomputer cost \$7 million, weighed 4762 kg and had a 115 kilowatt power supply. 1 CPU.
- The Raspberry Pi costs around \$70 (CPU board, case, power supply, SD Card), weighs 50g uses a five watt power supply and is more than 4.5 times faster than the Cray-1. 1 CPU = 4 Cores





Virtual Computing and Terminals

- ⇒ How did users access the mainframe computers?
 - Terminals!
 - Originally, a physical device (teletypewriters “TTYs”) that connected to a remote computer.
 - The term “terminal” comes from it being the endpoint of a session on a mainframe.
 - Allowed users to interact with powerful shared systems (early UNIX mainframes, minicomputers).
 - Eventually replaced by terminal emulators running inside modern operating systems.



Virtual Computing and Terminals

- What were the mainframes running?
- Rise of UNiplexed Information and Computing System (Unix)/Linux:
 - 1960s-1970s: UNIX was developed at Bell Labs.
 - 1980s: UNIX becomes the foundation of many academic and enterprise systems.
 - 1990s-Present: Linux, an open-source UNIX-like OS.



Virtual Terminals

- With the advent of personal computers and super computers, terminals were no longer useful.
 - Telnet (1970s-1990s) – Early network-based remote access (unencrypted, insecure).
 - SSH (1995-present) – Secure, encrypted alternative for remote terminal access.
- What is the Shell?
 - The shell is a command-line interpreter that sits between the user and the operating system.

User Space

Remote Space



What is The Cloud?

- ⇒ A shift from local to remote computing
- ⇒ Key characteristics:
 - On-demand resource allocation
 - Pay-as-you-go models
 - Scalability and flexibility
- ⇒ How old is it?
 - 1960s!
 - Data centres at DARPA
 - 1970s: IBM Time sharing
 - 1990s: VPNs
 - 2000s: Amazon etc

Is this it?



* The only generative AI content in these slides...



The Cloud

It's these things...

→ **You have all been using it for most of your lives...**

- Social Media
- Email
- Storage
- Computing

→ **Amazon Web Services**

- Early 2000s - Scalability issues running e-commerce
- Decouple components - sped up development & reduced bottlenecks
- Infrastructure as a service: 2006+ = EC2, S3 = Cloud!



Data Centres

- ⇒ Cloud servers are located in data centers all over the world
- ⇒ Amazon Web Services
- ⇒ Microsoft Azure
- ⇒ Google Compute
- ⇒ Other regions may have their own solutions...




Data Centres

- Cloud servers are located in data centers all over the world
- A data center is a facility housing many networked computers that work together to process, store, and share data
- By using cloud computing, users and companies do not have to manage physical servers themselves or run software applications on their own machines



AWS, AMI and EC2

- ↪ **Amazon Web Services (AWS)**
 - comprehensive cloud platform offering on-demand services
- ↪ **Amazon Elastic Compute Cloud (EC2)**
 - core compute service within AWS, providing resizable virtual servers on demand
- ↪ **Amazon Machine Image (AMI)**
 - a preconfigured virtual template used to launch EC2 instances, containing an operating system



AMI, Instance, Virtual Machines?!

↪ **Virtual Machine (VM)**

- A software-based computer that runs an entire operating system (OS) in an isolated environment.
- The VM shares the physical server's resources (CPU, memory, storage), but appears to the user like a standalone machine.

↪ **Amazon Machine Image (AMI)**

- A template provided by Amazon Web Services (AWS) used to create EC2 instances.
- Contains the operating system, software, and initial configuration.
- Acts as the base "blueprint"; you launch instances from AMIs.

↪ **Instance (in AWS context)**

- A running virtual server in AWS, launched from an AMI.
- Receives compute resources (e.g., CPU, RAM, storage) allocated by AWS.
- You can start, stop, reboot, or terminate an instance as needed.

Connecting to the Cloud



Remote Desktop
(Guacamole)



Secure
Shell (ssh)


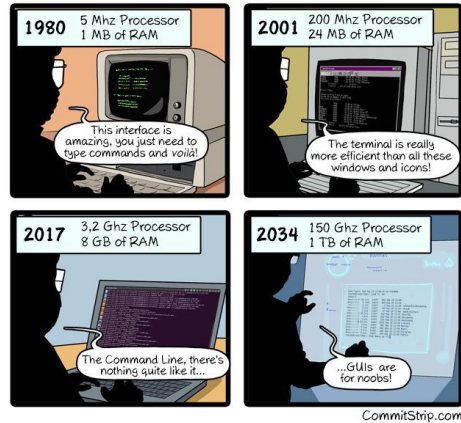



R Studio
Server



Jupyter
Notebook

Takeaways



Connecting to your Instance

↪ <https://evomics.org/2025-workshop-on-genomics/>

EVOLUTION AND GENOMICS

Intensive and immersive training opportunities

WORKSHOPS

LEARNING

PEOPLE

APPLY

INFORMATION

Get ready for May 2023!

Many important details can be read on our [FAQ](#) page and our [Housing and Transportation](#) page

[Instance addresses](#) (check every day for your new one!)

Check [here](#) to view Faculty / Organiser / Instructor arrival and departure dates. Check out our Faculty for 2023 [here](#) and our Instructors for 2023 [here](#)

As is tradition, we will be having a [T-shirt competition!](#) Best T-shirt design will feature on the Workshop on Genomics 2023 T-shirts – get designing!

Workshop on Genomics 2023 [BINGO!](#)

Faculty lunches sign-up sheet

Our [Code of Conduct](#) contact points are Josie Paris & Joan Ferrer Obiol

Connecting to your Instance (AMI)





- Find your name and check your instance address
- Check how to connect to your instance using guacamole, ssh and the RStudio server

	A	B	C	D	E	F	G	
1	Instance number	First Name	Last Name	Instance address	Guacamole connection	ssh connection	RStudio server connection	
2	1	Joan	Ferrer Obiol	3.238.107.169	3.238.107.169:8080/guacamole	ssh genomics@3.238.107.169	3.238.107.169:8787	

Connecting through SSH

- Open your preferred terminal on your laptop
- Type `ssh genomics@[instance address]`
- Enter the password

```
[apples-MacBook-Pro.local@apple[~]$ ssh genomics@3.238.107.169  
genomics@3.238.107.169's password: [?]
```





Connecting through SSH

```
#####  
##          Workshop on Genomics 2023          ##  
##      🍷🍷🍷 Spring Edition 🍷🍷🍷          ##  
##          Cesky Krumlov          ##  
##          @evomics #evomics2023          ##  
#####  
  
Welcome to Ubuntu 22.04.2 LTS (5.19.0-1024-aws).  
  
System information as of Sun May 14 15:33:29 CEST 2023  
  
System load: 0.85107421875      Processes:          220  
Usage of /: 58.1% of 484.63GB   Users logged in:   0  
Memory usage: 29%              IPv4 address for docker0: 172.17.0.1  
Swap usage: 0%                 IPv4 address for ens5: 172.31.10.159  
Last login: Sun May 14 11:08:55 2023 from 194.228.207.170  
genomics@ip-172-31-10-159: [~]$ █
```

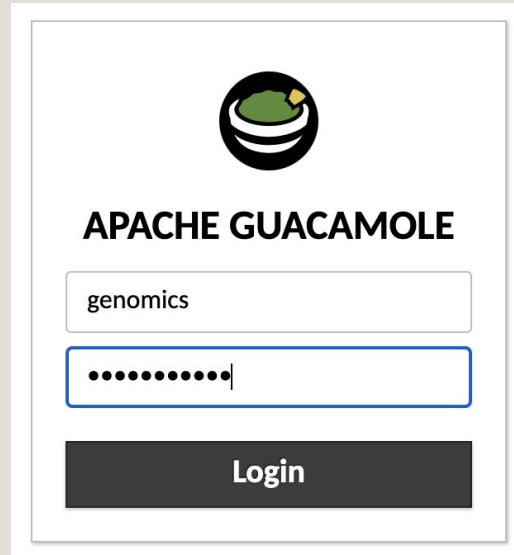
Connecting to Guacamole

- ↪ Open your preferred internet browser (i.e. chrome, firefox)
- ↪ Paste your instance address followed by :8080/guacamole



Connecting to Guacamole

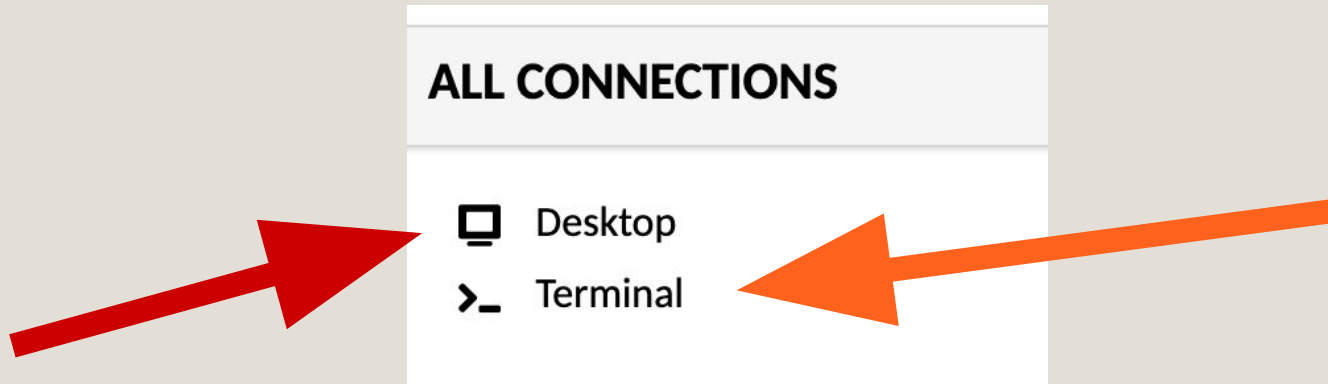
- Enter the username “genomics” and the password (on the flip-chart/whiteboard)



The screenshot shows the Apache Guacamole login page. At the top center is the Guacamole logo, which consists of a black circle containing a green bowl with a white spoon and a yellow pencil. Below the logo, the text "APACHE GUACAMOLE" is displayed in a bold, black, sans-serif font. Underneath this text are two input fields: the first is a white text box with a thin grey border containing the username "genomics"; the second is a white password box with a thin blue border, containing ten black dots and a vertical cursor. At the bottom of the form is a dark grey rectangular button with the word "Login" written in white, bold, sans-serif font.

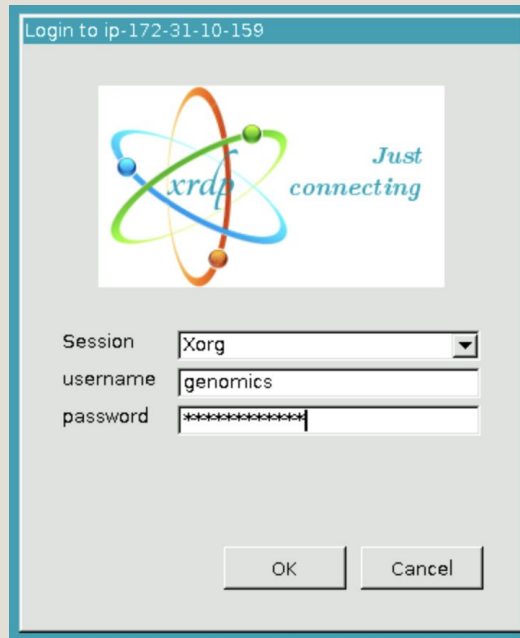
Connecting to Guacamole Desktop

- Open your preferred internet browser (i.e. chrome, firefox)



Connecting to Guacamole Desktop

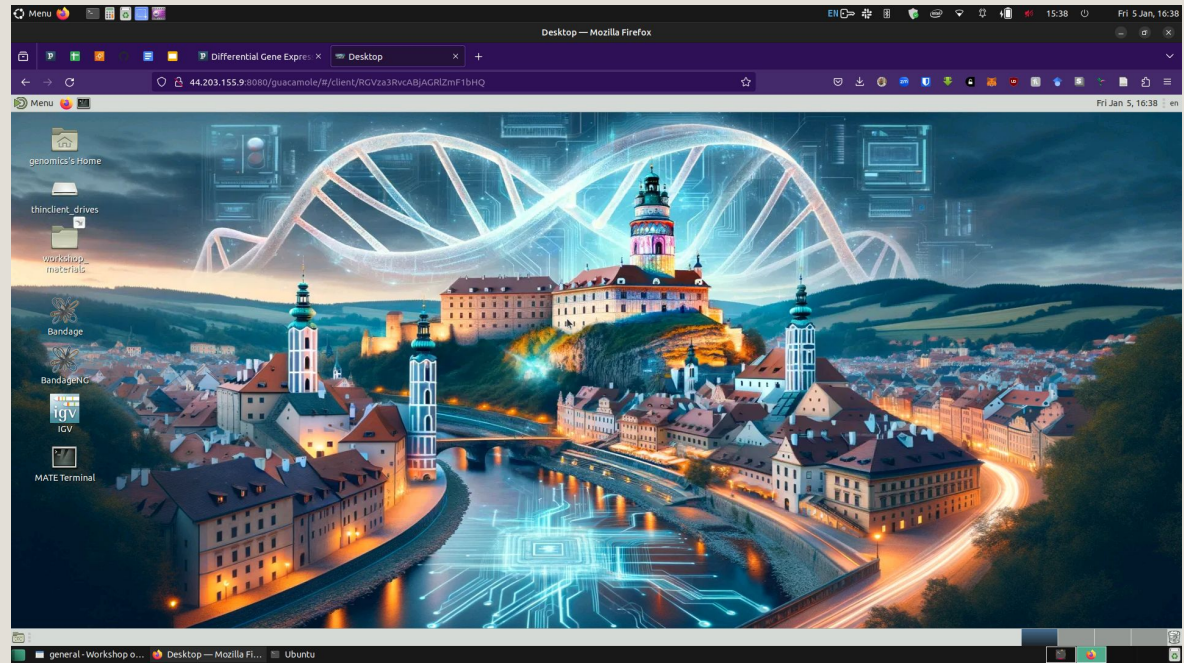
- Enter the username “genomics” and password again



The screenshot shows a login dialog box titled "Login to ip-172-31-10-159". The dialog features a central graphic with the text "xrdp" and "Just connecting". Below the graphic are three input fields: "Session" (a dropdown menu showing "Xorg"), "username" (a text field containing "genomics"), and "password" (a text field with masked characters). At the bottom of the dialog are "OK" and "Cancel" buttons.

Connecting to Guacamole Desktop


→ Open a terminal window using the terminal icon



Copying and Pasting

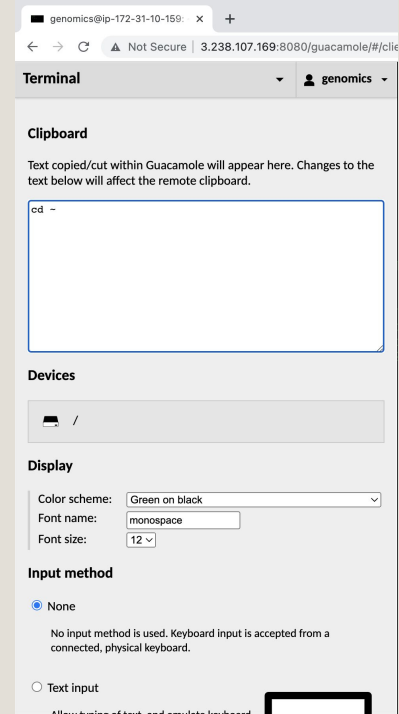


Copying and Pasting



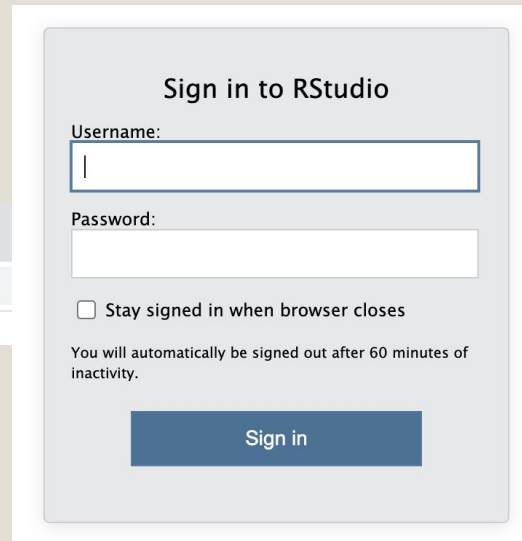

But if you need it:

- ↪ Press Ctrl+Alt+Shift (Mac: Ctrl+Opt+Shift)
- ↪ Paste the text in the pop-up box
- ↪ Press Ctrl+Alt+Shift (Mac: Ctrl+Opt+Shift) again
- ↪ Paste into the instance using right click



Connecting to R-Studio Server

- ⇒ Open your preferred internet browser (e.g. Chrome, Firefox)
- ⇒ Paste your instance address followed by :8787
- ⇒ Username: genomics
- ⇒ Password: On the whiteboard in the Prelate

A screenshot of the RStudio Sign In form. The form is titled "Sign in to RStudio" and contains the following elements:

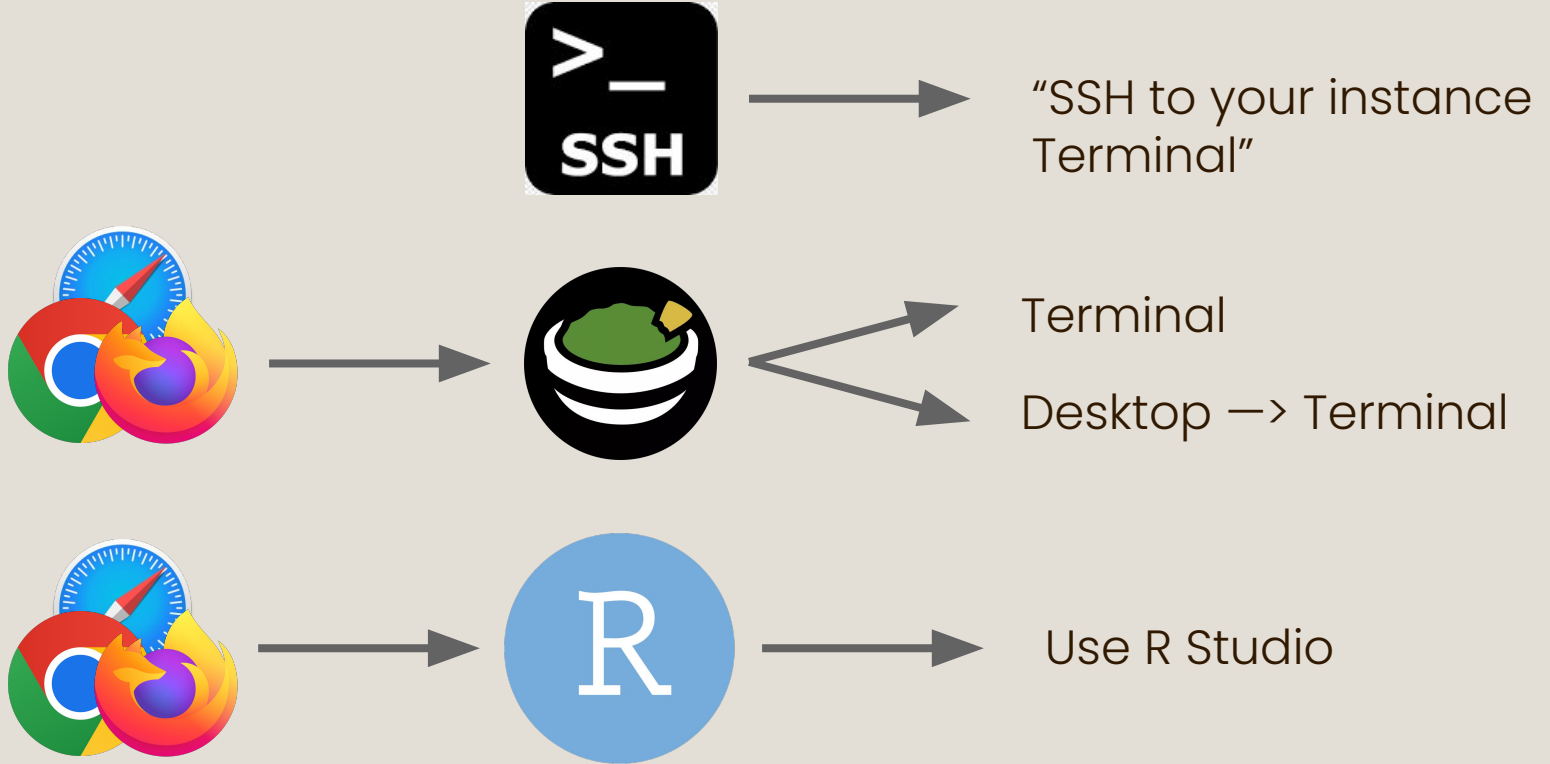
- A "Username:" label followed by a text input field.
- A "Password:" label followed by a password input field.
- A checkbox labeled "Stay signed in when browser closes".
- A warning message: "You will automatically be signed out after 60 minutes of inactivity."
- A blue "Sign in" button.

Your Daily AMI



- ↪ The instance IP address will **change** every day after we stop and restart the instances.
- ↪ Each morning, you will need to return to the “Instance address spreadsheet” on the webpage, retrieve your new address and login again using the type of connection you need

I forgot how do I connect?



Keyboard Troubles!


↪ **Make sure that you can type the following characters:**

↪ tilde (~)

↪ backslash (\)

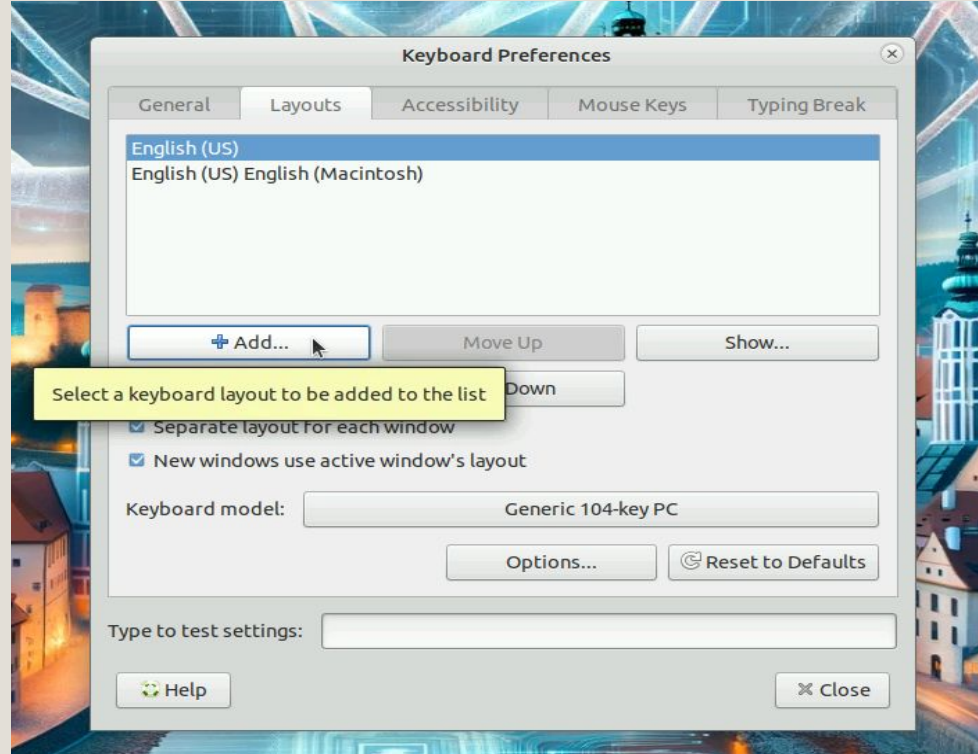
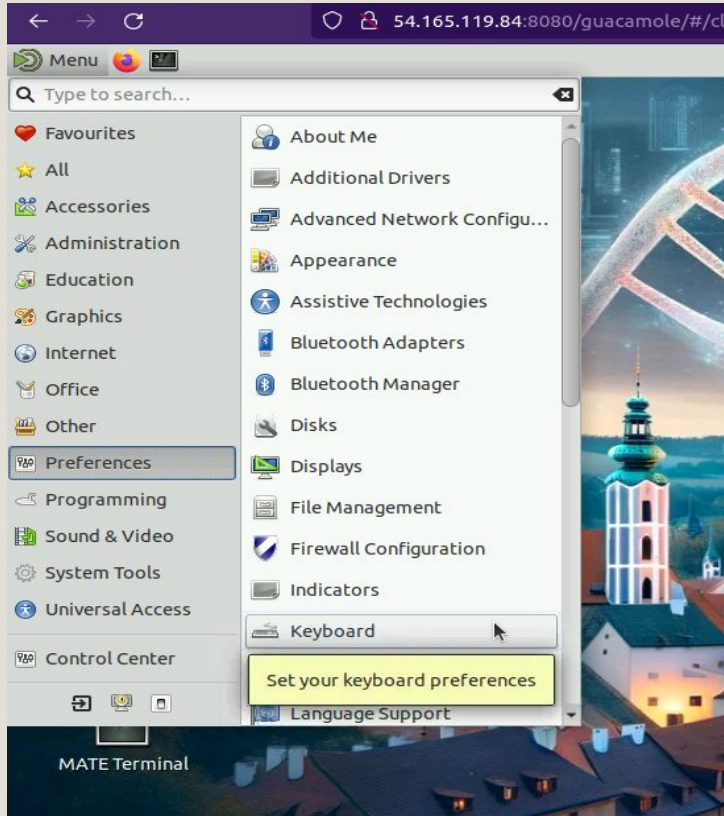
↪ pipe (|)

↪ carat (^)

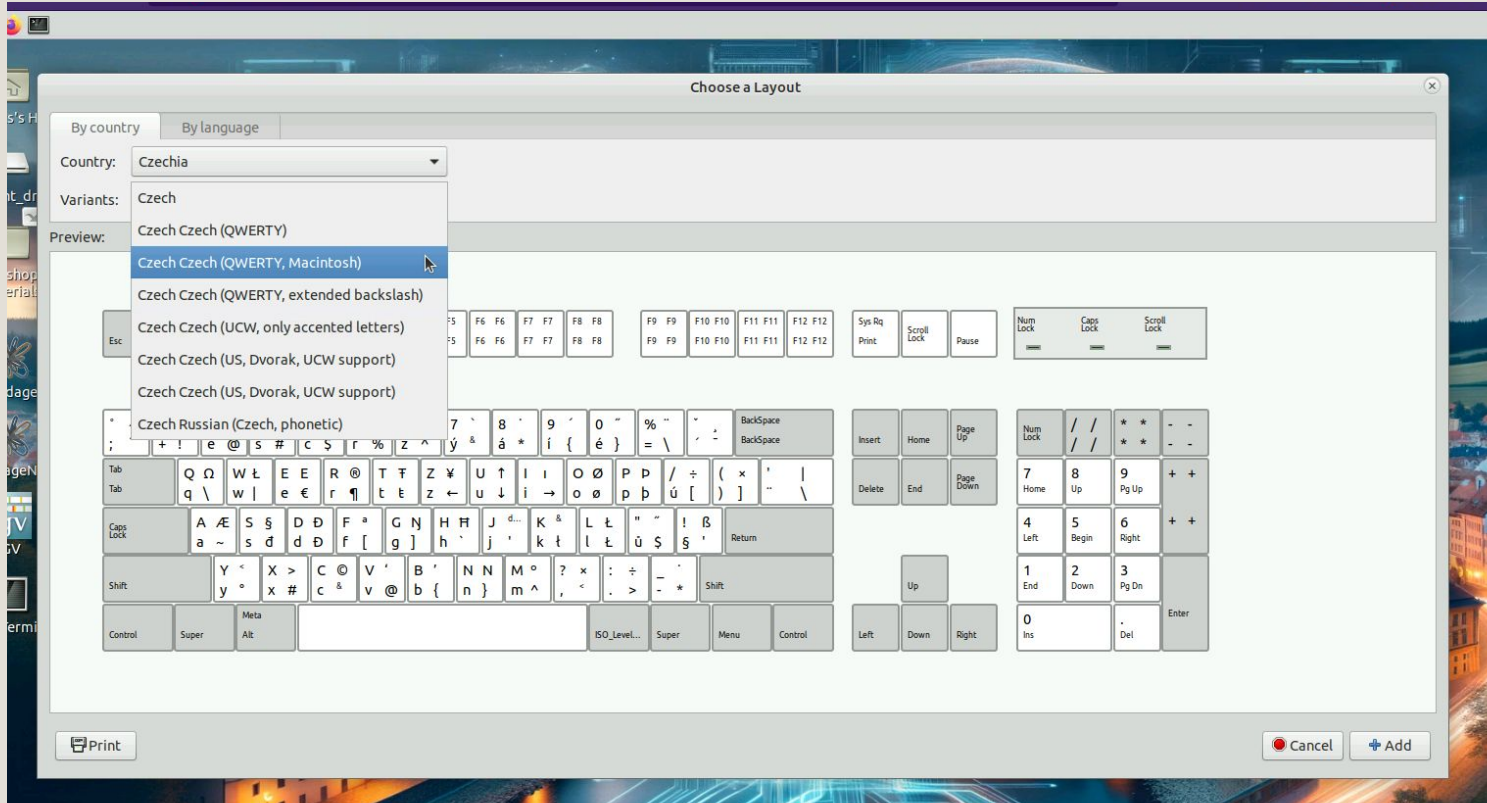


↪ If you can't type these characters, please get our attention!

Keyboard Troubles!



Keyboard Troubles!



Tilde ~



tilde

/'tɪldə, 'tɪldi/



Backslash



backslash
"bakslaf/



Pipe |

vertical bar
pipe



Caret ^



Caret
/kɛrət/

