The Research Desktop

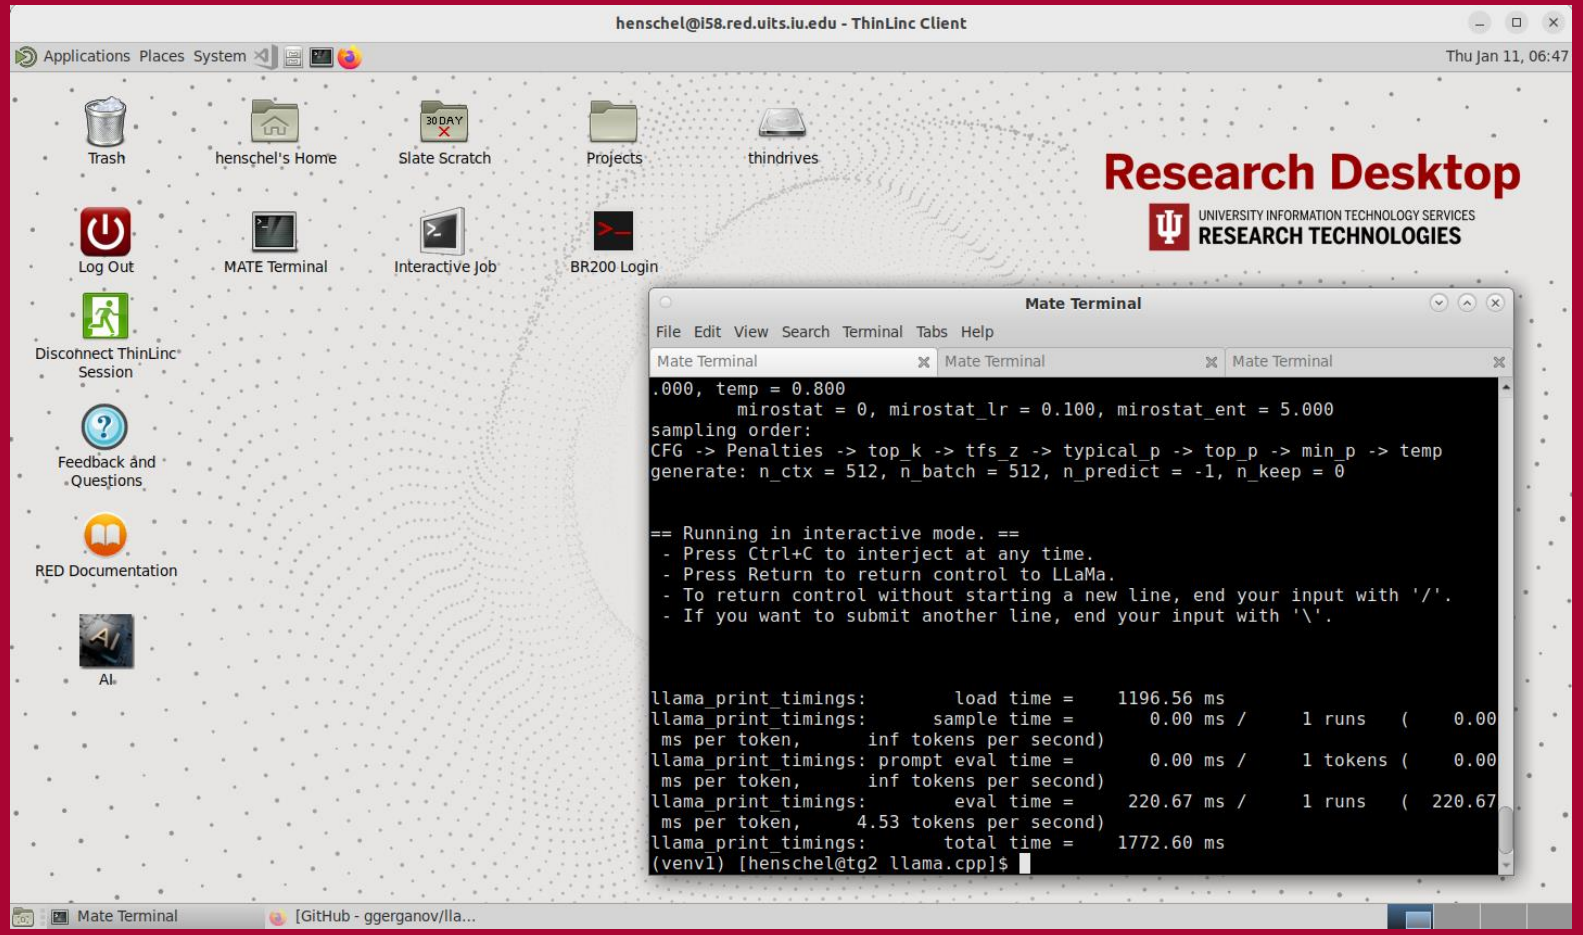

Robert Henschel

Project Director, Research Engagement

Research Technologies

henschel@iu.edu

UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Thank you!

Sunita and Samantha

Questions:

- Have you used a Linux desktop?
- How have you used HPC? (researcher, operator, educator,…)

The next hour…

- Introduction
- Motivation for creating a Research Desktop
- Definition and History of the Research Desktop
- Current State and Future Developments

- Q&A and Demo

Introduction

- Robert Henschel
  - At Indiana University since 2008.
  - Project Director for Research Engagement.
  - Architect of the IU Research Desktop.
  - Working 60% at IU; 40% HPC and Research Computing consulting.
  - Currently consulting for Cendio, maker of ThinLinc.
- Trying to cover everything that I promised in the abstract of the talk, but always open for questions!

Indiana University

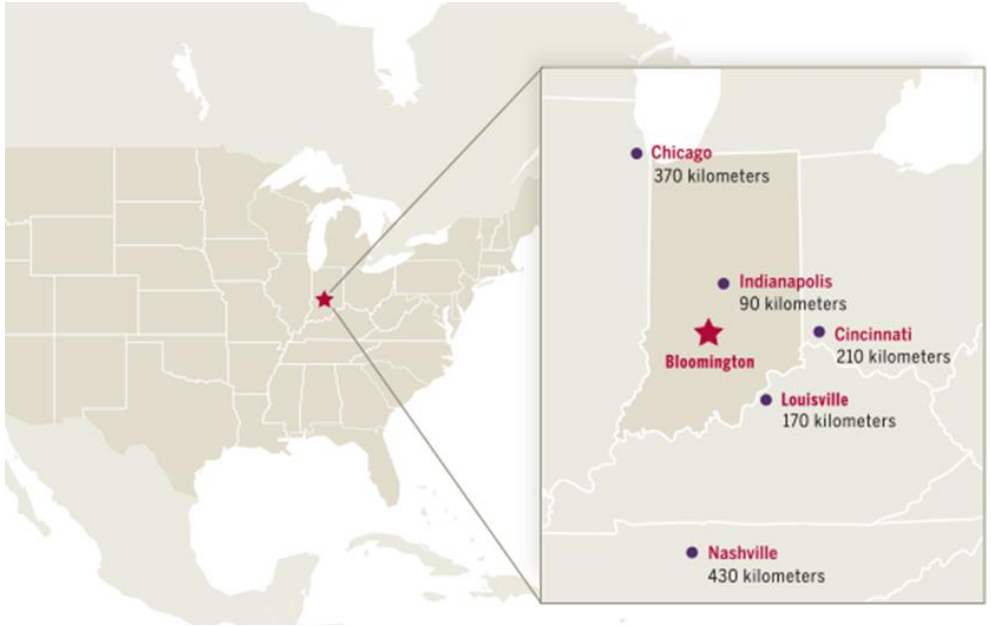

Fall 2022	Number
Research Campuses	2
Regional Campuses	7
Undergraduate	71,000
Graduate + Doctoral	19,000
Total Students	90,000
Staff	11,000
Faculty	9,000
Grand Total	110,000

Overall

- Founded in 1820
- Annual operating budget - $4B
- Grant Awards of $772M in 2023

Centralized IT - UITS

- 650+ professional staff
  - 100 Research Technologies
- 500+ part time staff

HPC @ Indiana University

- **Quartz** – HPE/Cray system; 92 dual CPU nodes; 36 4-way V100 GPU nodes
- **Big Red 200** – CRAY Shasta system; 640 dual CPU nodes; 64 4-way A100 nodes
- **Research Desktop** – Dell servers; 34 dual CPU servers

- **Jetstream 2** – NSF's second production cloud environment, main system at IU and small regional clouds in 4 regions. 8 petaflops of VM capacity (CPU + GPU)

The next hour…

- Introduction
- Motivation for creating a Research Desktop
- Definition and History of a Research Desktop
- Current State and Future Developments

- Q&A and Demo

History of Supercomputing

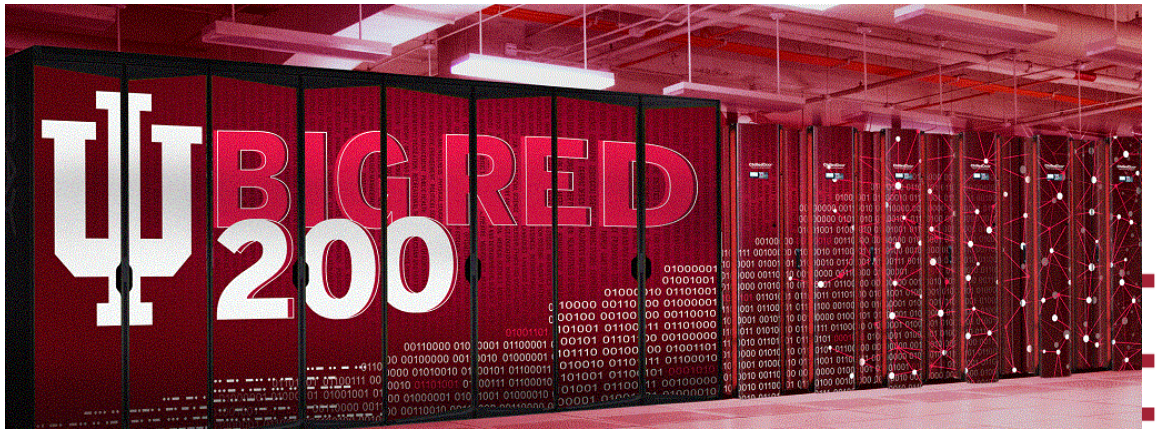

History of Supercomputing

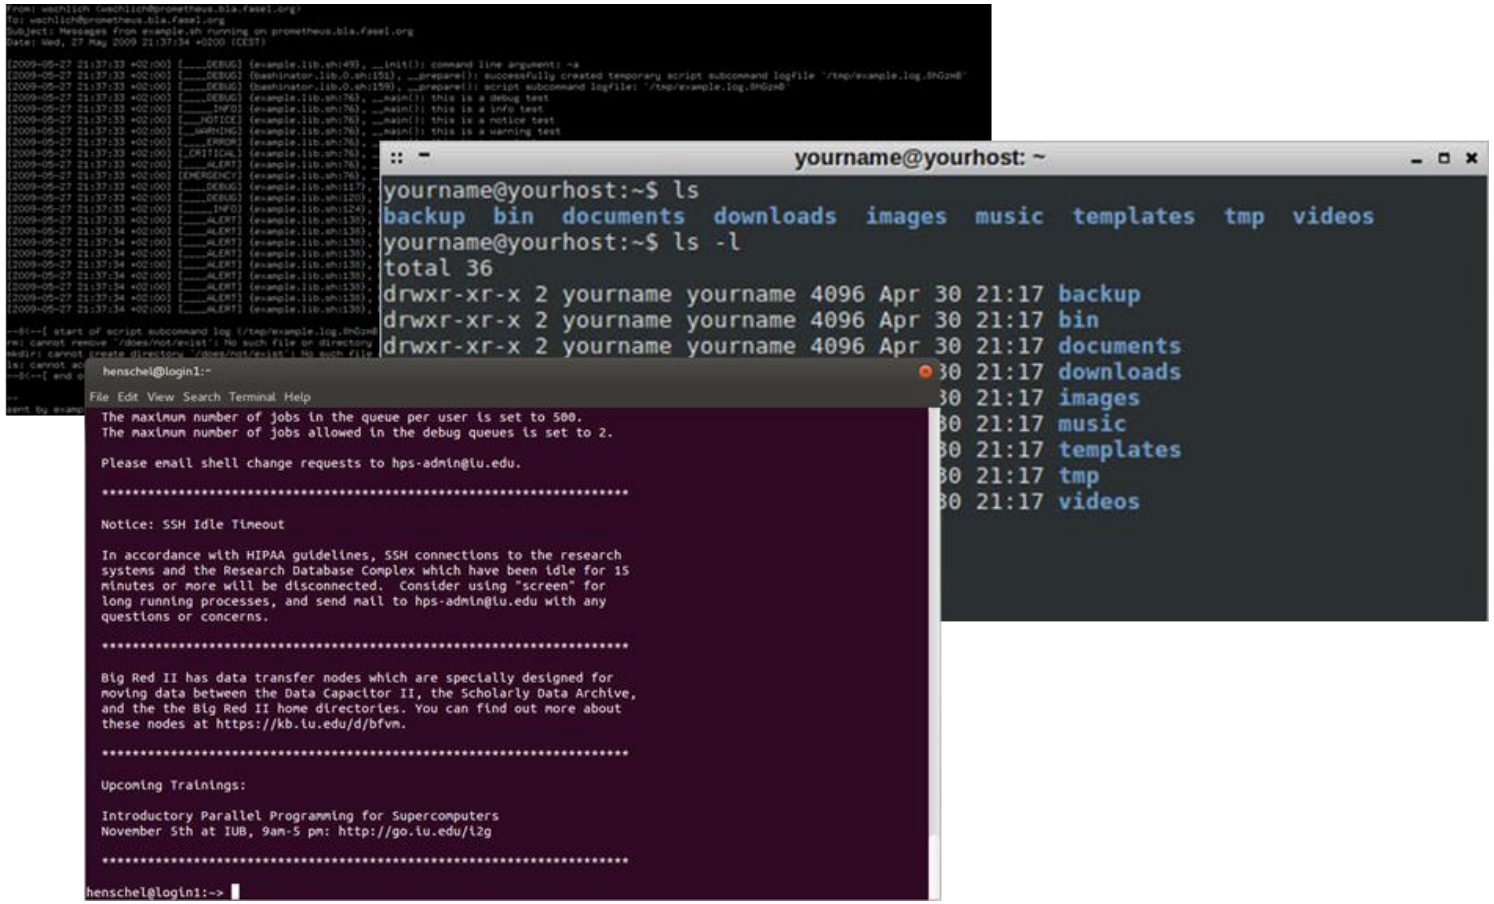

History of Desktop Computing

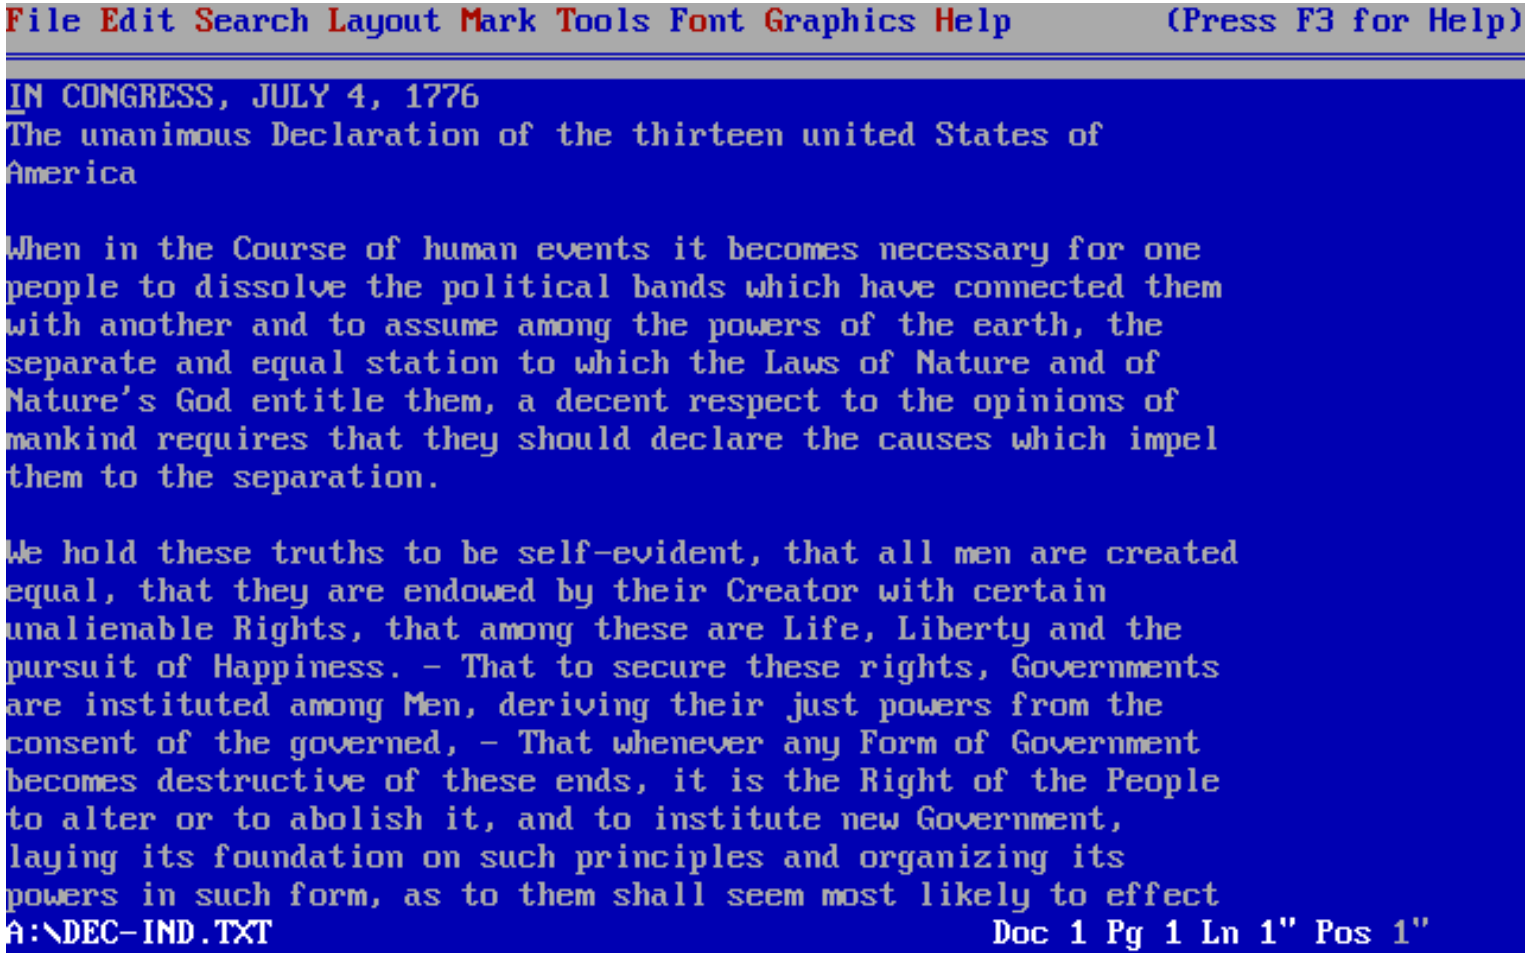

History of Desktop Computing

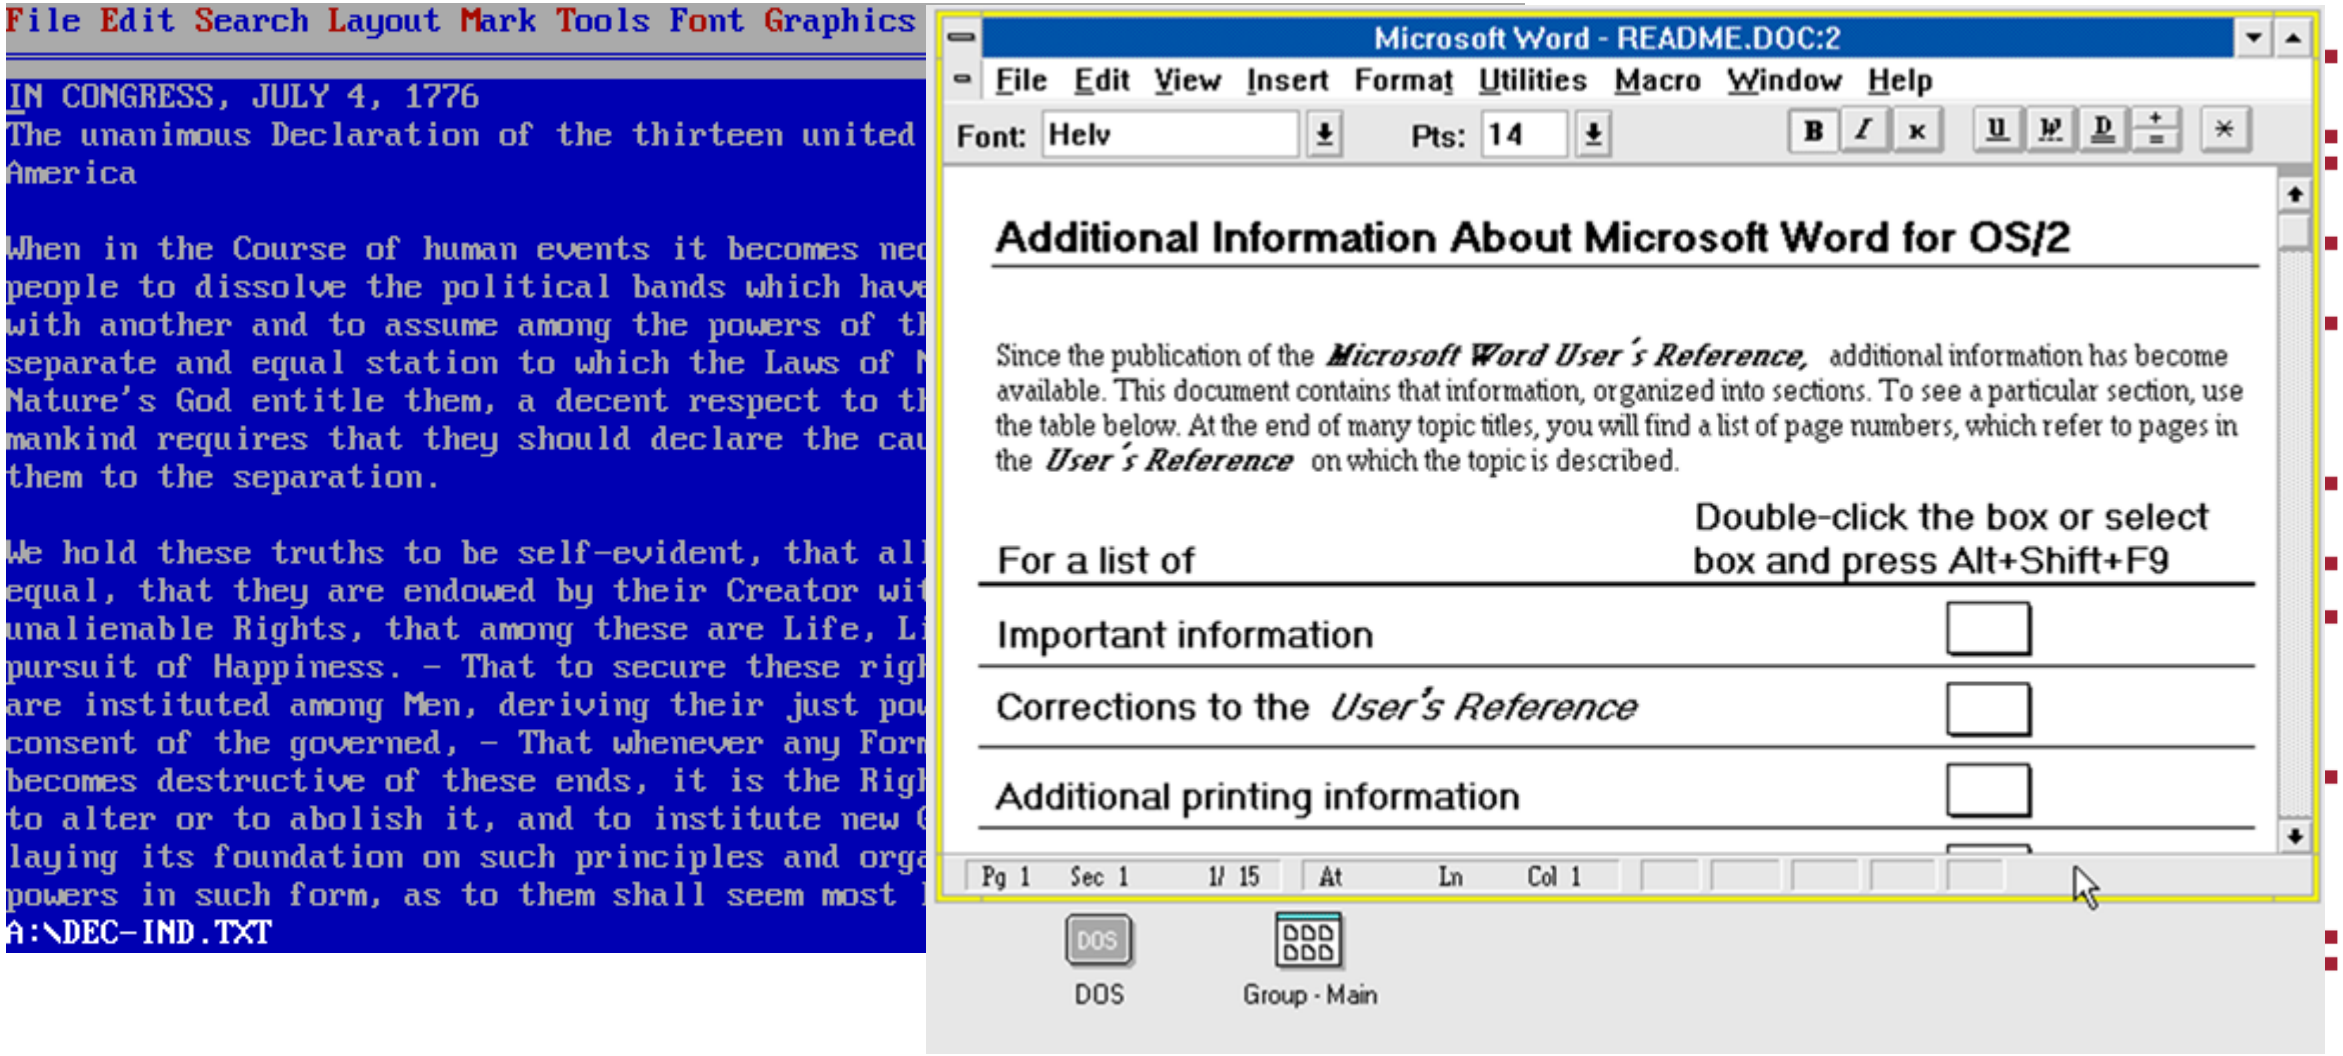

History of Desktop Computing

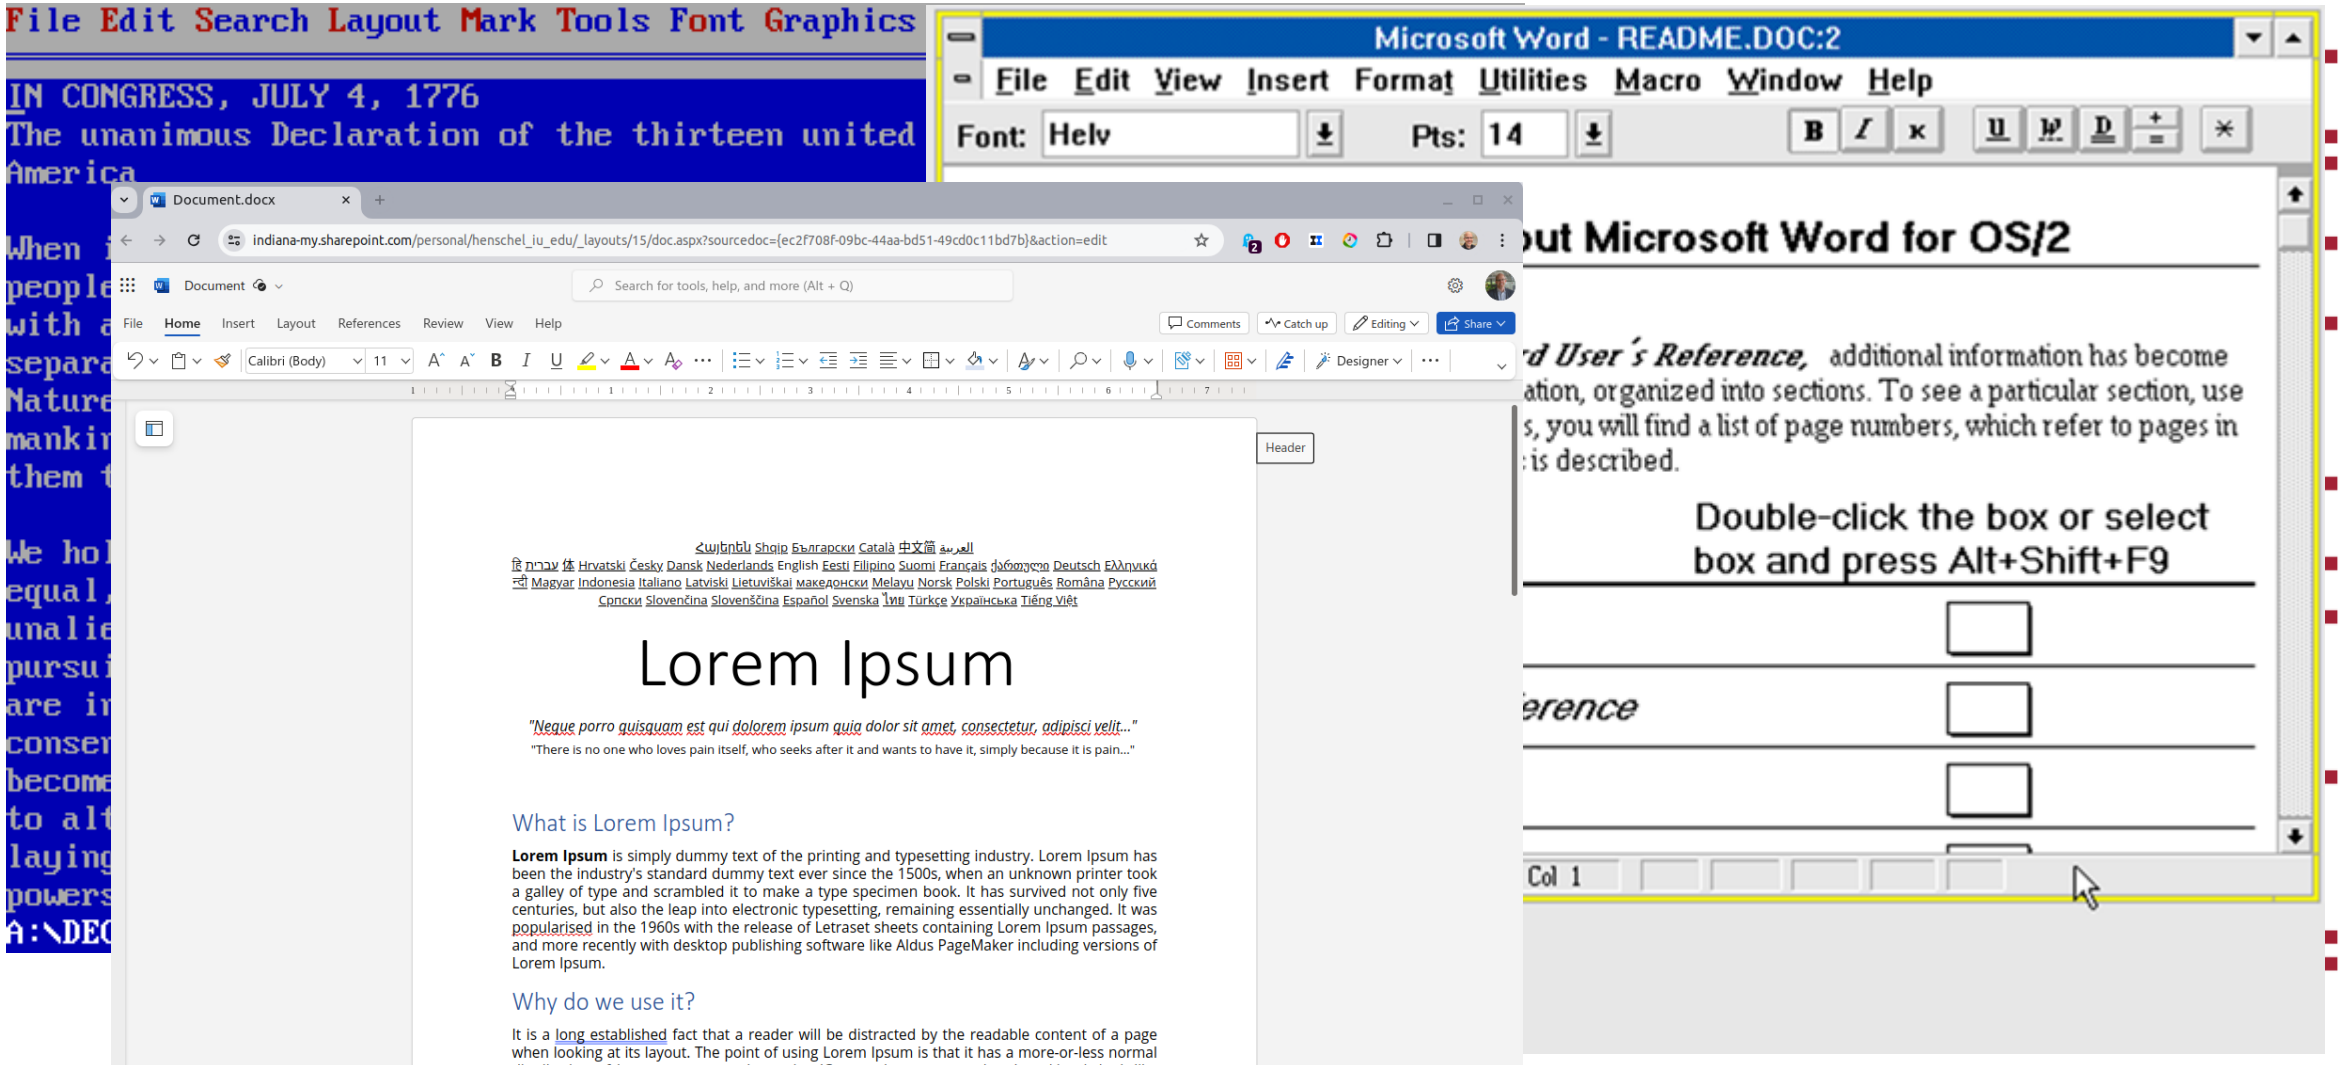

Expectations

- Today, new users are coming with a **different skill set than 10 or 20 years ago**.
- What do “new users” really want and what do HPC centers offer?
  - Change in workload for HPC systems.
- IU has an open access policy for all HPC systems, thus needs to support a very diverse user base. IU experiences this problem sooner than others, and at a higher magnitude.

Leveraging HPC for Research Computing

- **Competing with cloud providers** PR is hard. Their sales pitch is very convincing and should not be underestimated!
- HPC centers have to **balance convenience and efficiency**.
- Can HPC centers provide what new users really want? Can HPC centers create something that is truly compelling, both from a performance point of view and a convenience point of view… and that also offers a path to efficient HPC usage?

Summary: Motivation

- Attract more users and new users.
  - Broaden the reach of HPC systems.
  - Lower the barrier of entry.
- Offer a service that users truly want to use.
- Stay relevant in the era of cloud computing.
- Graduate Research Desktop users to HPC users and train the next generation of HPC users.

The next hour…

- Introduction
- Motivation for creating a Research Desktop
- Definition and History of a Research Desktop
- Current State and Future Developments

- Q&A and Demo

Challenge: Research Desktop

What if you had access to a **high performance desktop from any endpoint device**? What if this desktop allowed you to **launch interactive graphical applications** like MATLAB, Jupyter Notebooks, Ansys and VS Code, as well as **compose scientific workflows** using graphical and command line tools? What if this desktop had **enough performance to run the applications** and also provided **tools to quickly manage and submit jobs** to an HPC system? What if this desktop was **available to you for weeks or months**, allowing you to reconnect whenever you need it? Indiana University built such a desktop in 2017 and is evolving it ever since.

Architecture (HPC)

SSH

Login Nodes

Batch System

Architecture (HPC + RED)

SSH

Login Nodes

Batch System

ThinLinc
(Application or Web Interface)

Research Desktop

Point and Click

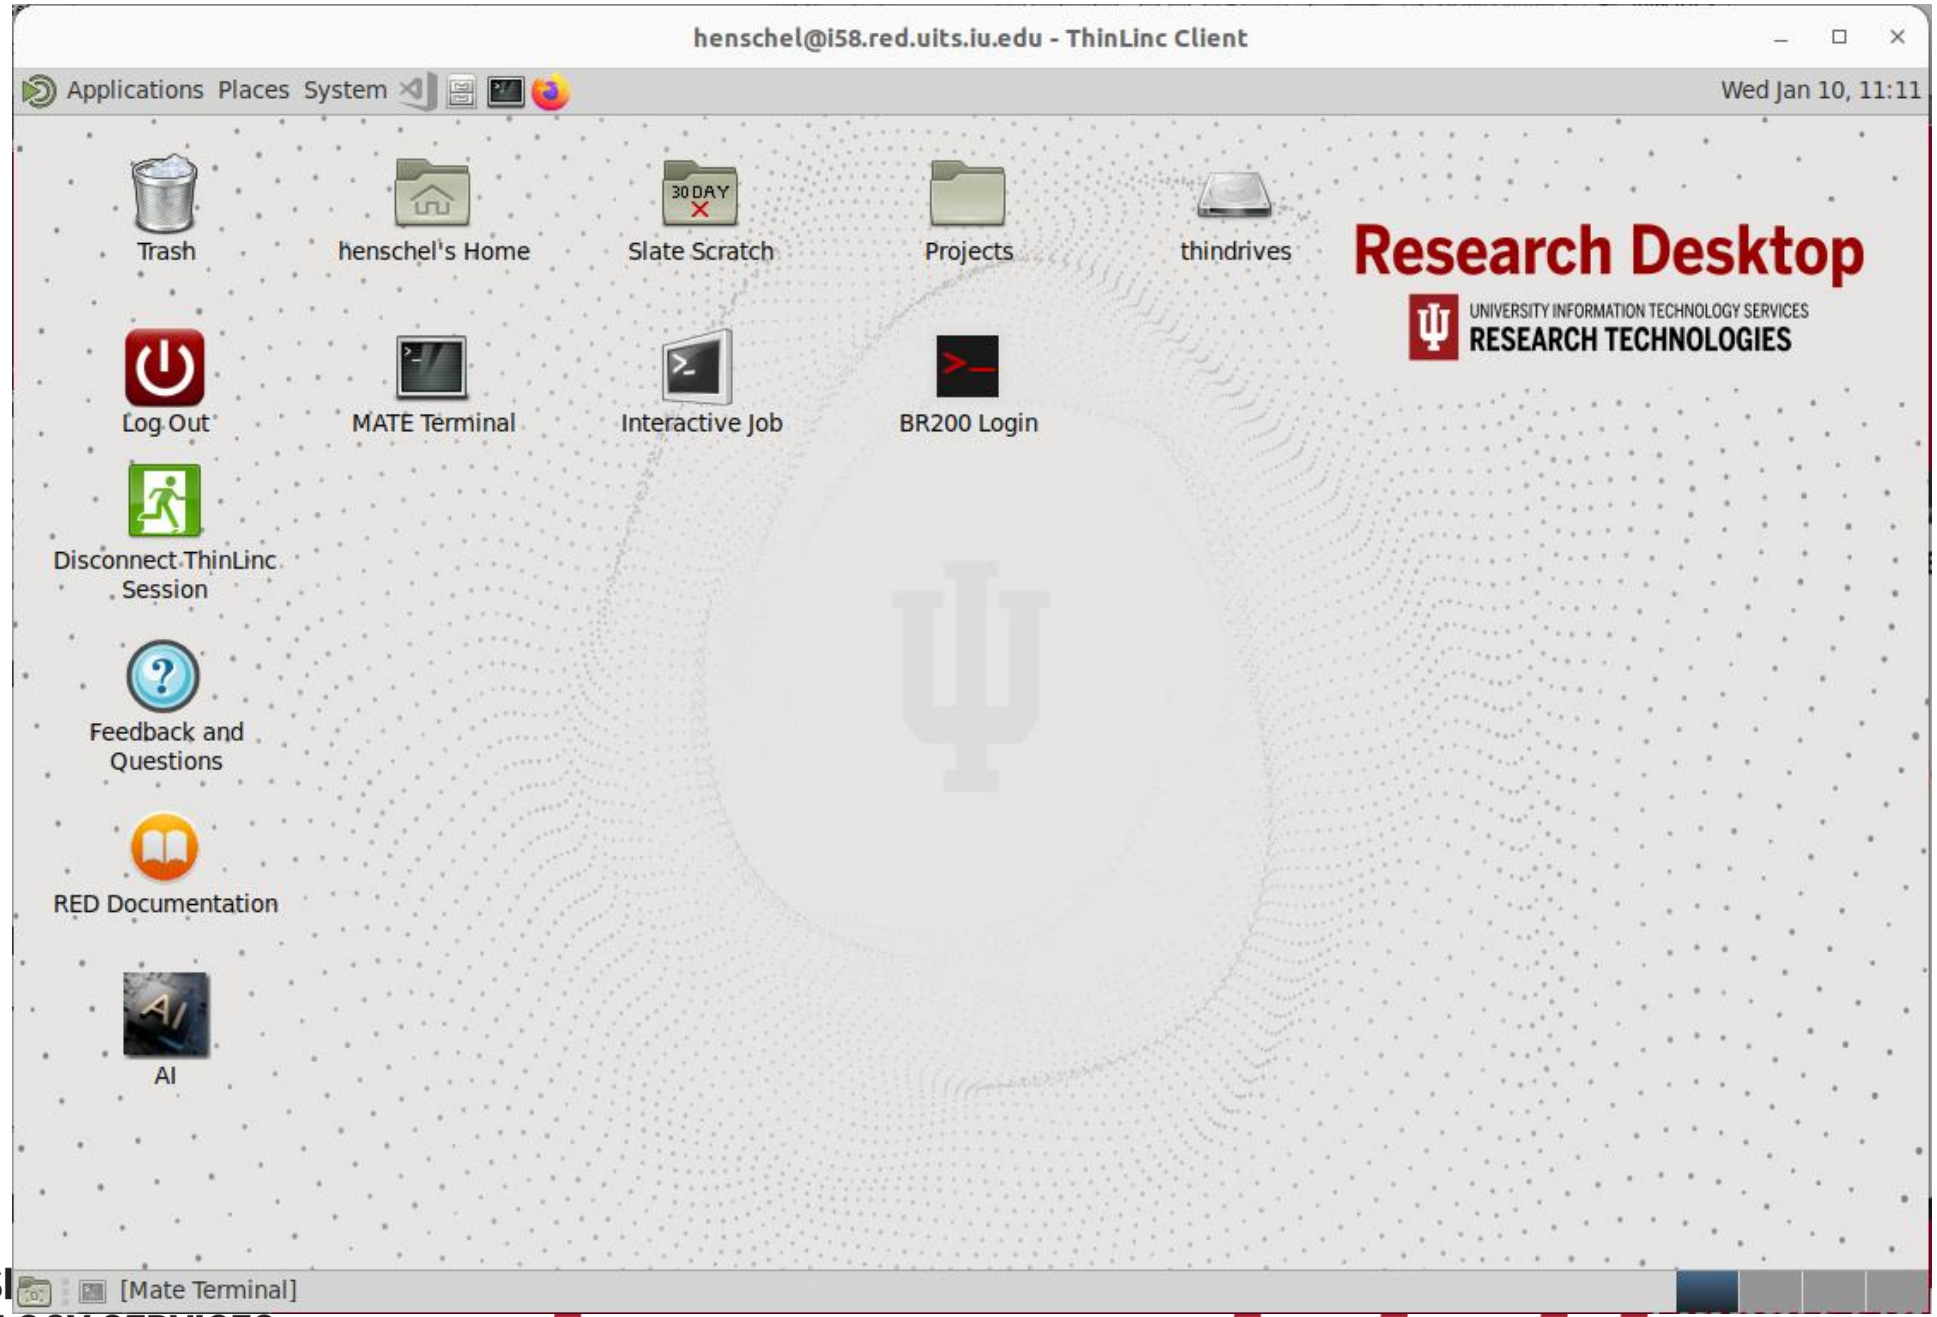
henschel@i58.red.uits.iu.edu - ThinLinc Client
Applications Places System
Wed Jan 10, 11:11
Trash
henschel's Home
30 DAY
Slate Scratch
Projects
thindrives
Research Desktop
UNIVERSITY INFORMATION TECHNOLOGY SERVICES
RESEARCH TECHNOLOGIES
Log Out
MATE Terminal
Interactive Job
BR200 Login
Disconnect ThinLinc Session
Feedback and Questions
RED Documentation
AI
[Mate Terminal]

History

- Launched as “Karst Desktop” in 2015.
- Closed Beta for 6 months, open beta for 2 years.
- General Available since July 2017.
- Renamed to Research Desktop in August 2018.
- Expanded into OpenStack managed VMs in 2021.
- Migration to new hardware and VMs in 2023.

History (cont'd)

- Desktop servers increased from 10 to 34.
- Concurrently active users increased from 100 to 530.
- Upgrade from RHEL 6 to 7 to 8 (GNOME 2.28 to MATE).
- Substantial customization:
  - Heavily customized desktop
  - rclone for cloud storage (Google Drive, Box, DropBox, OneDrive,…)
  - Zenity for graphical dialogs
  - Dynamic cgroups CPU limits
  - Usage statistics scripts
- Enabling web interface and 2-factor authentication.

Policies and Customizations

- 100 GB of RAM per user session.
- Soft limit of 8 cores per user.
- Long runtime of disconnected sessions. (7 days)
- Only full desktops, only one desktop environment, only one session per user.
- An interactive systems needs more care and feeding than a batch system.
- Custom desktop (background, menus, icons)
- Tools
  - Request/Help form
  - Batch system GUI
  - Interactive job icon/tool
  - rclone for cloud storage
- Usage statistics scripts

Summary: Definition and History

- The Research Desktop is “a place where you want to do all your computational research”.
  - Convenient enough that you want to use it every day.
  - Performant enough that it can handle most demanding workflows.
- Growing service at IU - default way how to access IU HPC systems.
- Always trying new things, new customizations, new opportunities, new use cases.

The next hour…

- Introduction
- Motivation for creating a Research Desktop
- Definition and History of a Research Desktop
- Current State and Future Developments

- Q&A and Demo

RED - Yearly Unique Users

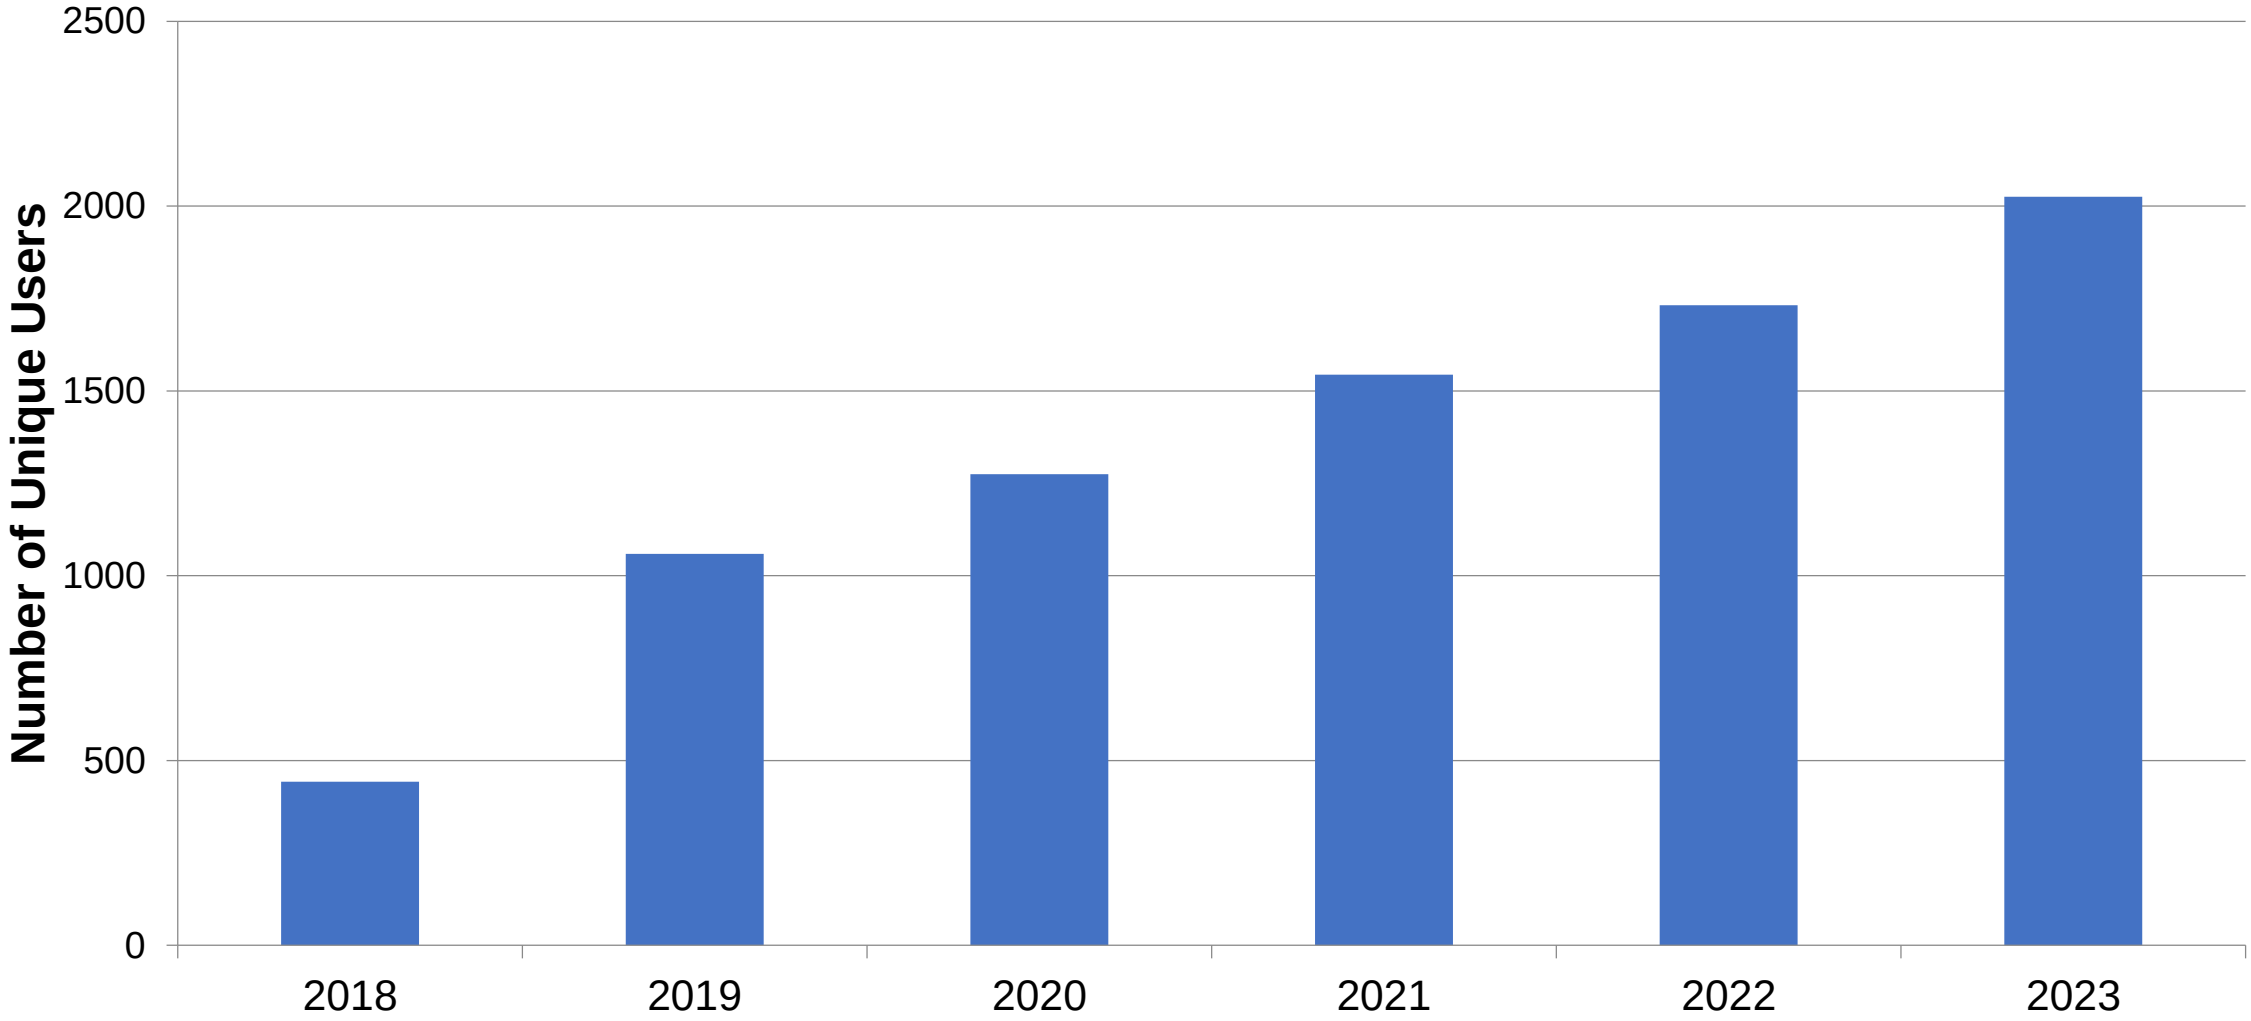

RED - Monthly Unique Users

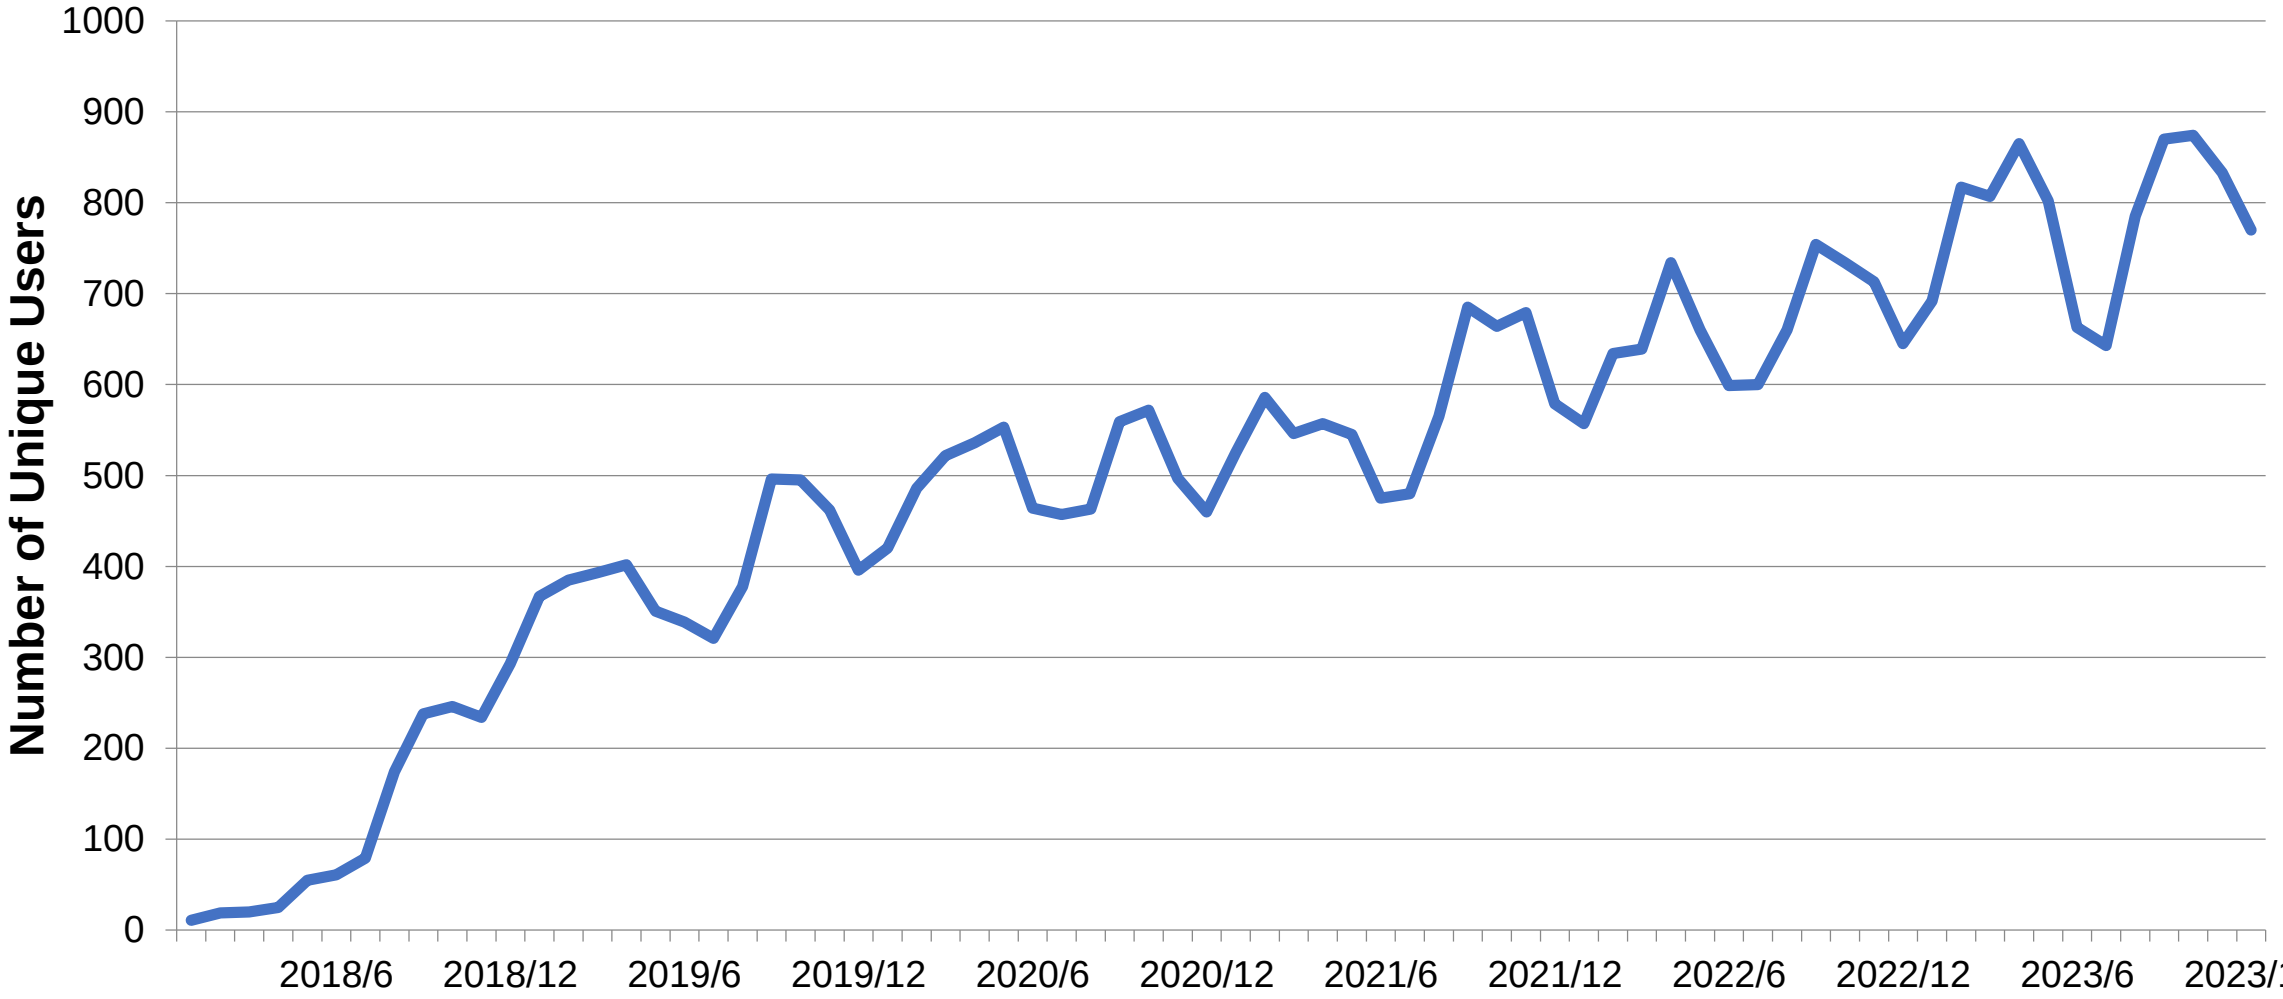

Red – Daily Active Users

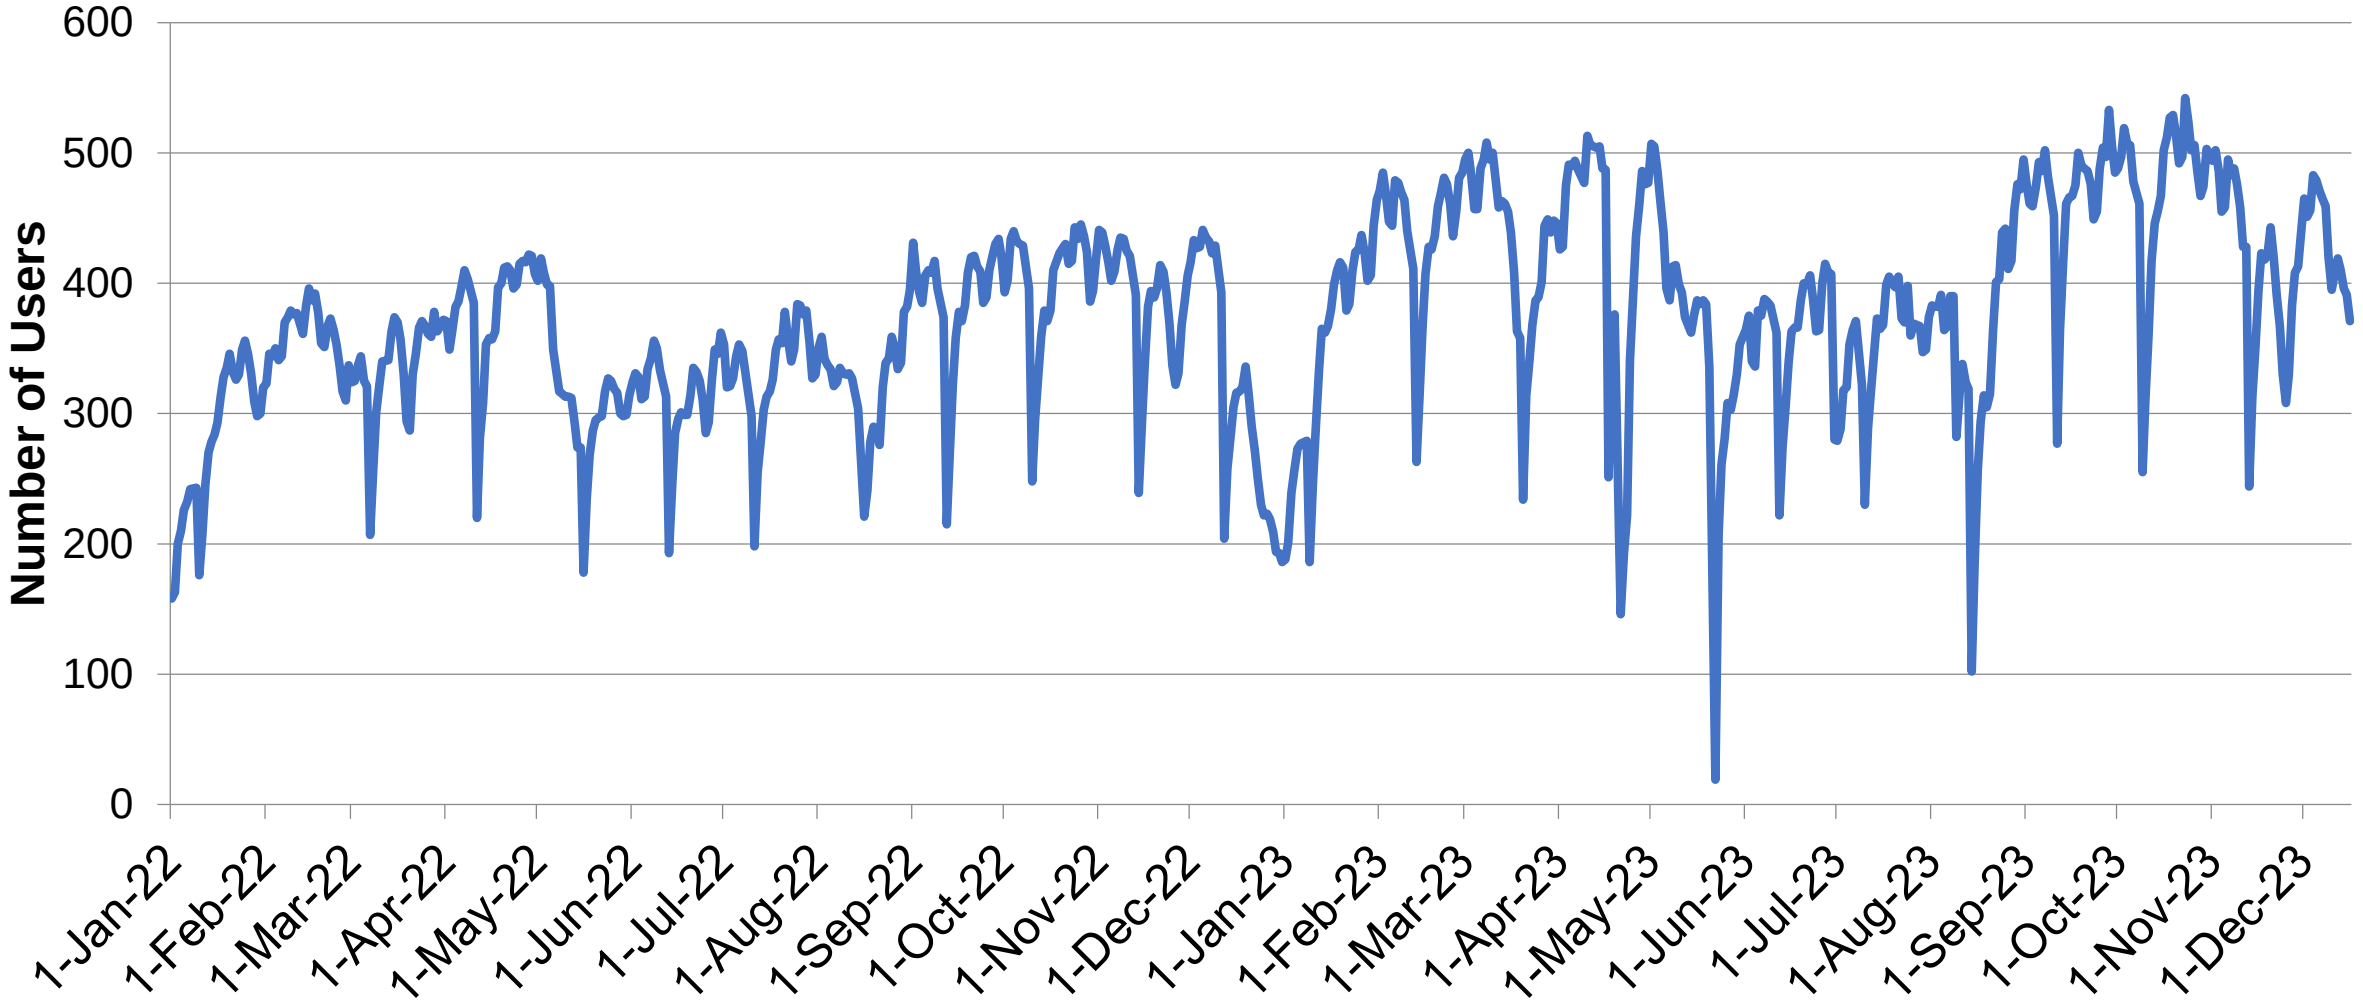

Use Cases

- Mathematical and statistical applications like SPSS, SAS, STATA, MATLAB
- Graphical tools, IDEs, etc.
- Visualization applications
- GUIs of HPC performance analysis applications such as Vampir, Allinea MAP, TotalView
- Client/Server applications like COMSOL Multiphysics, Schroedinger, Ansys
- Replacement for screen / nohup
- Teaching (Easy access to compilers and editors for classes)
- HPC training
- Secure Data Enclave (FISMA and HIPAA)
- Long running data movement jobs
- Support collaboration across institutions
- User Support


allinea
MAP


R Studio


MATLAB


allinea
PERFORMANCE
REPORTS


jupyter


Wolfram
Mathematica


sas


STATA


TOTALVIEW
TECHNOLOGIES


python


SPSS


VMD
Remote

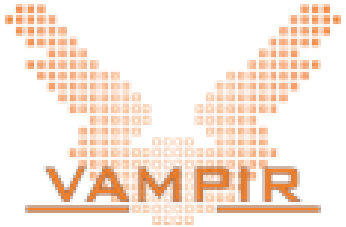
VAMPIR


COMSOL
MULTIPHYSICS


Schrödinger

Research Desktop vs. HPC vs. OOD

Category	Research Desktop	HPC	OOD
Run text mode applications	Run in terminal	Write batch script, submit to batch system	Batch job via Web UI
Run Graphical applications	Pick application from menu	X11 Forwarding SSH connection + X11 Forwarding batch jobs + X-Server on client	App in Web UI
Up/Download files	File browser / ThinDrives	SCP, SFTP	Web UI for upload/download
Edit text documents	Any Linux editor	VI/EMACS/NANO	Build in editor
Look at images, PDFs	Double click in file browser	X11 Forwarding	App in Web UI
Run multiple applications	Just launch them on the desktop	Separate batch jobs	Separate tabs in the browser
Long running application	Launch, disconnect/reconnect	Batch job, tail output file	Batch job, observe output file

The Future - Building a Community

- Explore new concepts of how to integrate an HPC system into a modern desktop environment. (GNOME 3 and extensions, gfxLauncher, VirtualGL)
- Interactive and Urgent HPC workshop/BoF at SC and ISC.
- IEEE CiSE special issue paper:
  - Interactive And Urgent HPC: Challenges And Opportunities
  - https://arxiv.org/pdf/2401.14550.pdf

Summary: Today and Tomorrow

- Growing service at IU - default way how to access IU HPC systems.
- New use cases every year.
- Explore new Desktop environments and abstractions for HPC.
- International movement for more “interactive HPC”.
  - Community building
  - Maybe UD should join the community?

The next hour…

- Introduction
- Motivation for creating a Research Desktop
- Definition and History of a Research Desktop
- Current State and Future Developments

- Q&A and Demo

Q&A and Demo

- Connect
- Applications in menu, VS Code
- How to import data using ThinDrives
- Disconnect / Reconnect with web app
- Modules
- Connect to BigRed 200
- Interactive Job