3-7 March 2025, Yokohama, Japan

Experiences from the National Research Platform

- Innovations and Challenges of Scientific Workloads -

59th Meeting of the Asia Pacific Advanced Network (APAN59)
Asia Pacific Research Platform (APRP) WG Session 1

Seungmin Kim
ehf (at) yonsei.ac.kr
Yonsei University College of Medicine, Republic of Korea
(Ex-) San Diego Supercomputer Center, University of California, San Diego

The National Research Platform (NRP)

The **NRP** has moved beyond its original vision as a **"ScienceDMZ data freeway"** and **evolved into a distributed cloud supercomputer.**

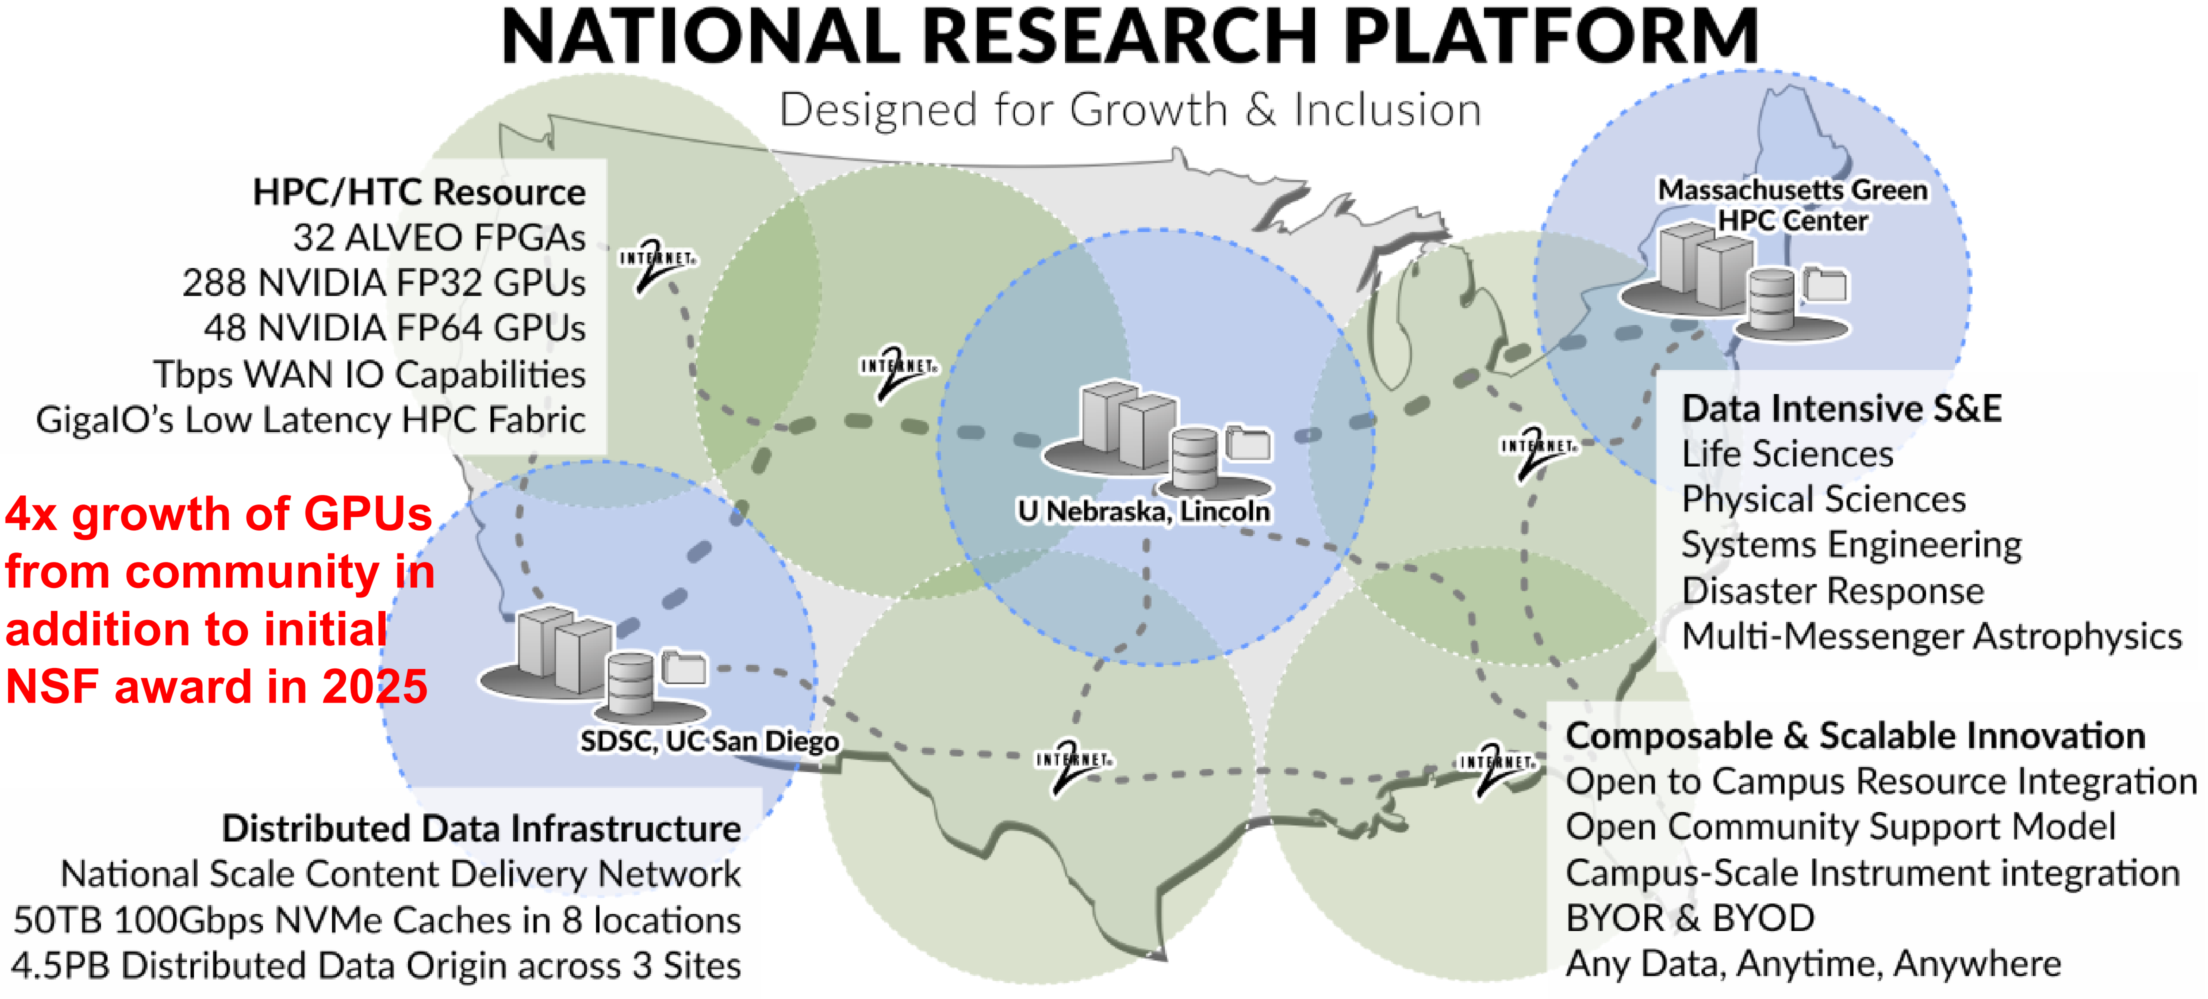

4x growth of GPUs from community in addition to initial NSF award in 2025

Reference: https://www.hpcwire.com/off-the-wire/nrp-accelerating-science-with-bold-national-research-platform/

Nautilus Hypercluster

- **Nautilus** is a geo-distributed and federated **Kubernetes** cluster
- Institutions contribute **hardware** within **campus premises** and join through **ScienceDMZ**
- **Current GPUs (1459) including community contributions: >4x the initial GPU purchase from the PNRP (OAC-2112167) NSF award (336, previous slide) – Rapid community involvement**

Why Nautilus?

- **Kubernetes** (K8s) automatically distributes many **containers** to many **nodes** under a unified API
- **Containers** isolate the Linux user-space OS (Ubuntu, Fedora, …) for each workload
- Any **Docker** container runs on K8s; deploy your own containerized service like AWS or GCP!
- Maximum flexibility, in between traditional HPC (SLURM) and Virtual Machines (OpenStack)
- **DevOps** platform + job scheduling: **JupyterHub** (**PyTorch**, etc.), **LiteLLM** (secure in-cluster **LLAMA 3**, **Mixtral**, etc.), Nextcloud, GitLab, Coder (VS Code Server), **Remote Desktop**
- Diverse hardware: CPU, **GPU**, **distributed storage** (Ceph), **FPGA** (Xilinx), **DPU** (Bluefield, GigaIO)
- Simple exposure of **HTTP(S)** web applications using HAProxy Ingress to `*.nrp-nautilus.io`

Centralized & Large-Scale Operation

https://grafana.nrp-nautilus.io

4x GPUs vs NSF award: **Community contributed!**

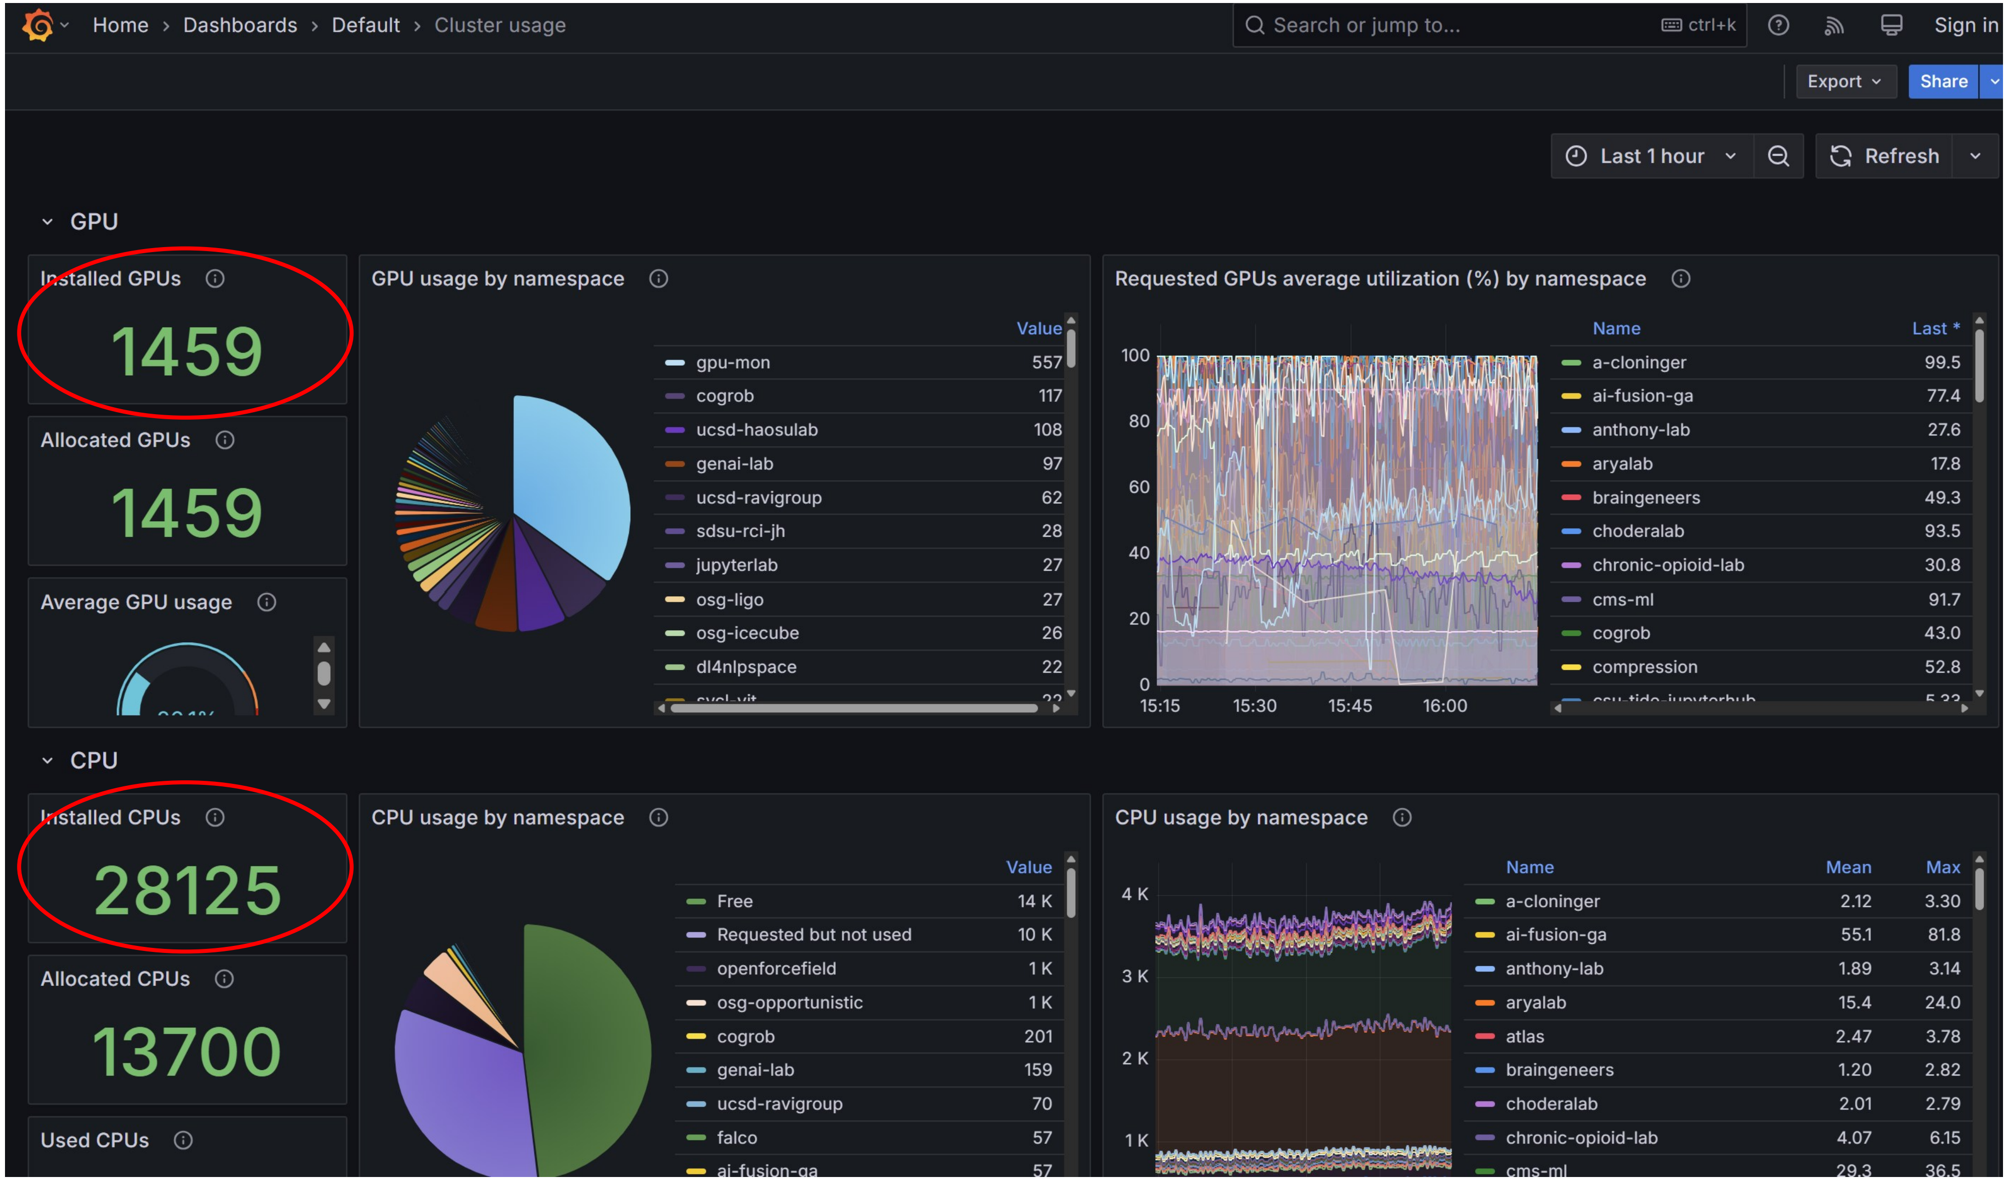

Centralized & Large-Scale Operation

https://grafana.nrp-nautilus.io

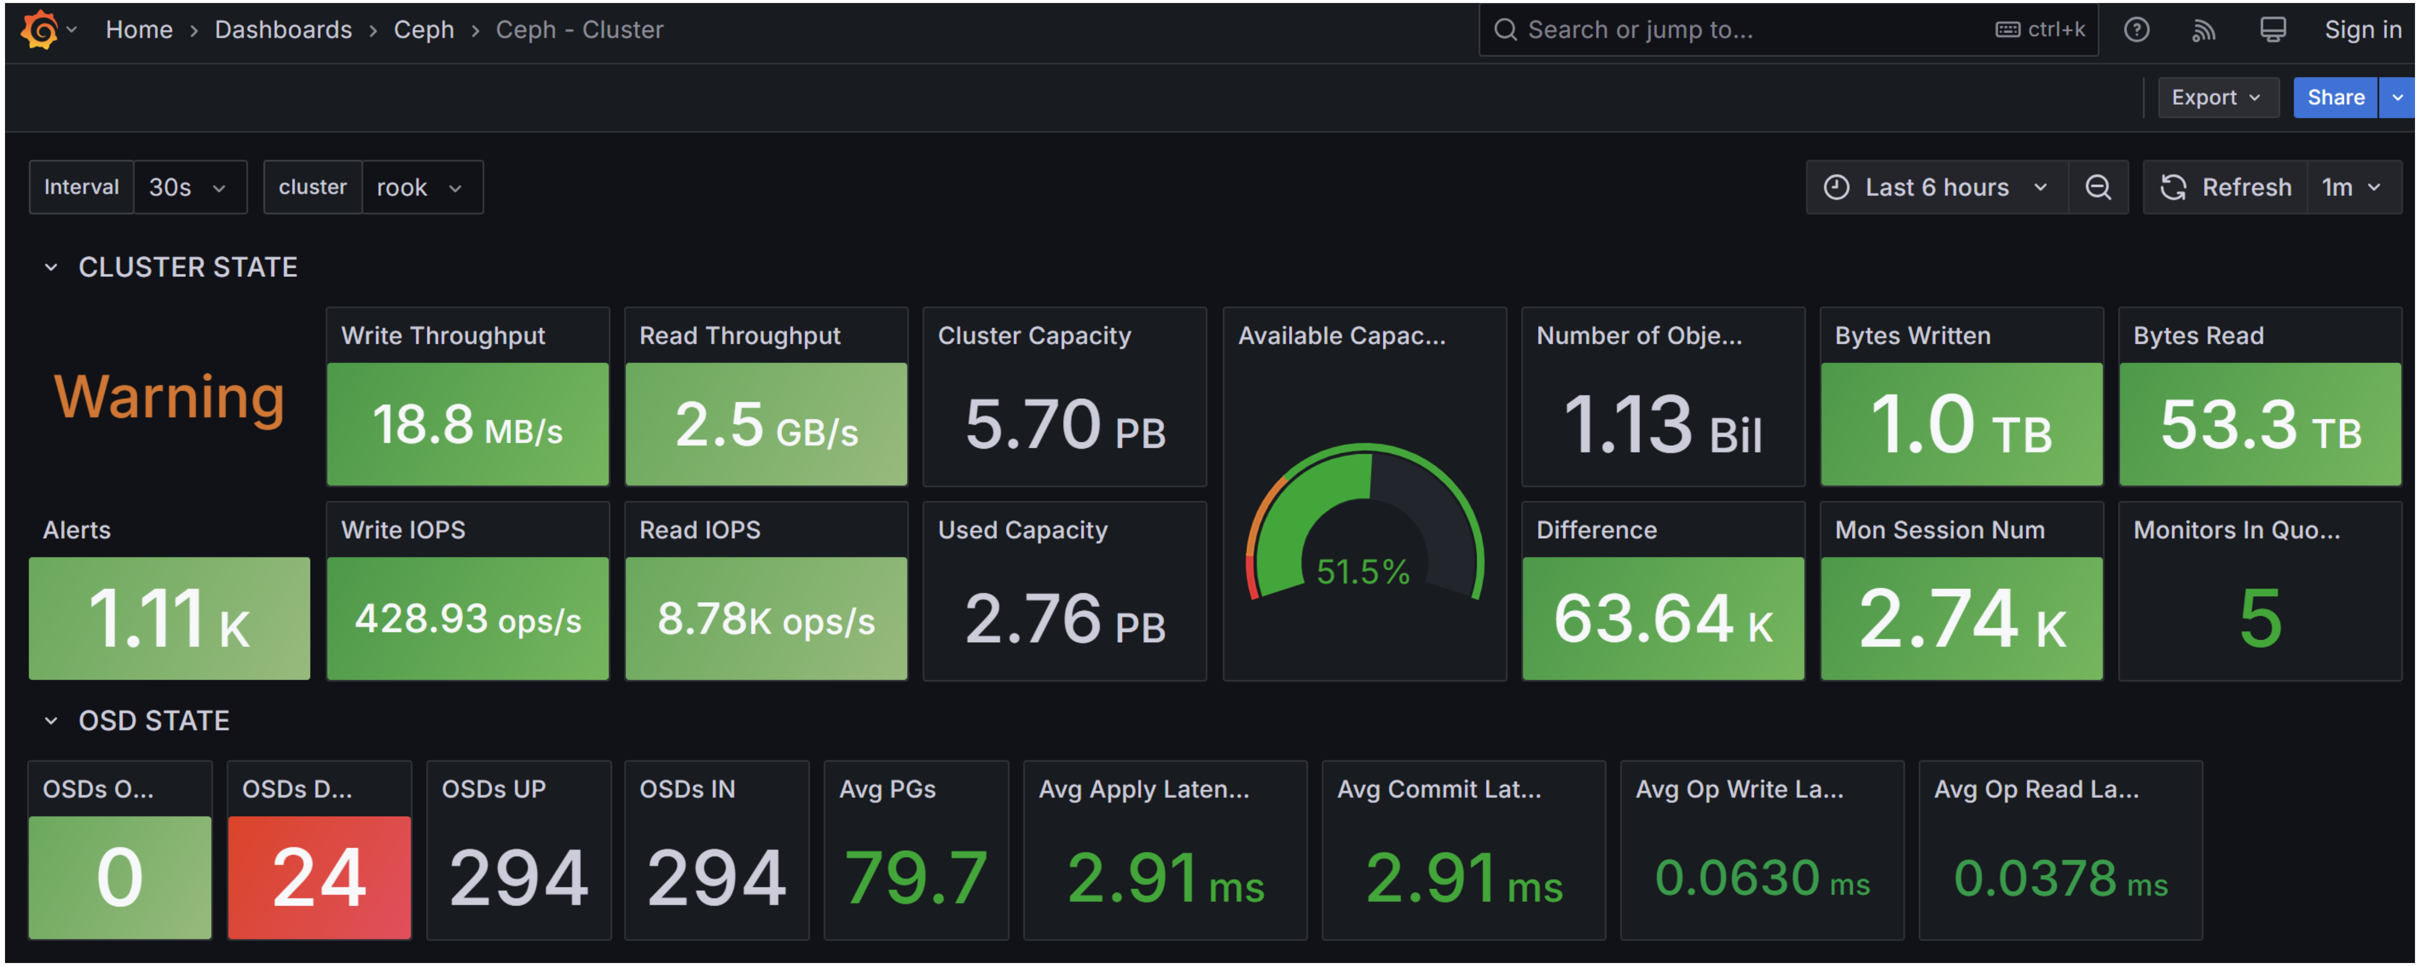

Nearly **20 PB** of geo-distributed storage across the cluster

Centralized & Large-Scale Operation

https://portal.nrp-nautilus.io/resources

Nautilus | Resources | Login

Name	Taints	GPUType	CPU Free	GPU Free	FPGA Free	Mem Free	Disk Free	CPU Total	GPU Total
aarch64.calit2.optiputer....	nautilus.io/ar...		46.722	0	0	64 GB	540.8 GB	48	0
alveo.sdsu.edu	nautilus.io/gi...		61.797	0	0	529.3 GB	2.7 TB	64	0
bcdc-gw-40g-ps.gatech....			12.392	0	0	69.3 GB	941.2 GB	16	0
bois.ps.internet2.edu			36.152	0	0	92.9 GB	794.4 GB	48	0
ceph-1.gpn.onenet.net	nautilus.io/ce...		35.028	0	0	108 GB	250.2 GB	48	0
ceph00.nrp.hpc.udel.edu	nautilus.io/ce...		31.512	0	0	15.8 GB	637.8 GB	48	0
clu-fiona2.ucmerced.edu		NVIDIA-G...	17.587	8	0	119.7 GB	297.2 GB	20	8
controller0.calit2.optiput...	node-role.ku...		60.568	0	0	258.5 GB	943.3 GB	64	0
cph-dgx-node2.humbol...		Tesla-V10...	29.797	4	0	360.1 GB	3.2 TB	80	8
cph-dgx-node7.humbol...		Tesla-V10...	57.797	3	0	426.2 GB	1.7 TB	80	8
cph-dgx-node8.humbol...		Tesla-V10...	13.797	1	0	194.3 GB	1.7 TB	80	8
cph-dgx-node9.humbol...		Tesla-V10...	37.797	2	0	417.6 GB	1.7 TB	80	8
cph-gpu01.humboldt.edu			50.797	1	0	52.6 GB	6.9 TB	56	4

Focus on your research, not infrastructure

- **Infrastructure as Code**: Ansible **auto-manages** the **OS and drivers (+ GPU)** of **your contributed servers**, and NRP provides priority to your own nodes when you need it – use diverse resources provided by other institutions as well
- Dramatically reduce maintenance load due to **standardized setup** and **centralized maintenance** without compromising **flexibility and versatility**
- **Abstract** computing resources where everything is pre-configured, just like AWS/GCP/Azure public cloud platforms
- Develop and deploy in one place – Pre-configured containers with **CUDA** and your **LLM, machine learning, deep learning, data science** library
- Submit larger-scale jobs (just like SLURM or HTCondor) after small-scale prototyping in the **same interface** with Jupyter, GitLab CI/CD, Coder, etc.
- **Focus on your research, not the hardware and operating system!**

Research Productivity Tools in Nautilus

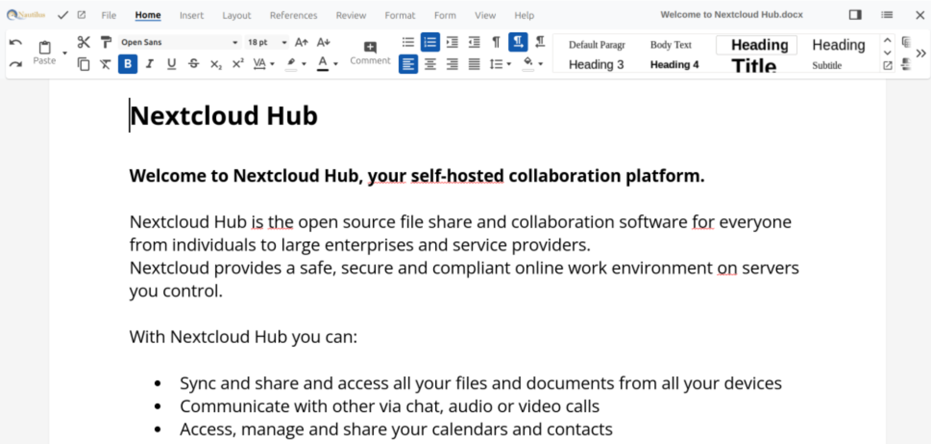

Nextcloud / Collabora Office / Matrix

Self-Hosted Office 365, Google Workspaces, Slack
Collaborative Calendar and Time Management
Team Messaging, RSS Feed for Journals

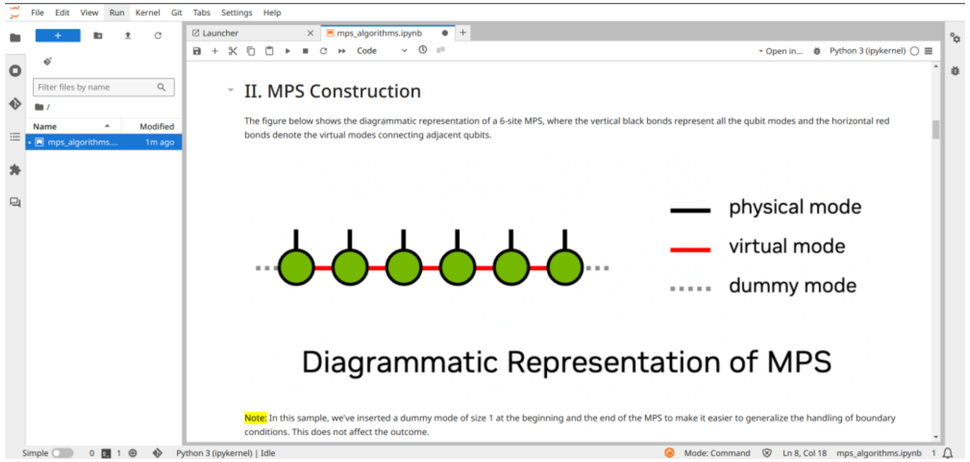

JupyterHub

Jupyter AI, VS Code Server, PyTorch, TensorFlow, FastAI, JAX, Hugging Face, DeepSpeed

- Jitsi Meet: Self-Hosted Zoom or Microsoft Teams
- Self-Hosted Overleaf (LaTeX editor)
- Coder: Self-Hosted CodeSpaces/GitPod

- LiteLLM: OpenAI-compatible LLM API and AI Chatbot
- Self-Hosted GitLab CI/CD
- PerfSONAR & MaDDASH

https://portal.nrp.ai/documentation/userdocs/start/resources/

All provided by the team without additional setup or subscriptions

OpenAI-like Local LLM API – Safe & Secure

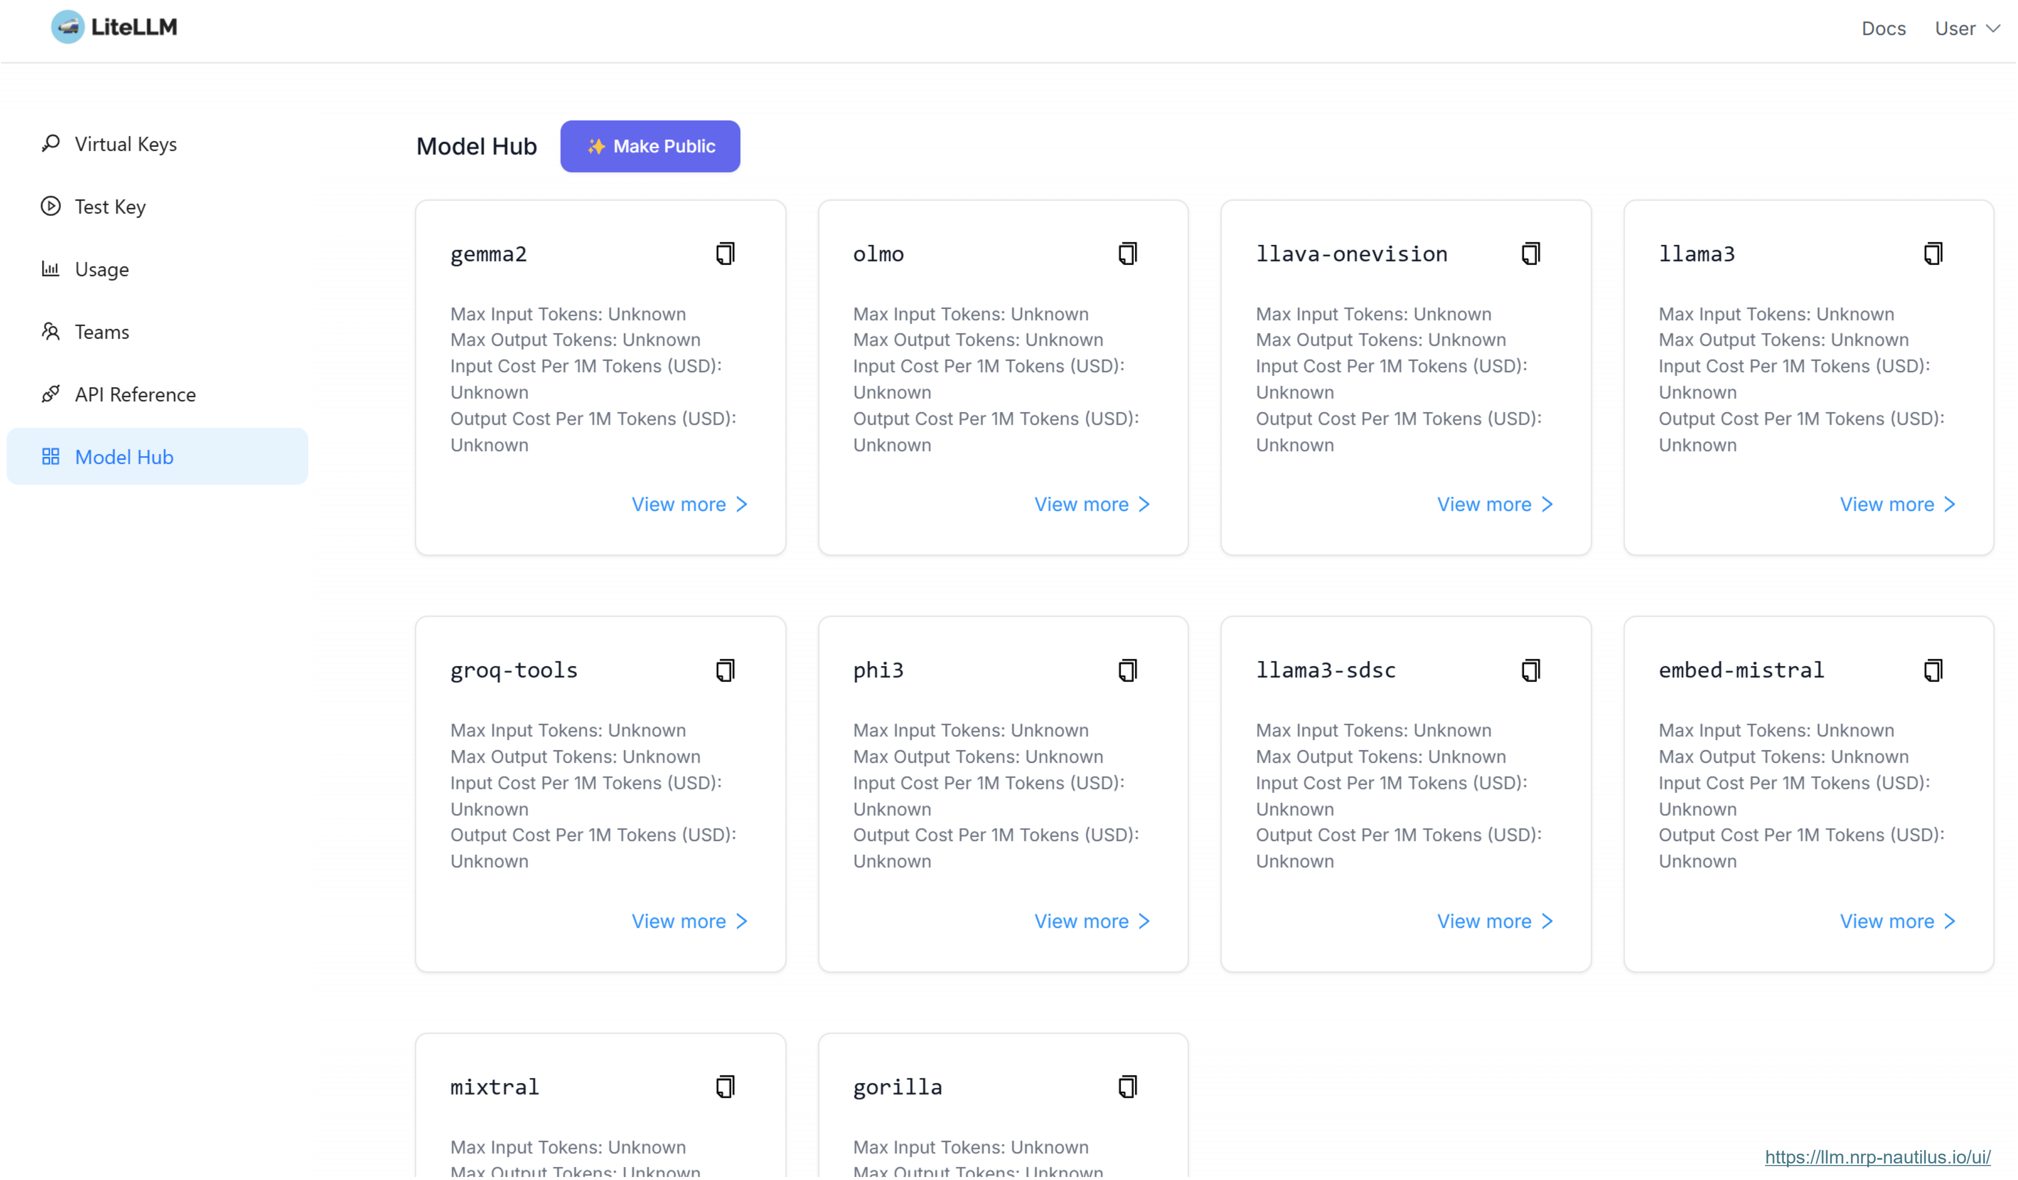

https://llm.nrp-nautilus.io/ui/

OpenAI-like Local LLM Chat – Safe & Secure

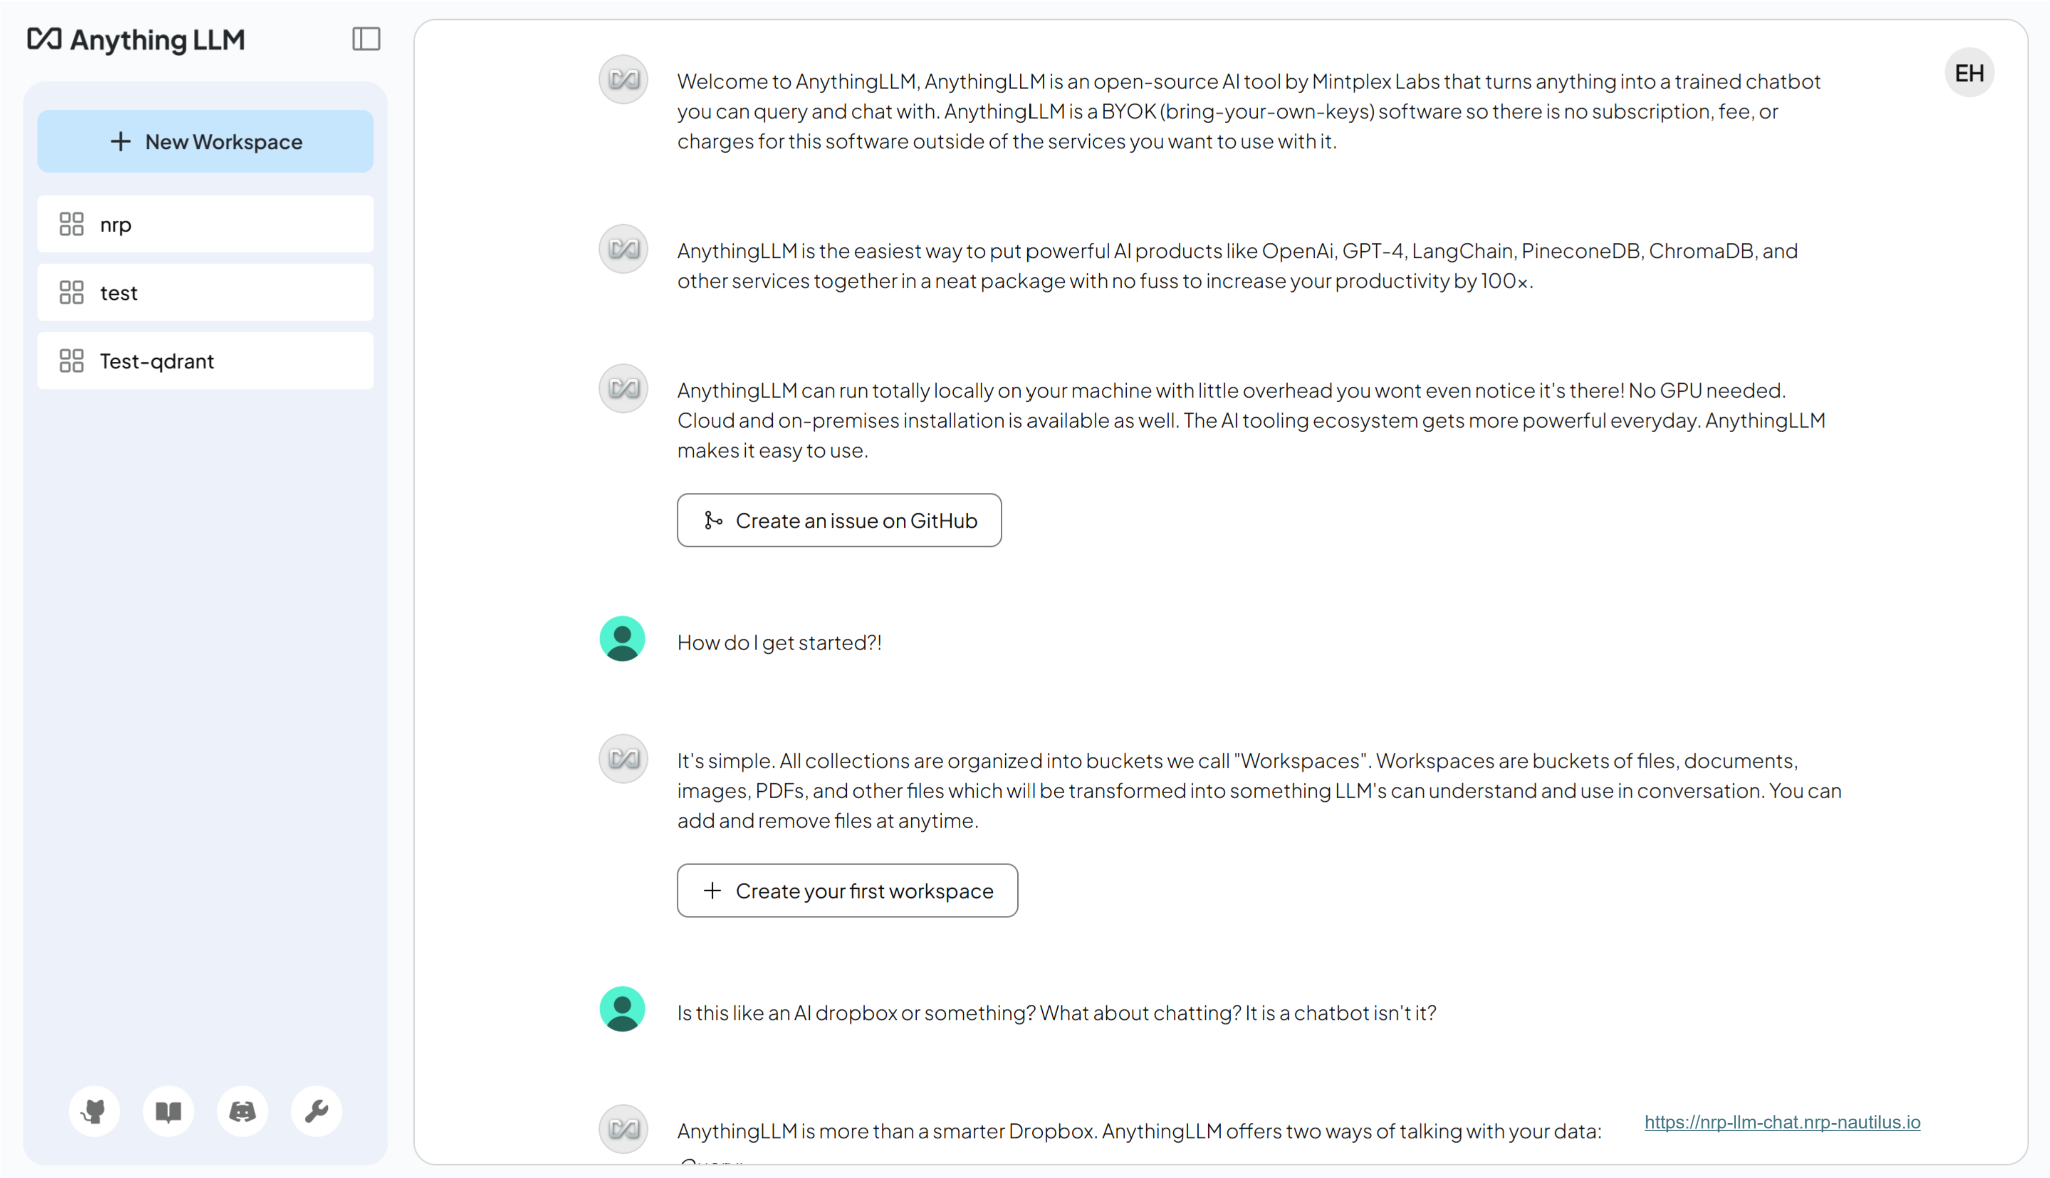

https://nrp-llm-chat.nrp-nautilus.io

Rise of 3D Graphics for Cloud/HPC

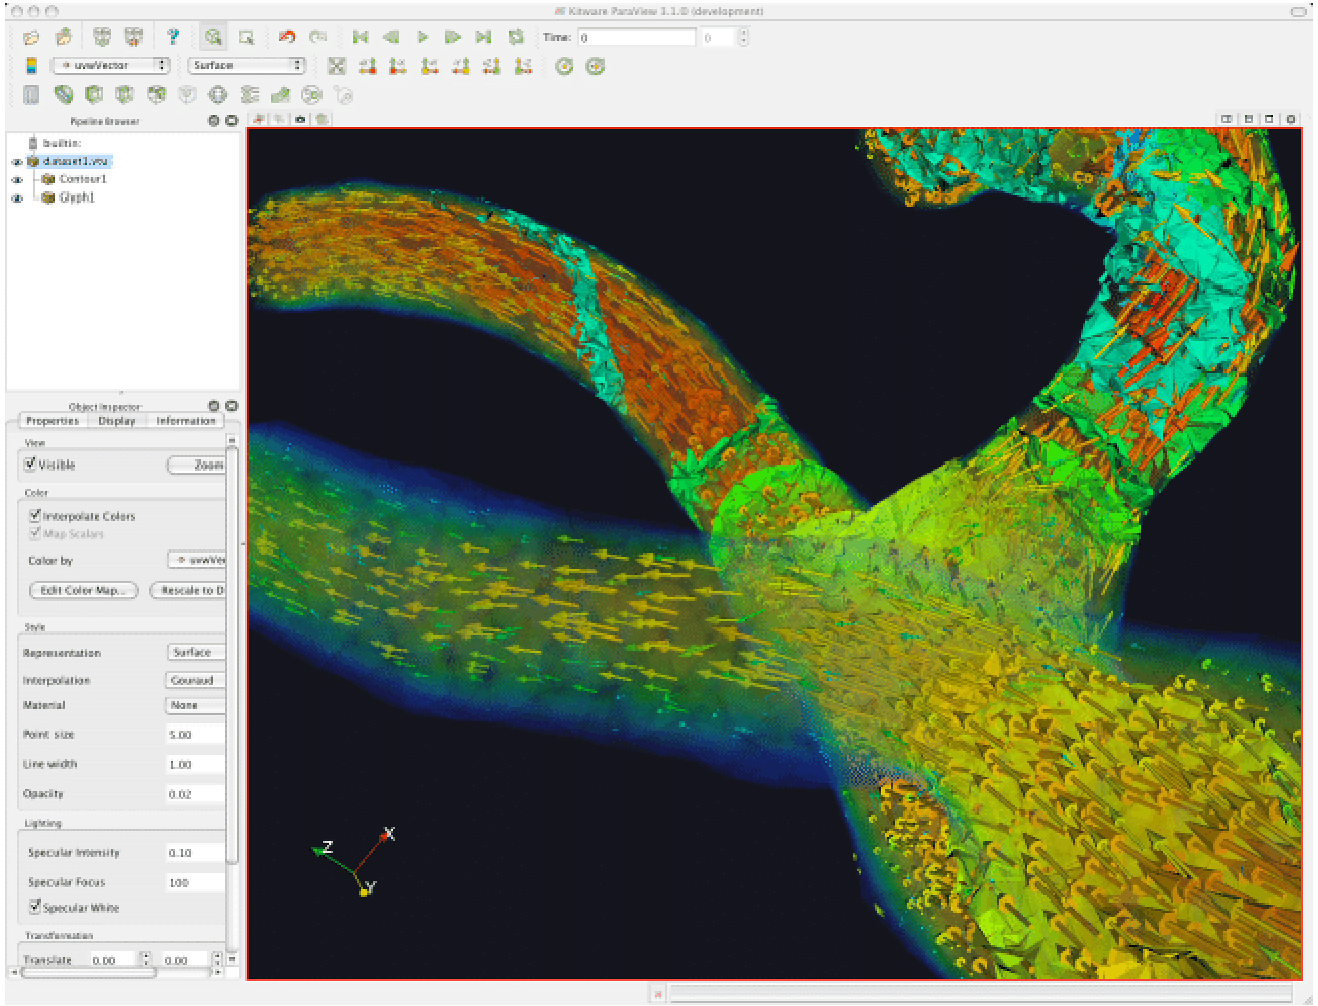

ParaView (VTK) – Medical Visualization

https://www.kitware.com/expertise/

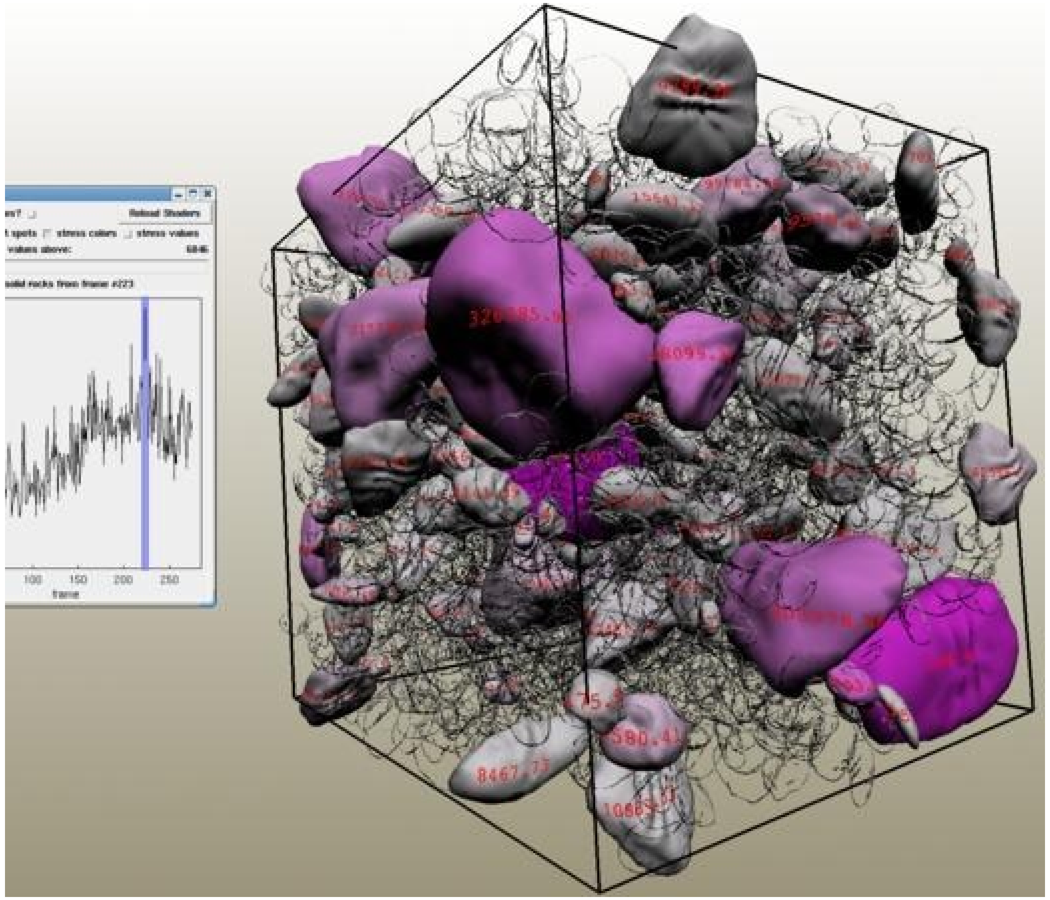

Visualization of High-Performance Concrete (NIST)

https://math.nist.gov/mcsd/savg/vis/concrete/index.html

Massive Emerging 3D Demand for Linux

- **AI-based Drug Discovery & Structural Biology with VMD, PyMOL, ChRIS**
- Early wildfire detection with WIFIRE (**Interactive and Urgent HPC**)
- Robotics Simulators with ROS/Gazebo
- Computer Vision with OpenCV, Autonomous Driving with CARLA
- Fluid Dynamics and CAD simulations with ANSYS, Solidworks
- FPGA Development Environments with Xilinx Vivado and Vitis
- 3D Modeling with Blender & Unreal Editor
- Geographic Information System (GIS) with QGIS/ArcGIS/OSGeo/Cesium
- Drone Mapping with OpenDroneMap
- The Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) with GnuRadio
- Atmospheric Sciences, Archaeology, Geology, Medical Imaging, Oceanography, and other sciences
- Other interactive decisions and tasks requiring the power of HPC clusters

Rise of Multimodal & Visual-Spatial AI

- AI used to be just numbers; input numbers and predict the output (Regression, Classification)
- ChatGPT (Transformer models) process and output text
- Now, AI models process Visual-Spatial data interactively

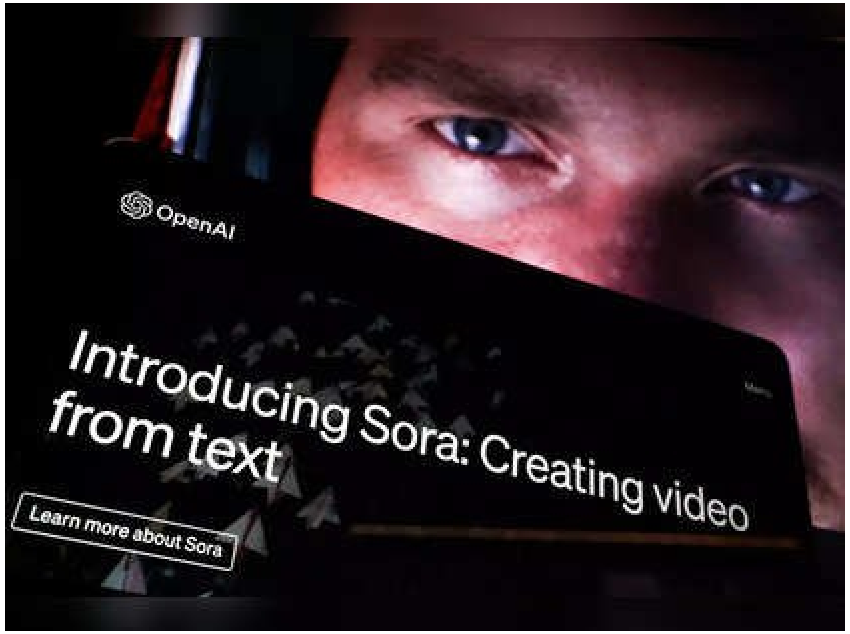

OpenAI Sora – Video from text

https://openai.com/index/sora/

NVIDIA Omniverse – Real-time visual modeling

https://developer.nvidia.com/blog/rapidly-generate-3d-assets-for-virtual-worlds-with-generative-ai/

Use Case: CRISPR Gene Editing

- Yonsei Genome Editing Laboratory (PI: Prof. Dr. Hyongbum Henry Kim)
- Large-scale high-throughput **deep-learning-based design of CRISPR gene editors**
- Published in Cell, Nature Biotechnology, Nature Methods, etc.
- **Predicting the efficiency and activity** of gene editors without lab experimentation; unforeseen scalability with time/cost savings
- Structural design of gene editors and utilizing large language models to analyze gene editors

Cell Supports open access

Prediction of efficiencies for diverse prime editing systems in multiple cell types

Goosang Yu[1,2,10] · Hui Kwon Kim[1,3,10] · Jinman Park[1,2] … · Jiyun Kim[3] · Jisung Kim[3] · Hyongbum Henry Kim[1,2,4,5,6,7,8,9,11] … Show more

Affiliations & Notes Article Info

Download PDF Cite Share Set Alert Get Rights Reprints

Show Outline

Highlights

- The selection of optimal prime editors and guide RNAs is difficult
- The prime editing efficiencies of 338,996 pairs of pegRNAs and targets were evaluated
- Deep-learning models (DeepPrime and DeepPrime-FT) predict prime editing efficiencies
- DeepPrime-Off, another model, predicts off-target effects of prime editing

nature biotechnology

Article https://doi.org/10.1038/s41587-023-01792-x

Deep learning models to predict the editing efficiencies and outcomes of diverse base editors

Received: 29 September 2022

Accepted: 13 April 2023

Published online: 15 May 2023

Check for updates

Nahye Kim[1,2,14], Sungchul Choi[1,14], Sungjae Kim[3], Myungjae Song[1,2], Jung Hwa Seo[4], Seonwoo Min[5], Jinman Park[1,2], Sung-Rae Cho[2,4,6,7] & Hyongbum Henry Kim[1,2,8,9,10,11,12,13]

Applications of base editing are frequently restricted by the requirement for a protospacer adjacent motif (PAM), and selecting the optimal base editor (BE) and single-guide RNA pair (sgRNA) for a given target can be difficult. To select for BEs and sgRNAs without extensive experimental work, we systematically compared the editing windows, outcomes and preferred motifs for seven BEs, including two cytosine BEs, two adenine BEs and three C•G to G•C BEs at thousands of target sequences. We also evaluated nine Cas9 variants that recognize different PAM sequences and developed a deep learning model, DeepCas9variants, for predicting which variants function most efficiently at sites with a given target sequence. We then develop a computational model, DeepBE, that predicts editing efficiencies and outcomes of 63 BEs that were generated by incorporating nine Cas9 variants as nickase domains into the seven BE variants. The predicted median efficiencies of BEs with DeepBE-based design were 2.9- to 20-fold higher than those of rationally designed SpCas9-containing BEs.

Reference:
https://www.cell.com/cell/fulltext/S0092-8674(23)00331-8
https://www.nature.com/articles/s41587-023-01792-x

AI-Based Protein Generation

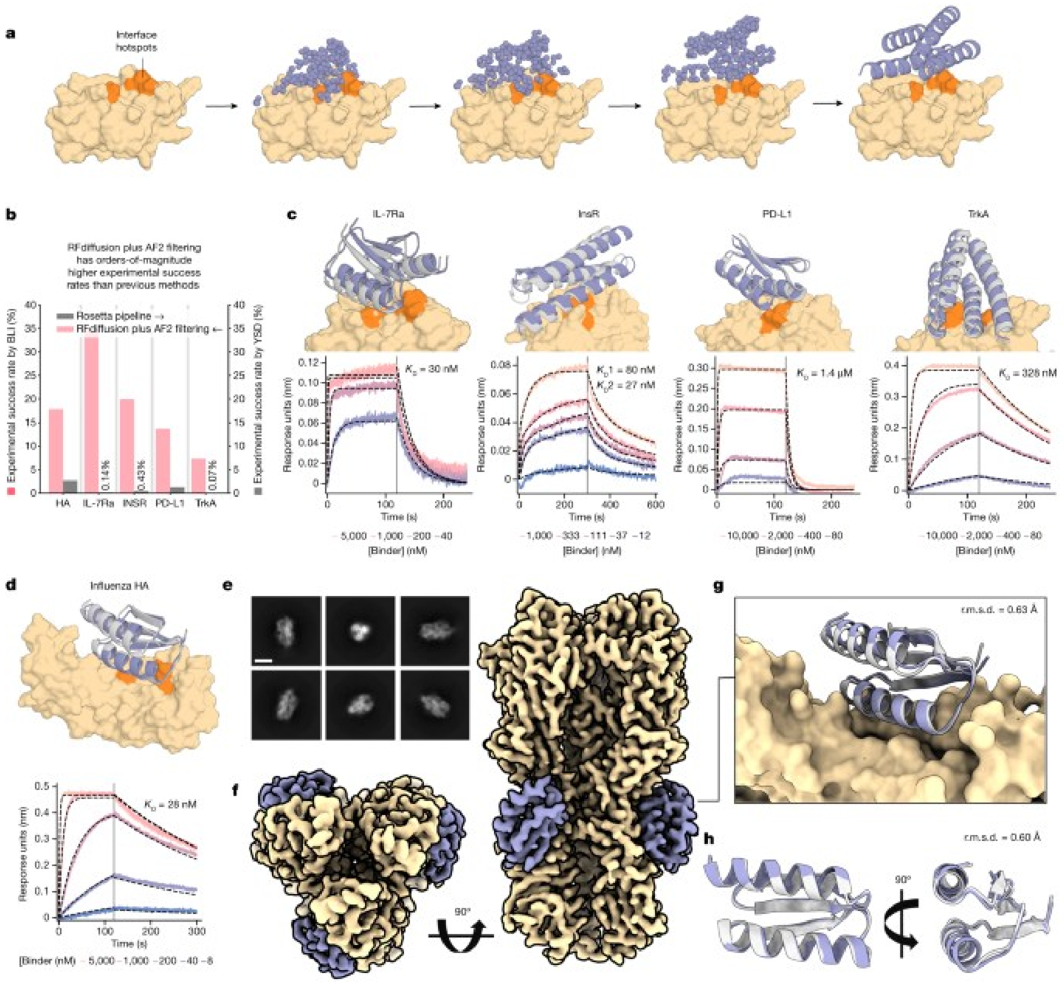

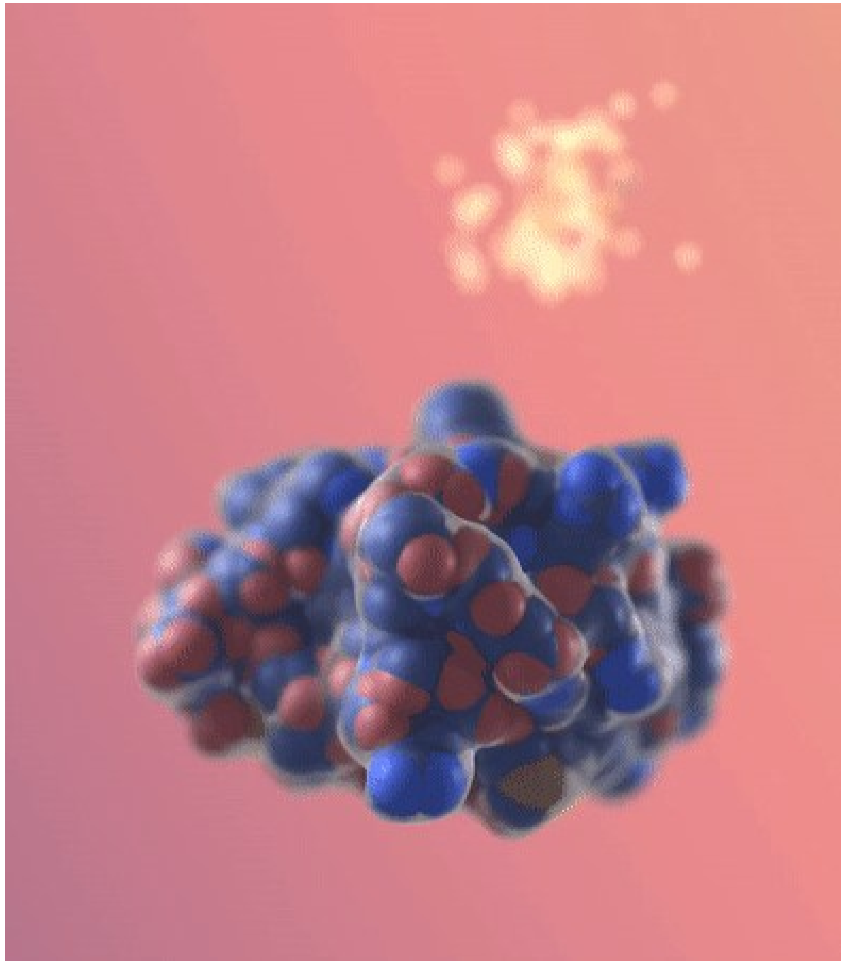

University of Washington, RFDiffusion

Reference:
https://www.nature.com/articles/s41586-023-06415-8
https://www.bakerlab.org/2023/03/30/rf-diffusion-now-free-and-open-source/

GUI on Kubernetes and HPC

- More scientific workloads require graphical interaction
- **GPU-accelerated graphical Kubernetes containers** were unexplored; required host access and Linux X11 display sockets
- **Selkies** (https://github.com/selkies-project) by Dan Isla and Seungmin Kim
- First developed by Google Cloud engineers (Dan Isla, Jan Van Bruggen)
- Subsequently evolved to provide fully isolated (unprivileged) **GUI desktop containers** for HPC (**Apptainer/Singularity**) and Kubernetes workloads
- Now also utilized in emerging **public AI cloud** services including Vast.ai, RunPod, TensorDock, CoreWeave, etc. (Utilizing Docker or Kubernetes)
- **Use HPC like your own PC in front of you**

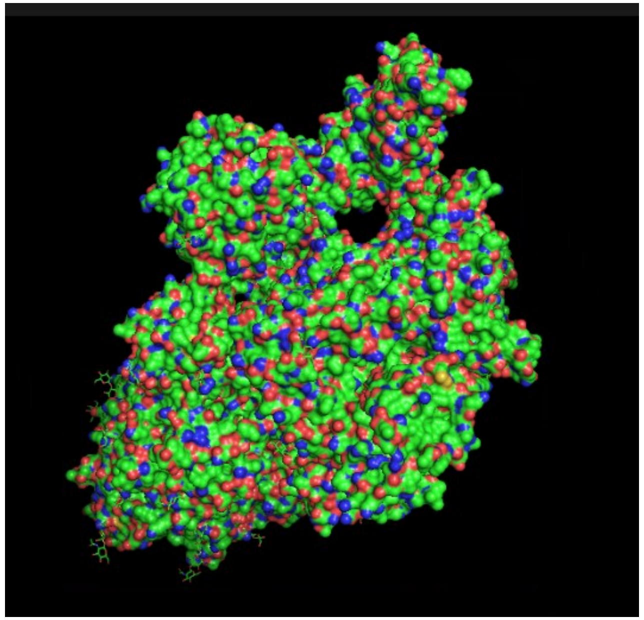

https://www.youtube.com/watch?v=zagBgED7JOI

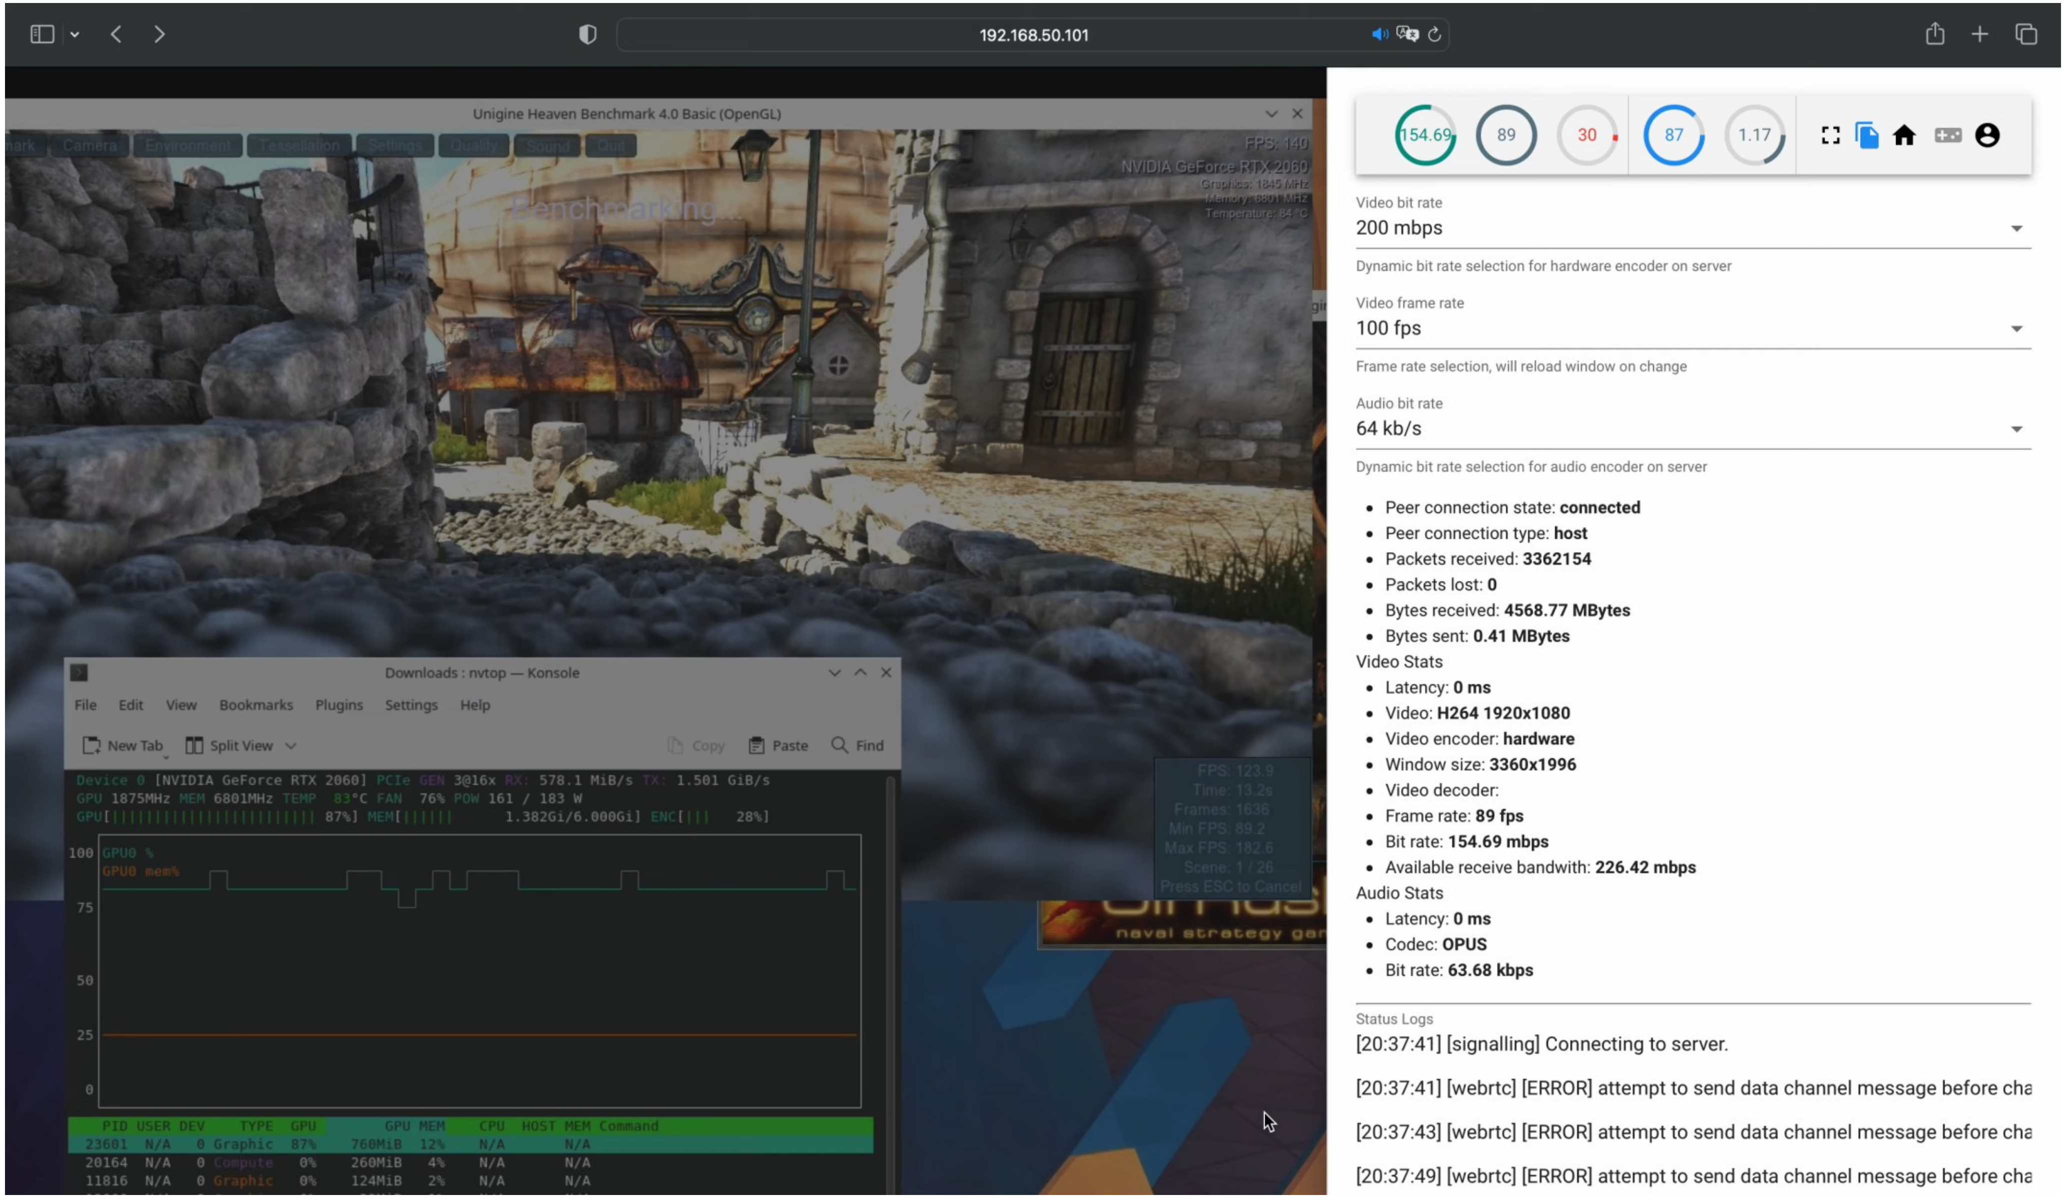

Please watch from YouTube or the PowerPoint for real-life quality

Coder: VS Code IDE in your web browser

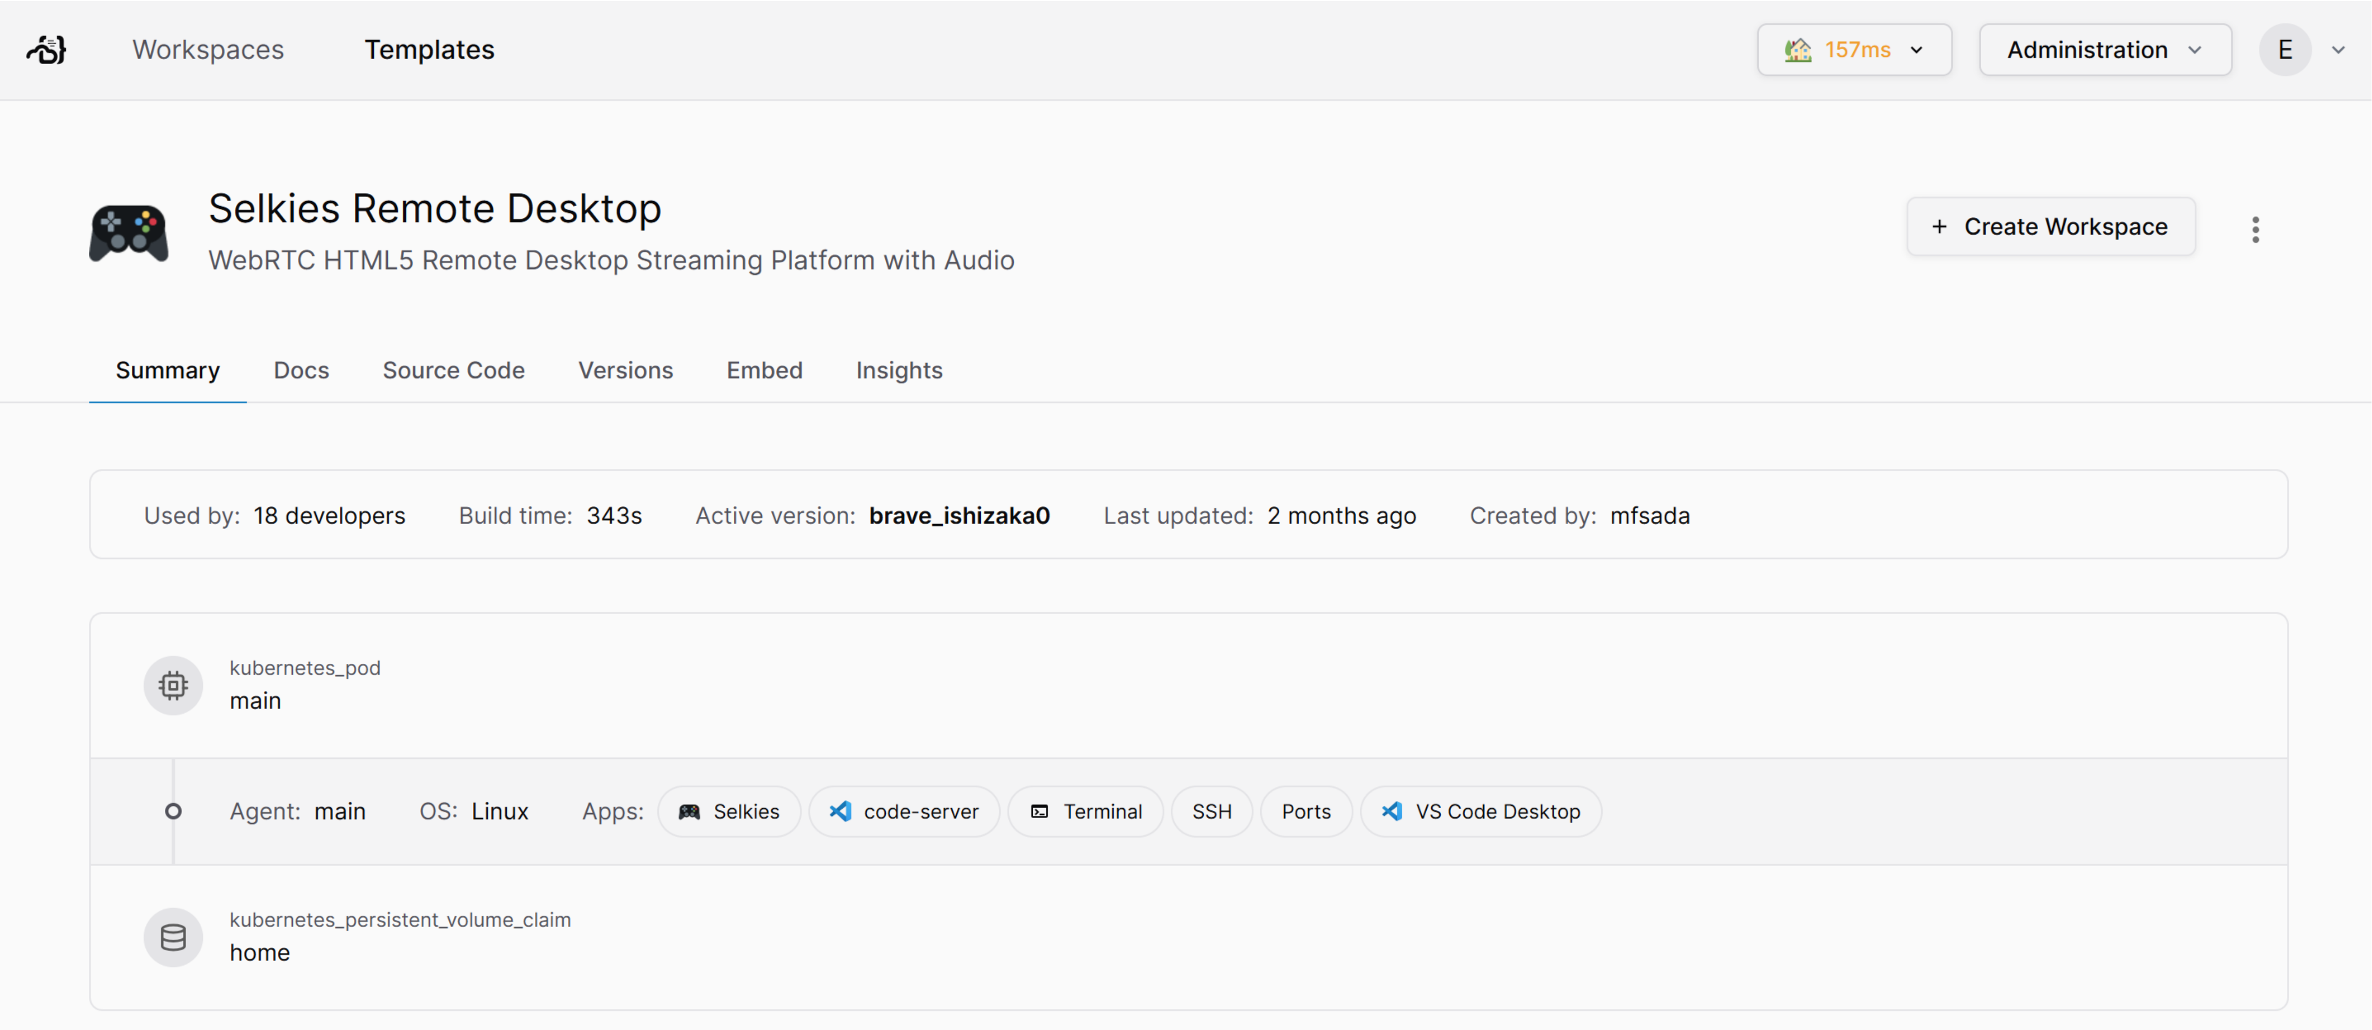

https://coder.nrp-nautilus.io/templates/selkies

Extra Resources

- Supplementary slides attached
- Introduction to Nautilus: https://docs.google.com/presentation/d/1GMvaZr9Nm6LhYUU_E0E0LdoebPpk0dgb2Z6oS9v2Ww8/edit?usp=sharing
- Resources: https://portal.nrp.ai/documentation/userdocs/start/resources/
- NRP CUDA/Desktop Container Images: https://gitlab.nrp-nautilus.io/nrp/scientific-images
- NRP JupyterHub: https://portal.nrp.ai/documentation/userdocs/jupyter/jupyterhub-service/
- Selkies-GStreamer: https://github.com/selkies-project/selkies-gstreamer
- docker-nvidia-glx-desktop: https://github.com/selkies-project/docker-nvidia-glx-desktop
- docker-nvidia-egl-desktop: https://github.com/selkies-project/docker-nvidia-egl-desktop

Thank you!

This project has been developed and is supported in part by the National Research Platform (NRP) and the Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI) at the University of California, San Diego, by funding from the National Science Foundation (NSF), with awards #1730158, #1540112, #1541349, #1826967, #2138811, #2112167, #2100237, and #2120019, as well as additional funding from community partners, infrastructure utilization from the Open Science Grid Consortium, supported by the National Science Foundation (NSF) awards #1836650 and #2030508, and infrastructure utilization from the Chameleon testbed, supported by the National Science Foundation (NSF) awards #1419152, #1743354, and #2027170. This project has also been funded by the Seok-San Yonsei Medical Scientist Training Program (MSTP) Song Yong-Sang Scholarship, College of Medicine, Yonsei University, the MD-PhD/Medical Scientist Training Program (MSTP) through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea, and the Student Research Bursary of Song-dang Institute for Cancer Research, College of Medicine, Yonsei University.

Supplementary Slides

Selkies

Bringing the governance of remote HPC back to the community

Seungmin Kim

Yonsei University College of Medicine

(Ex-) San Diego Supercomputer Center, University of California, San Diego

The story starts…

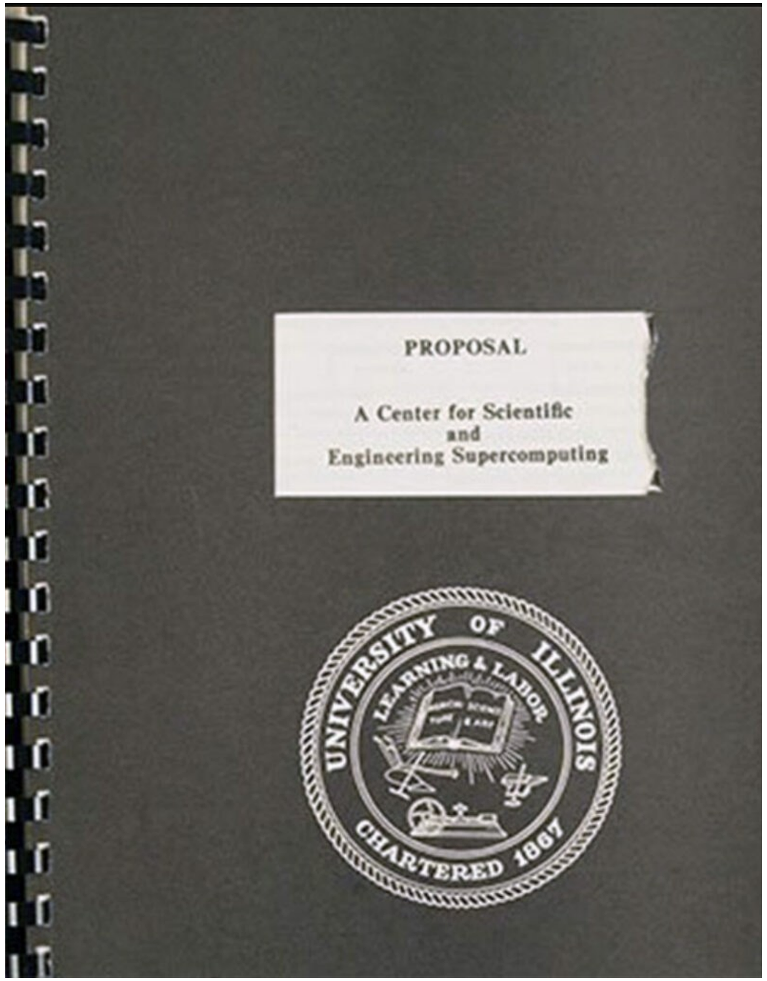

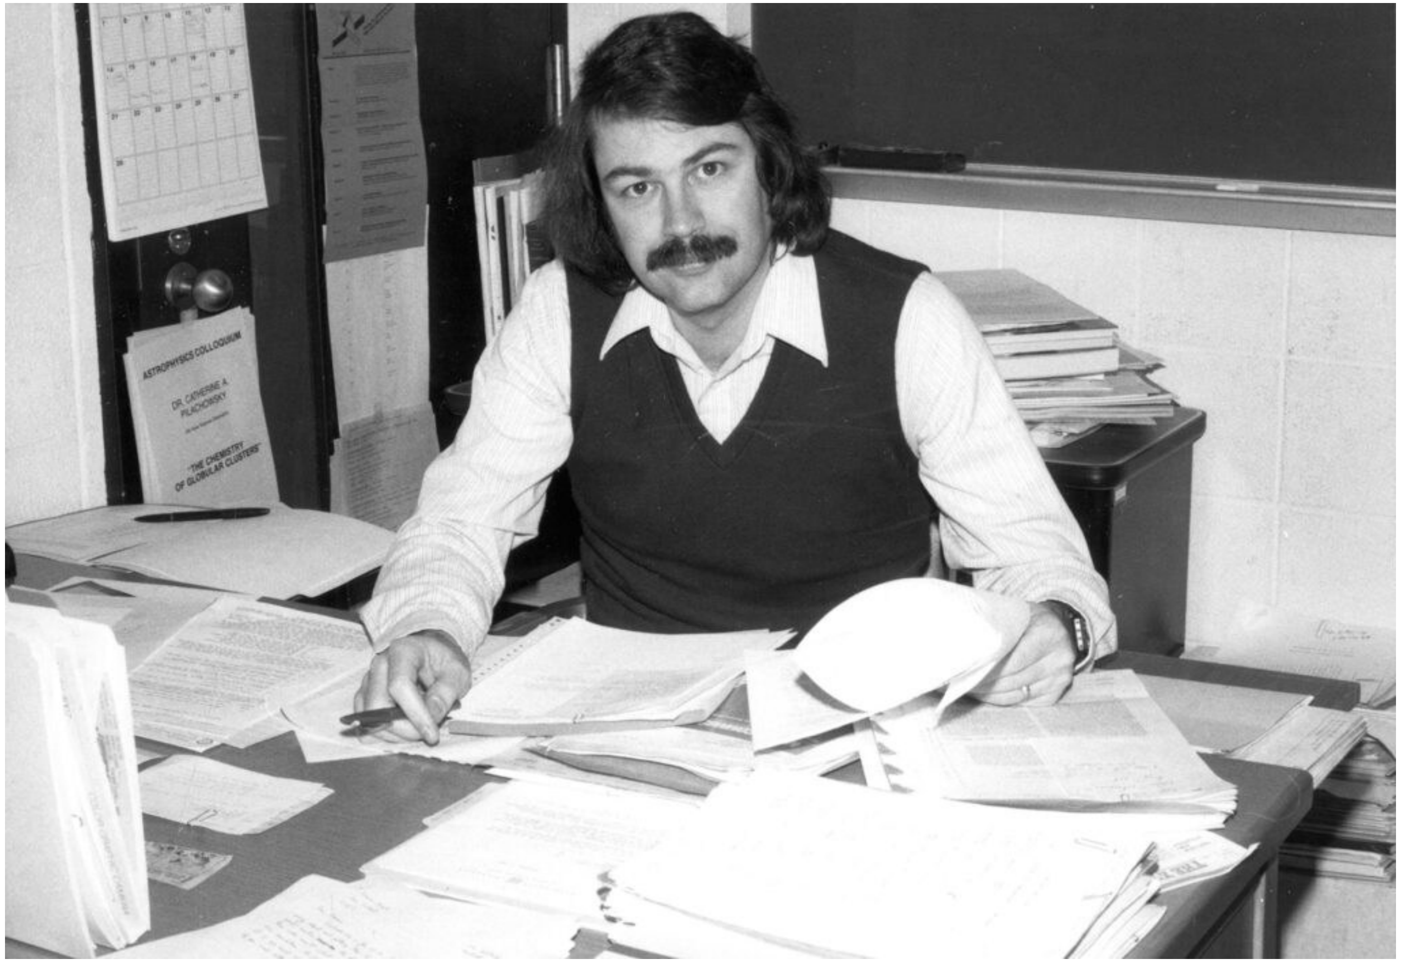

Dr. Larry Smarr, Ph.D., and the first five supercomputers

https://www.ncsa.illinois.edu/homecoming/

Now, when did everyone use computers at home?

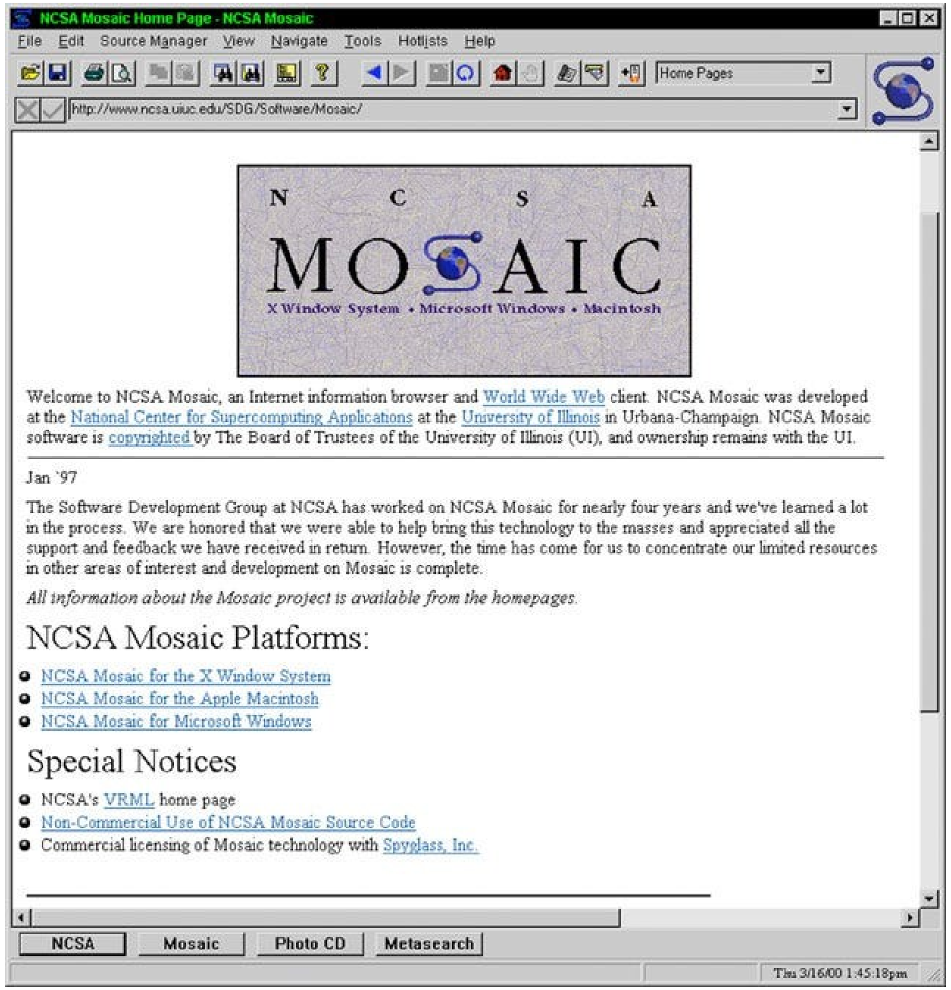

NCSA Mosaic, 1993

https://www.internethistorypodcast.com/2014/01/mosaic/

Netscape and Internet Explorer

- Mosaic was arguably the moment the computer became a commodity from an enthusiast machine
- The internet was a key traction in making computers an essential home appliance from a limited user pool
- The GUI was always the essential component for user accessibility
- But Netscape and Microsoft made the Internet proprietary by dumping non-standard extensions until Google Chrome overtook IE in market share
- The internet decomposed centralized clusters into individual hands

The Drift of the Proprietary Internet to Open Source

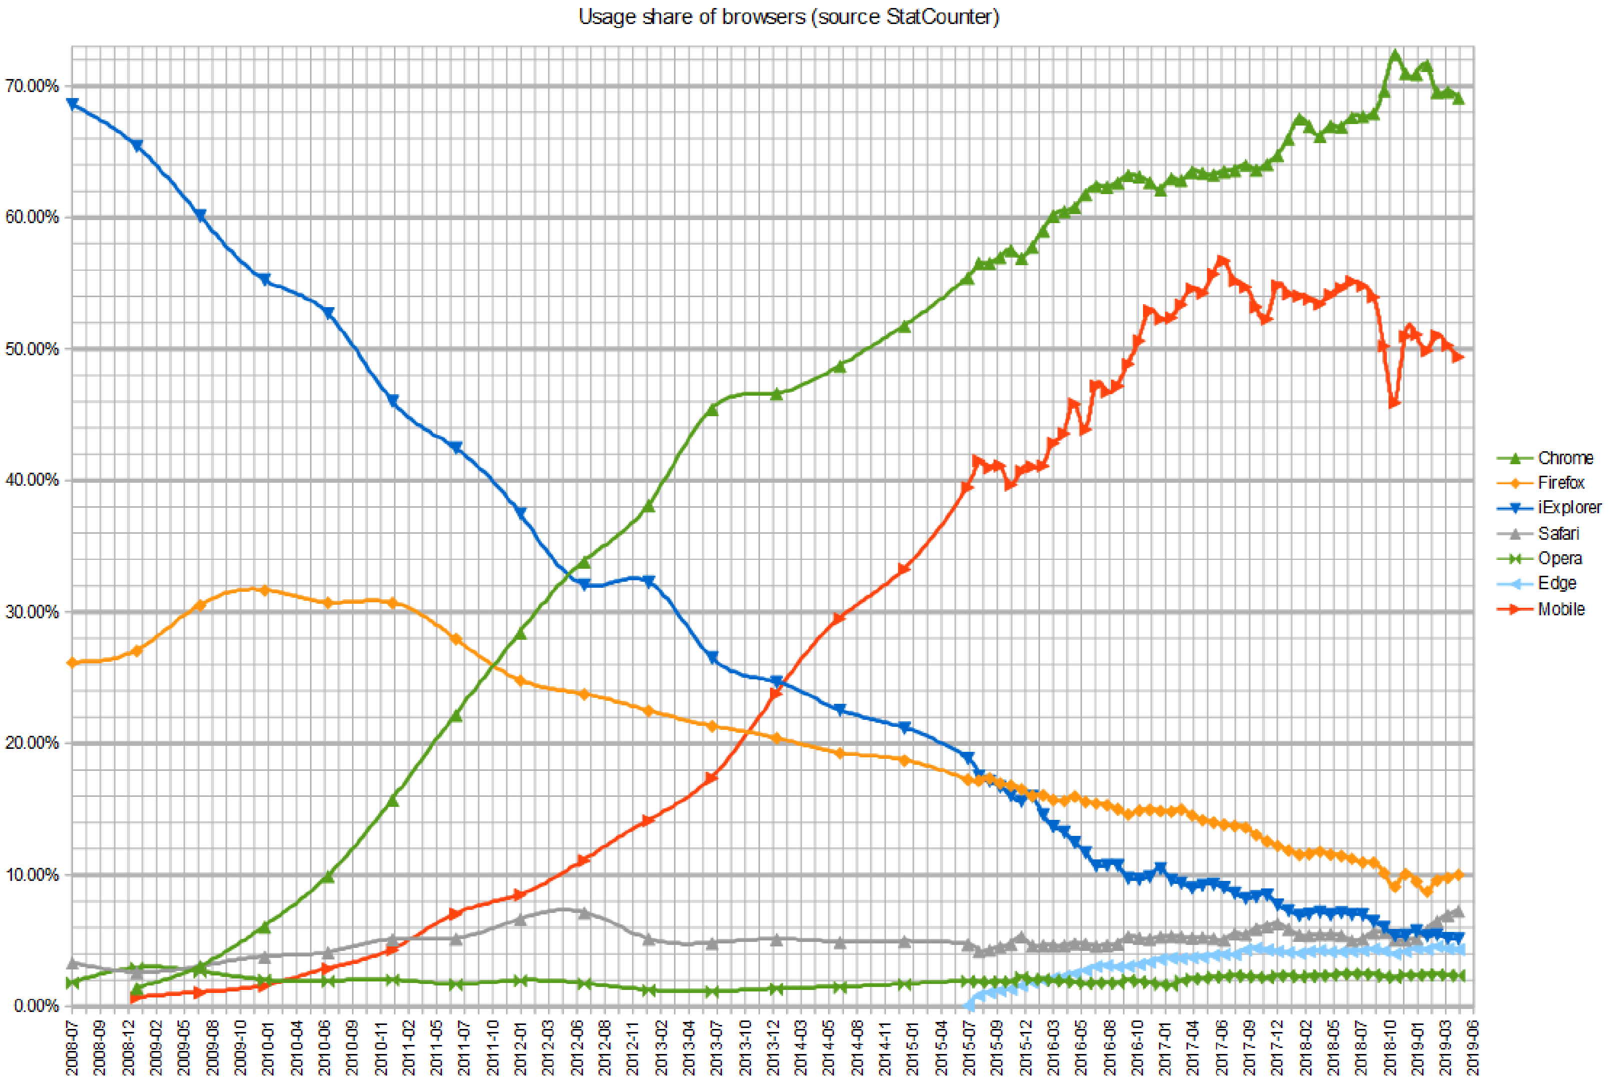

https://commons.wikimedia.org/wiki/File:BrowserUsageShare.png

Age of Clusters and Supercomputers… Again

1985-8, NSF Office of Advanced Cyberinfrastructure

https://x.com/NCSAatIllinois/status/1783463523823071742

2023, NSF Office of Advanced Cyberinfrastructure

National Artificial Intelligence Research Resource Pilot

https://www.nsf.gov/news/nsf-partners-kick-nairr-pilot-program

AI Revives Linux & Clustered Computing

- In the 2010s, the cloud era arrived, decomposing physical servers into logical abstractions, with Linux as its champion
- Large Language Models and other deep learning paradigms reward massive scale; not a personal workstation, a whole cluster
- Containers (Docker/Kubernetes/Apptainer) became the AI workload abstraction platform of choice
- AI workloads, therefore, are now commonly distributed from a central source such as ACCESS for individuals to access remotely
- HPC and supercomputers, with Linux, are now again, seriously relevant to everyone's lives

AI Revives Clustered Computing

Frontier Supercomputer, Oak Ridge National Laboratory

https://www.flickr.com/photos/olcf/52117623843/

Eos NVIDIA DGX SuperPOD, NVIDIA Corporation

https://resources.nvidia.com/en-us-dgx-systems/eos-blog

https://www.youtube.com/watch?v=G4P6cc3vq7I

WIFIRE (PI: Dr. Ilkay Altintas)

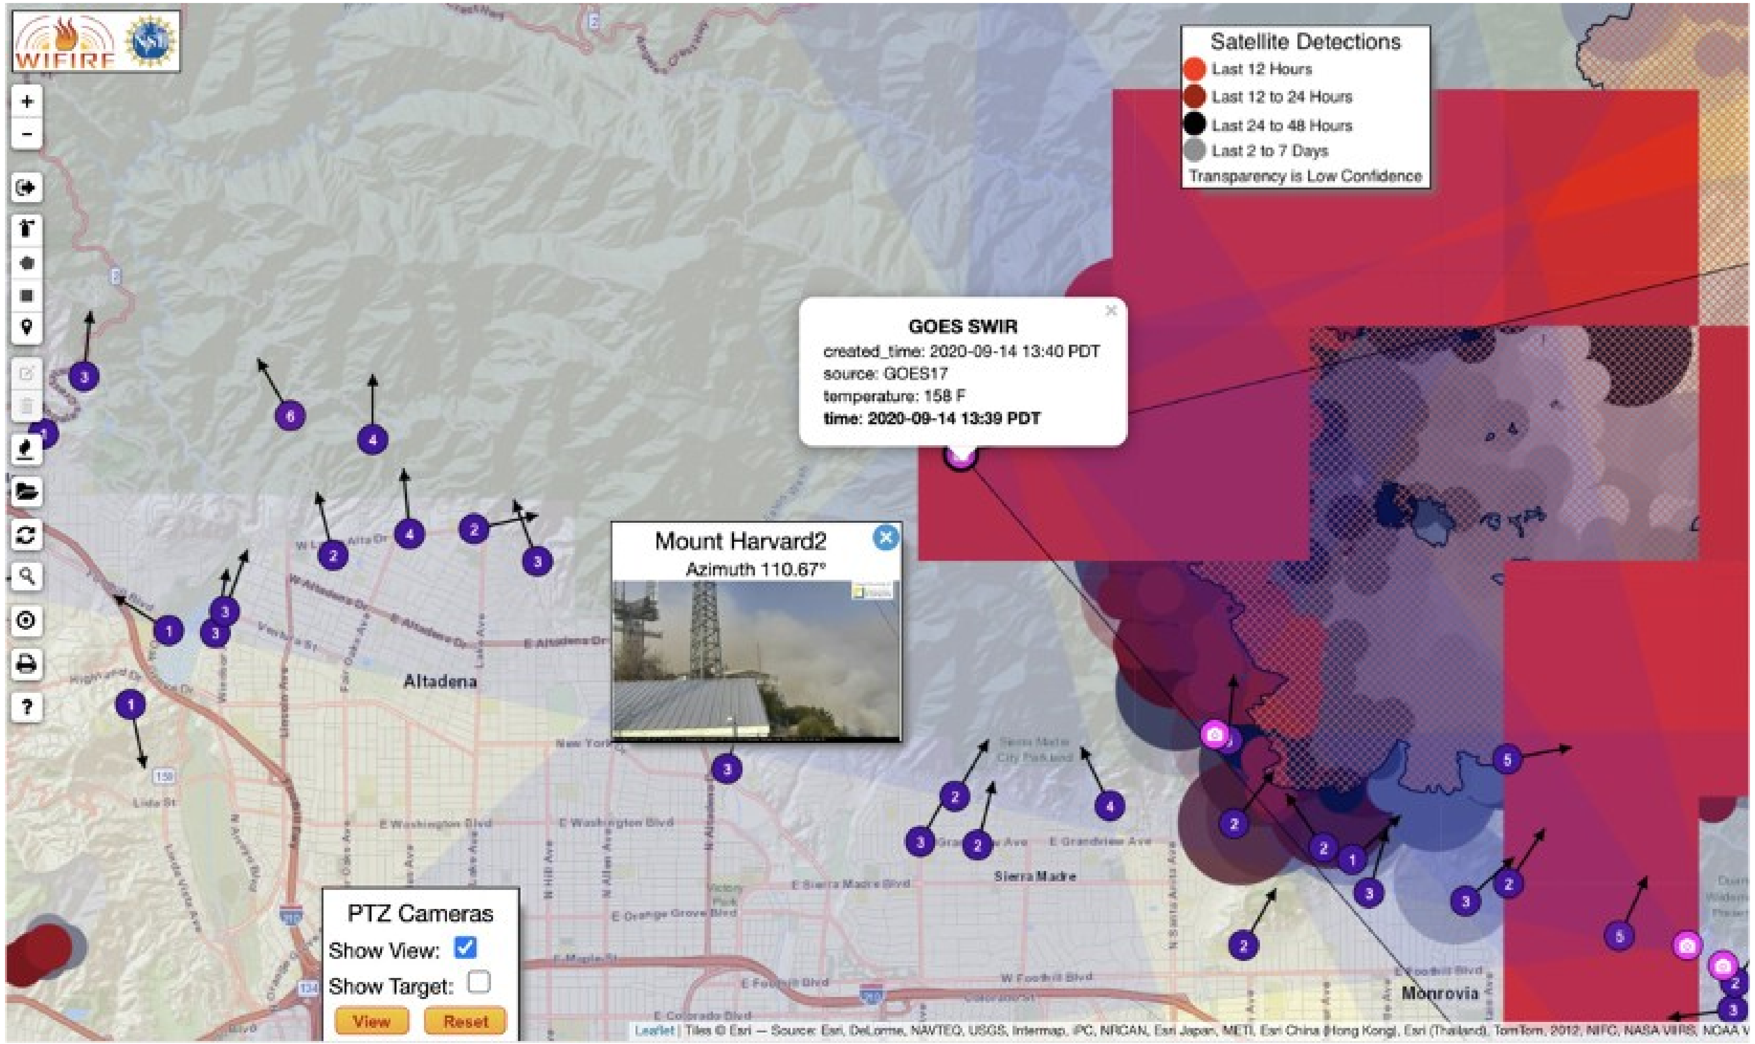

https://www.zdnet.com/article/how-ai-is-saving-homes-and-lives-in-california-during-wildfire-season/

SunCAVE (PI: Dr. Thomas DeFanti & Neil Smith)

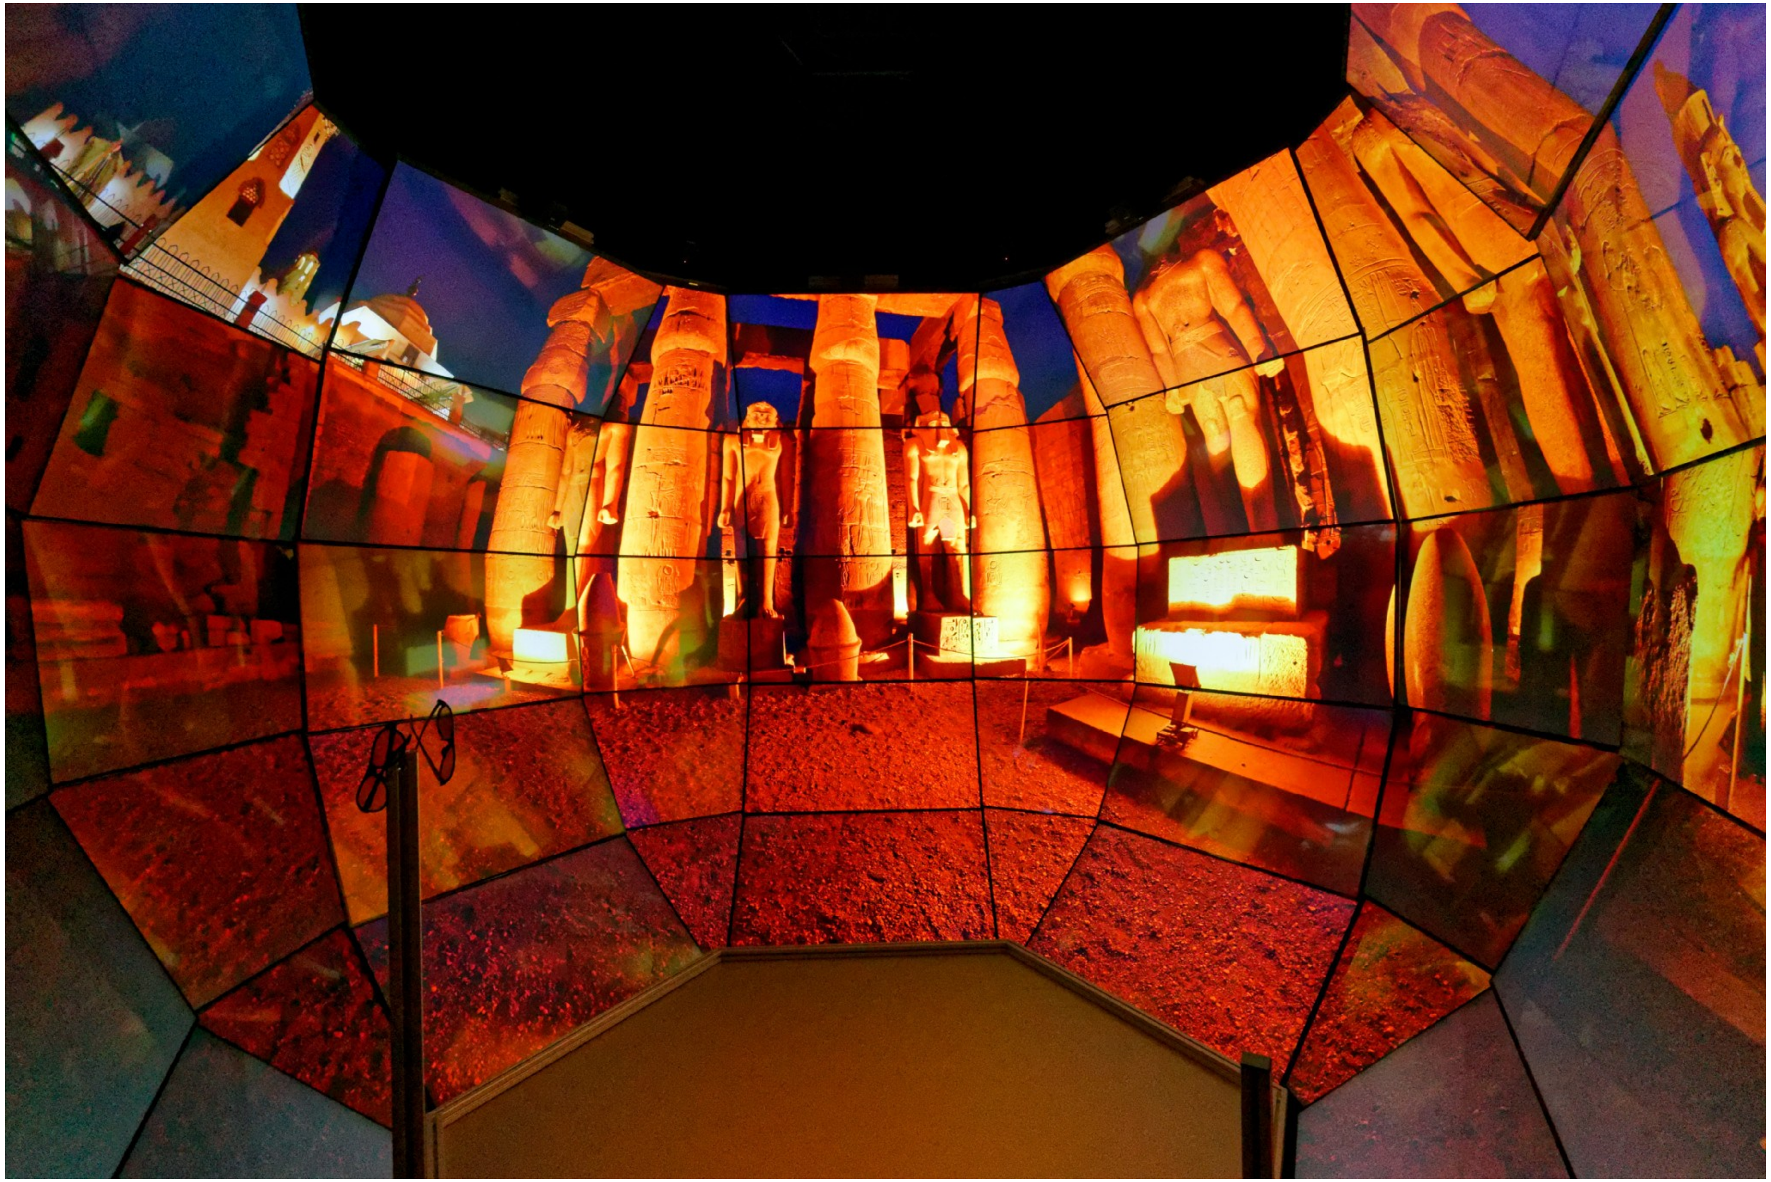

https://www.hpcwire.com/off-the-wire/qualcomm-institutes-cave-turns-two/

Remote Desktop Landscape

- People wanted to access mainframes from their house.
- Started from command-line UIs (Telnet, SSH, etc.).
- X11 Forwarding / Microsoft Remote Desktop Protocol (send primitive CopyRect draw requests directly to client)
- Virtual Network Computing (VNC) (mix CopyRect/Hextile with Zlib and JPEG, nowadays primarily JPEG)
- libjpeg-turbo enables live-speed JPEG framebuffer encoding
- Xpra (similar to JPEG, evolved to utilize WEBP/H.264 and hardware encoding)

The Usefulness of H.264 Encoding in a VNC Environment

- 2014, D. R. Commander (https://turbovnc.org/About/H264)
- Author of VirtualGL, TurboVNC, libjpeg-turbo
- There are still only a handful of datasets for which H.264 can be shown to compress better than the TurboVNC encoder, …, the improvement is now very compelling. The x264 library still takes up too much CPU time, however…
- This research did show that H.264 can be beneficial for certain types of apps…that generate somewhat game-like or video-like image sequences. X servers won't be around forever.
- Apps will increasingly stop using X primitives in favor of X RENDER and other image-based drawing methods…will become more H.264-friendly as time progresses.

The Remote Game Streaming Scene

- OnLive (2009), GamingAnywhere (2013), Parsec (2016)
- H.264 video encoding in both hardware (GPU) and software (CPU) overtook the efficiency of JPEG and X11 CopyRect
- Microsoft RDP RemoteFX Media Streaming (H.264) in 2016
- Google Stadia implemented capabilities for low-latency interactive computing for Chrome, later in other browsers
- NVIDIA GeForce Now, Microsoft Xbox Cloud, Amazon Luna (2018-2020)
- Subsequent adaptation to professional graphics & visualization during COVID-19
- Rise of the scene just when something was not enough with TightVNC (back in 2017-18) for remote accessing Molecular Dynamics simulations

The Remote Game Streaming Scene

OnLive, 2009

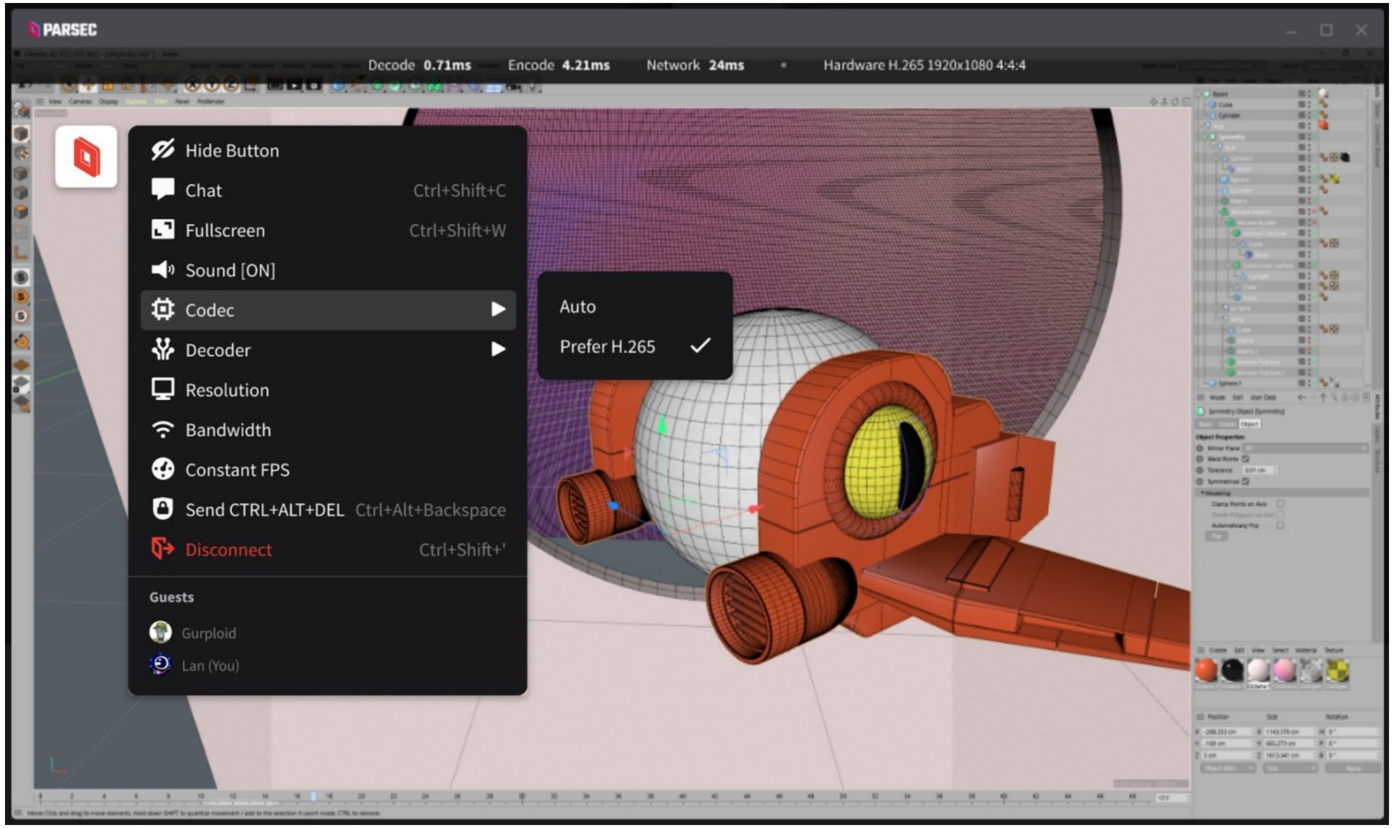

Parsec, 2016

https://techcrunch.com/2011/12/07/onlive-ios-android/

https://techcrunch.com/2020/12/15/parsec-series-b/

Current State of Linux Remote Desktops

- FOSS: RealVNC, rdesktop, TigerVNC, TurboVNC, Xpra, Xrdp, RustDesk…
- Either rely on VNC-like protocols or are meant for TeamViewer-like remote user support
- No FOSS solution in Linux maintaining server-client architecture to deploy on massive clusters AND providing modern performance, except for VNC – RFB is nearly 30 years old
- Paid: NoMachine, NICE DCV, Teradici, Reemo, Citrix…
- All have H.264 implemented for Linux 3D graphics
- Alternative: Moonlight + Sunshine – FOSS but designed for individual game streaming from desktops

Issue with Proprietary Remote Desktops

- License Servers and Seats! MATLAB is already painful enough
- The issue here is flexibility, not the fees
- Cannot bundle into a portable environment (Snapshot vs Container)
- Cannot customize deployment to various environments; forces a certain environment configuration the program wants
- OS and architecture support is also inflexible

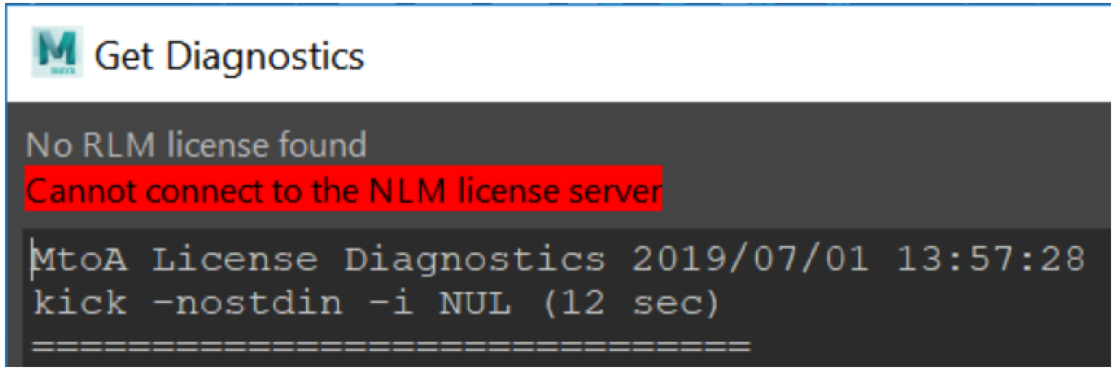

https://www.autodesk.com/support/technical/article/caas/sfdcarticles/sfdcarticles/No-RLM-license-found-Cannot-connect-to-the-NLM-license-server-doing-Get-Diagnostics-on-Maya-trying-to-get-an-Arnold-license.html

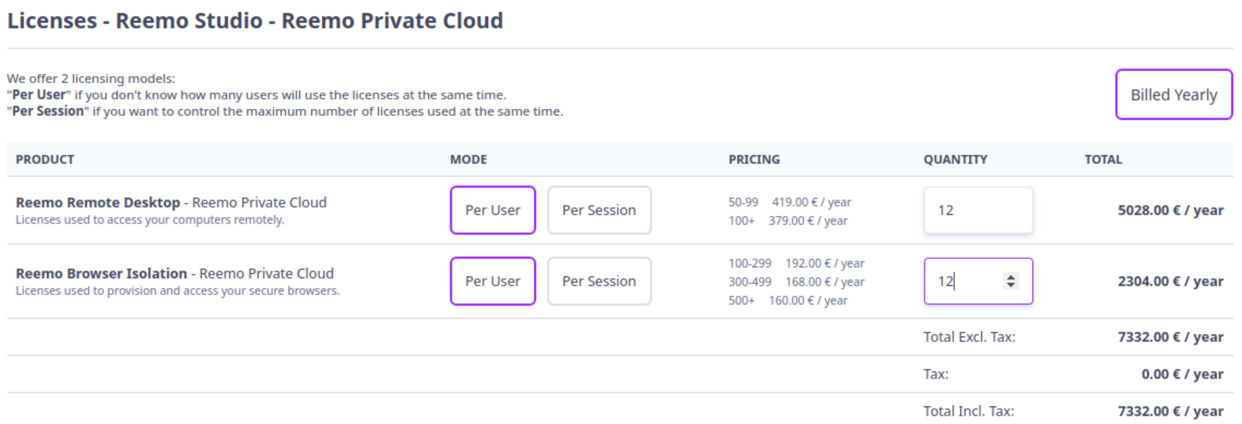

Licenses - Reemo Studio - Reemo Private Cloud

We offer 2 licensing models:
"**Per User**" if you don't know how many users will use the licenses at the same time.
"**Per Session**" if you want to control the maximum number of licenses used at the same time.

Billed Yearly

PRODUCT	MODE	PRICING	QUANTITY	TOTAL
Reemo Remote Desktop - Reemo Private Cloud Licenses used to access your computers remotely.	Per User / Per Session	50-99 419.00 € / year 100+ 379.00 € / year	12	**5028.00 € / year**
Reemo Browser Isolation - Reemo Private Cloud Licenses used to provision and access your secure browsers.	Per User / Per Session	100-299 192.00 € / year 300-499 168.00 € / year 500+ 160.00 € / year	12	**2304.00 € / year**
			Total Excl. Tax:	**7332.00 € / year**
			Tax:	**0.00 € / year**
			Total Incl. Tax:	**7332.00 € / year**

https://reemo.io

Current State of Linux Remote Desktops

https://web.archive.org/web/20210425211725if_/https://www.youtube.com/embed/XWqb3iTs1lw
https://web.archive.org/web/20210310083658/https://cloud.google.com/solutions/gpu-accelerated-streaming-using-webrtc

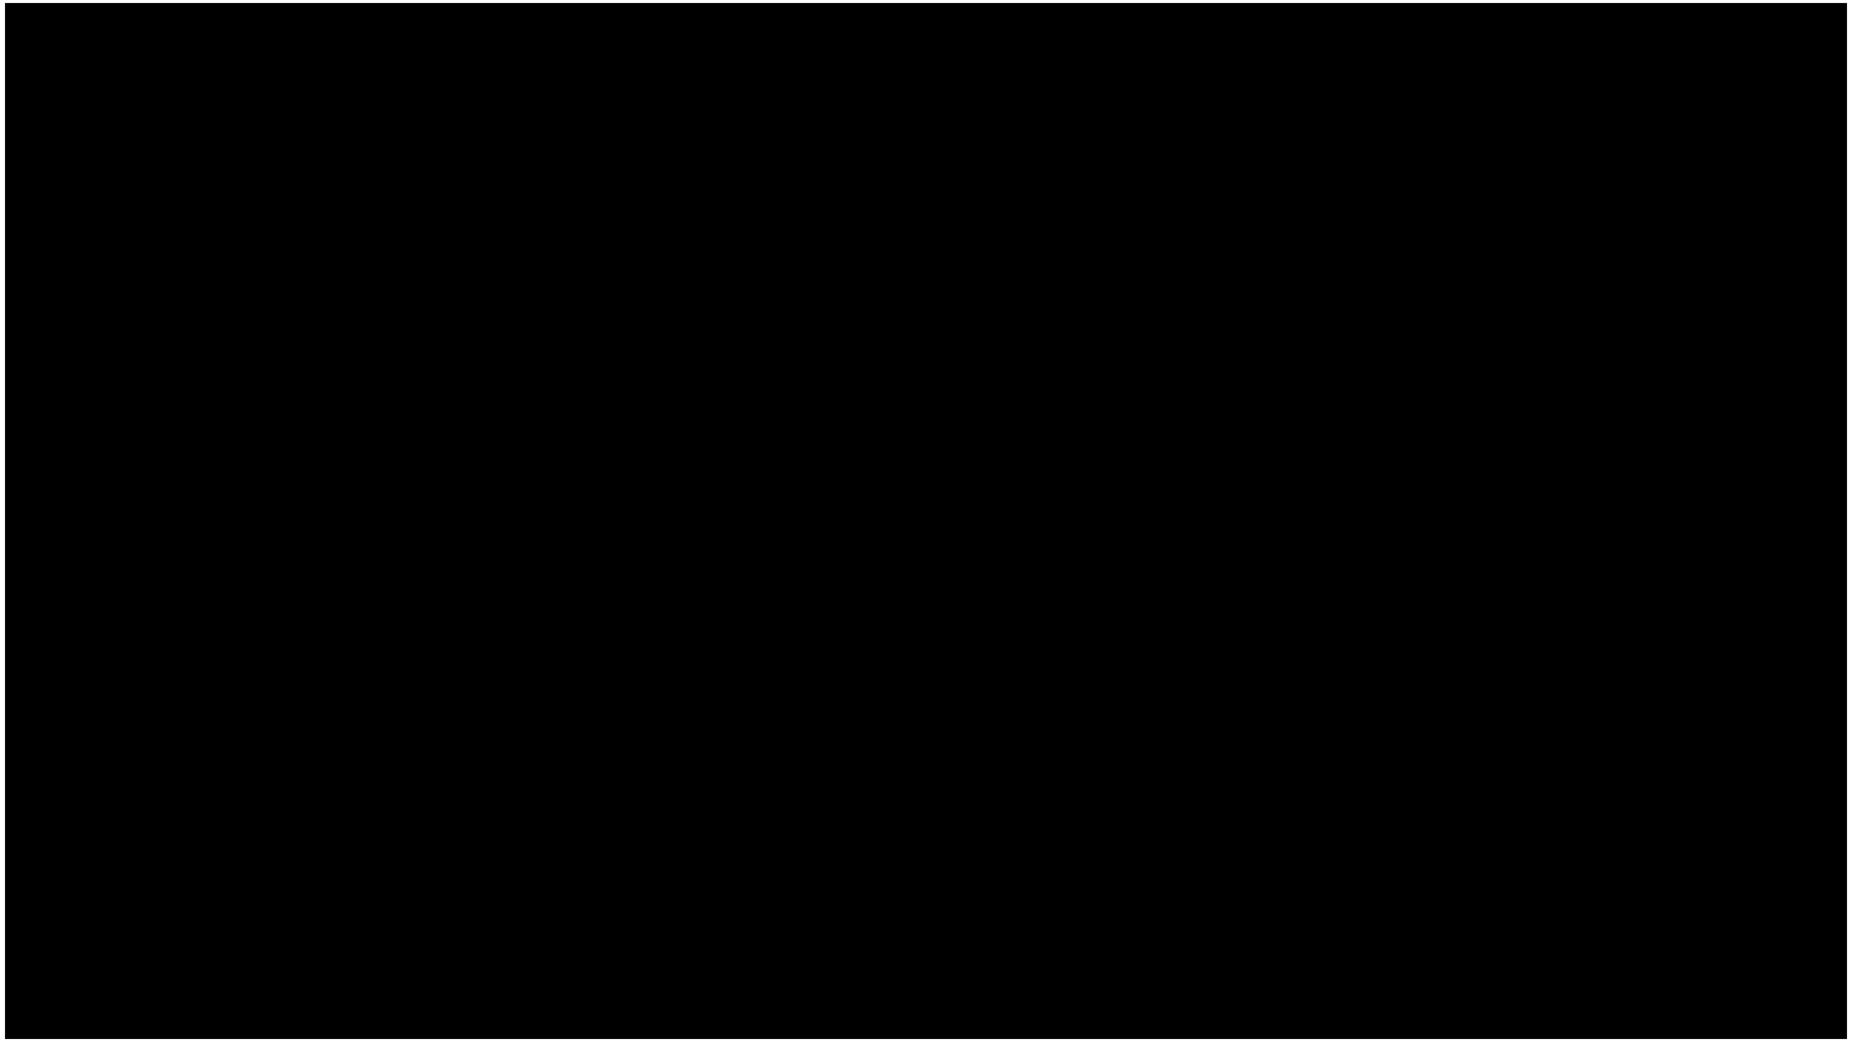

Google, Dan Isla, 2019-20

Design of Selkies

- Usage of HTML5 and WebRTC; anywhere Chromium, Firefox, Safari exists is a client (including embedded platforms)
- Utilizes the same techniques Google Stadia, NVIDIA GeForce Now, etc. used, obtaining equal or better latency and quality to commercial products
- Modularized: Web Client, Server, and Media Processor (GStreamer) are in separate components, compile everything anywhere
- Designed for Containers and HPC clusters, with minimum intervention from administrators (Kubernetes, (rootless) Docker/Podman, Apptainer)
- Open-Source License: MPLv2 allows flexibility to customize all code, but with a weak copyleft license
- Hardware GPU acceleration, both for the H.264 stream and the workload

Selkies Offerings

- docker-nvidia-egl-desktop (uses Xvfb + VirtualGL for HW acceleration of workloads), docker-nvidia-glx-desktop (X.Org server with NVIDIA X11 drivers inside the container)
- Portable distribution, compatible with any Linux OS glibc 2.17 or higher (based on the Conda build toolchain for distribution compatibility, but the user does not use Conda)
- Method to compile and/or install the Python Server, Web, and GStreamer Media Framework separately in any architecture
- Web interface is compatible with all major web servers and reverse proxies (HAProxy, NGINX, Jupyter, etc.), where noVNC is originally used

Selkies in modern Linux

- 10 years later from D. R. Commander, H.264, even in software x264 encoding, is capable of 60 FPS in Full HD resolution with only 130% CPU utilization
- NVIDIA and AMD/Intel GPUs are capable of live encoding QHD or 4K
- VNC (JPEG) struggles to encode 40-45 FPS with > 2000% CPU utilization with highly active or 3D visualization workloads
- Solves the underappreciation of GUI & Visualization in Linux, where multimodal AI is accelerating a new demand
- X11 is used primarily, Wayland has issues with NVIDIA drivers (needs to enable optional kernel modules)

Bringing Remote HPC Back to Everyone

- Freedom to bundle and customize to every environment
- Highly readable code for future contributors
- Sustained by consultancy contracts (for industrial contributors) and academic involvement
- Linux gamers are also our source of feedback in terms of usability and performance; Selkies is also the containerized, web-based choice
- Citations are the currency of academia, this is why primarily proprietary technologies can become open-sourced
- Future Potential: NSF Pathways to Enable Open-Source Ecosystems
- The way forward is to bring Selkies to a de-facto standard for Interactive and Urgent HPC across all of the states

The spirit of NCSA

- NCSA was the birthplace of the precursors of Netscape Navigator, Internet Explorer, Apache HTTPd, and many other tools we use
- Despite using modern technologies such as Docker, Kubernetes, Apptainer, Conda, etc., the core of our technologies and the hardest component is still tussling with X11, old-school to this day
- This leaves the door open for many senior-level HPC personnel with decades of experience to pitch their expertise
- Selkies must become a community-sustained project to survive; input from many people are always better than one