# King's e-Research Portal

Making things easier, securely

Matt Penn 7/12/2023

# King's e-Research

- In late 2021 funding was secured to hire and build a central research computing function at King's
- Prior to this there was no institution wide research computing facilities
- 16.5 FTE
  - Director
  - Infra, Data, Information Governance and Trusted Research Environment Leads
  - 2 x Research Software Engineers (4 posts being appointed)
  - 2.5 x Senior Infrastructure Engineers
  - 6 x Research Operations Engineers
  - Admin Manager

# Senior Infrastructure Engineers

- Skylar Kelty
- Xand Meaden
- Deep infrastructure knowledge
  - Linux administration at scale
  - Software development
  - Storage
  - Hardware
- Lots of python, Puppet, PHP, ...

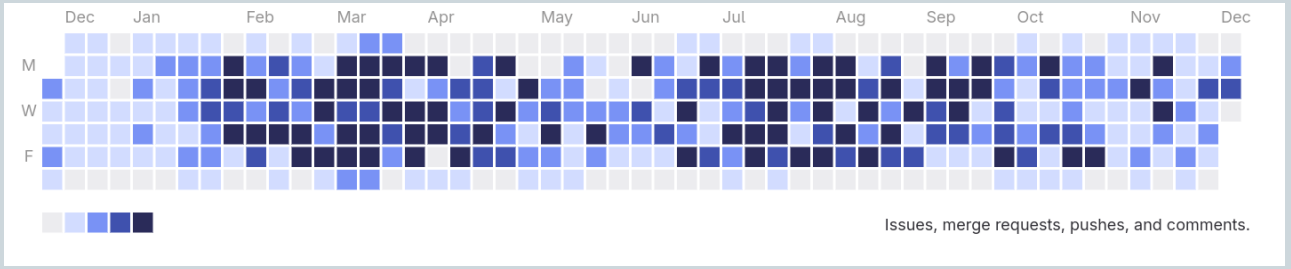


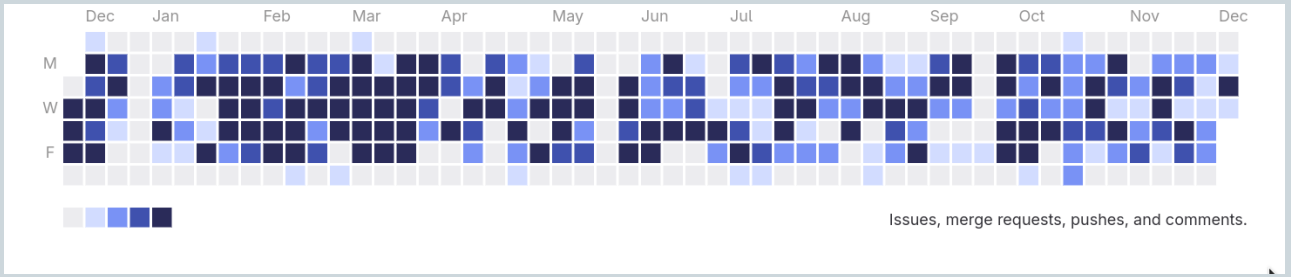

# King's CREATE

- HPC, private cloud, TRE, web farm
- Went live in April 2022
- Security is a priority
- Ease of access
- All deployment and maintenance in-house
- StackHPC for OpenStack support
- We use Ceph and Ubuntu for everything

# e-Research Portal

- In this talk
  - SSH MFA
  - Group management
  - Web proxy self-service
- Not in this talk
  - Guacamole based VDI (to Trusted Research Environments and Private Cloud)
  - Service account provisioning (for application access to shared storage)
  - Read-only account provisioning (for sharing of datasets via SFTP)
  - OpenVPN certificate generation
  - Static website, WordPress and PHP-FPM provisioning
  - OpenStack project provisioning

# e-Research Portal

- PHP/Laravel app
- Organically developed since 2021
- Very KCL specific
- Integrated with King's ID via SAML 2
  - SimpleSAMLphp
  - Portal -> idp.er -> MS "Entra" IdP
- Entra can be configured for tighter MFA timeouts on specific "apps", e.g.:
  - 7 days in general
  - 6 hours for TRE functionality
- Entra can implement "terms of use" acceptance flow

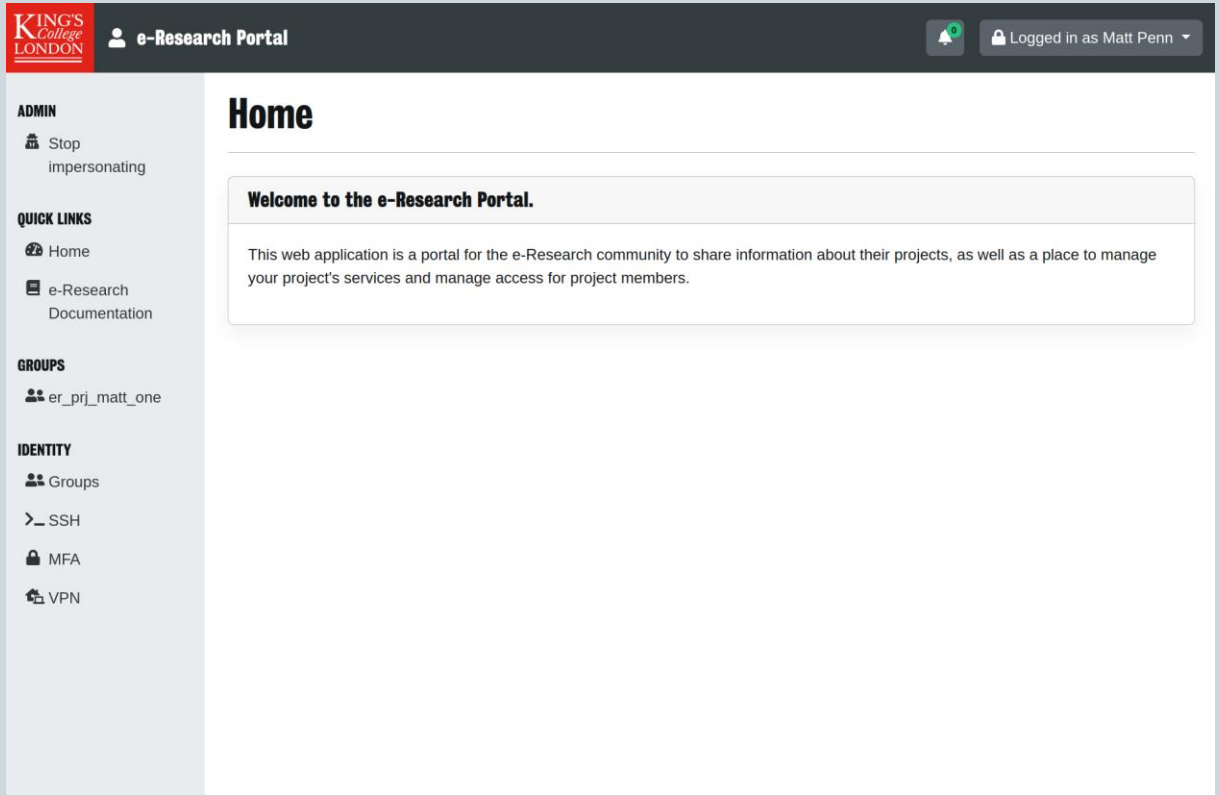

# Securing SSH access with UX in mind

- Started exploring SSH MFA in 2020
- Wanted to avoid
  - Managing OTP tokens: confusing for users, more to implement for us
  - Microsoft RADIUS: PAM plugin didn't support prompt for codes
- Suggestion from MSc student with industry experience: move the MFA to a web portal
- If we implement SSH MFA mechanism in e-Research Portal, we:
  - Utilise King's ID and token management
  - Support generic OTP clients in addition to MS proprietary
  - SSH client doesn't need to MFA on every connection

# sshd_config

AuthorizedKeysCommand

Banner


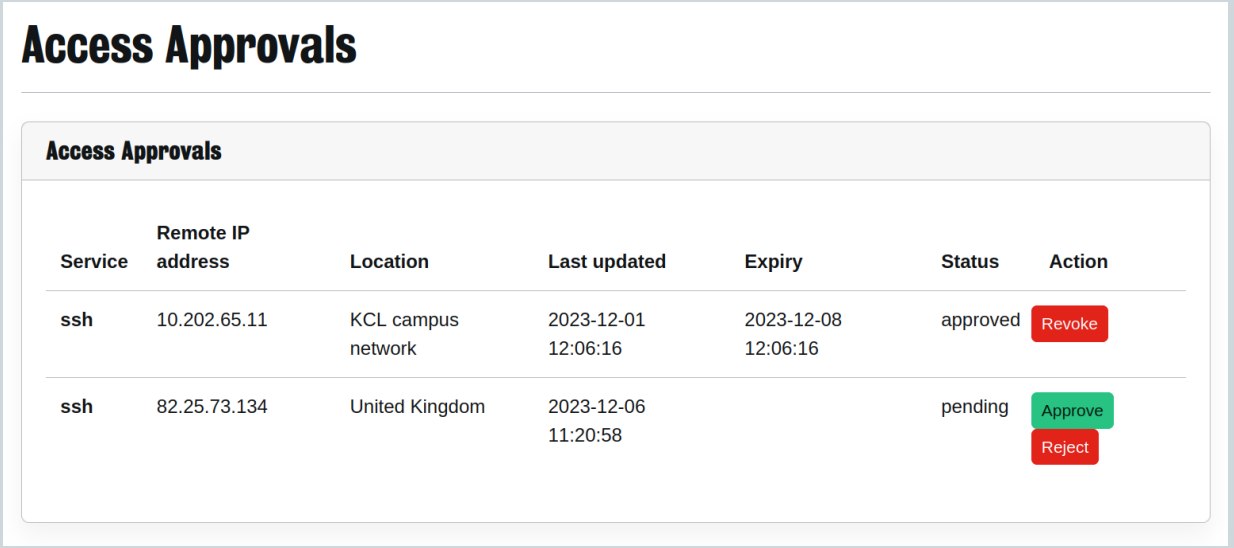

# auth_api

- python flask API w/ mysql database
- Tracks connection approve/deny actions made in Portal
- Store account specific loosenings, e.g.:
  - Removal of MFA (for limited access data ingress/egress use cases)
  - Restriction to specific commands (rsync, SFTP)
- Auxiliary scripts, e.g.:
  - bastion/logon nodes: determines source address of sshd connection and confirms MFA state in
  - Slurm nodes: check if user has a running job

https://github.com/kcl-eresearch/auth_api

# e-Research SSH/VPN Bastion Authentication

## A. Configuring credentials

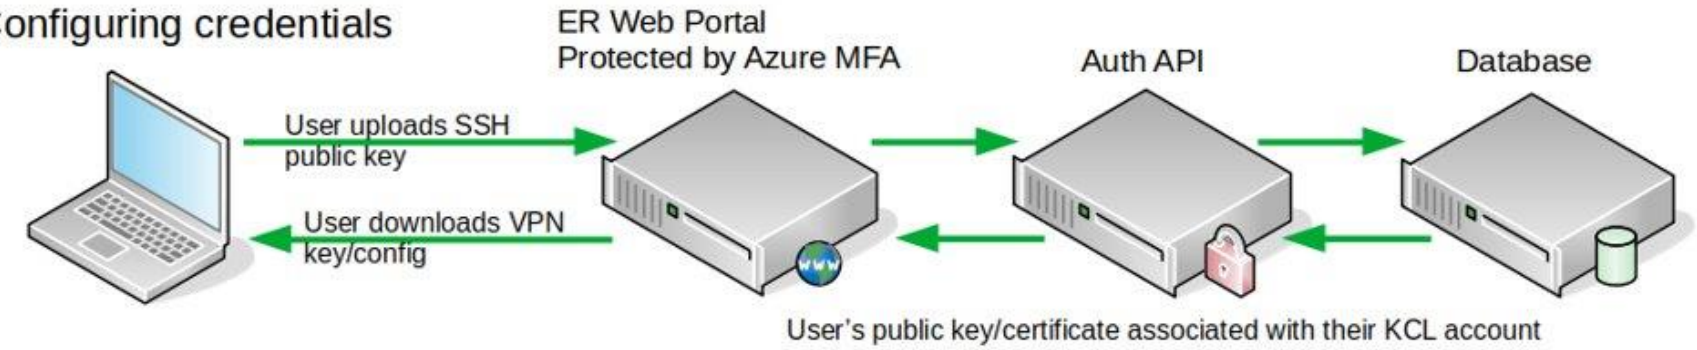


## B. Connecting by SSH/VPN

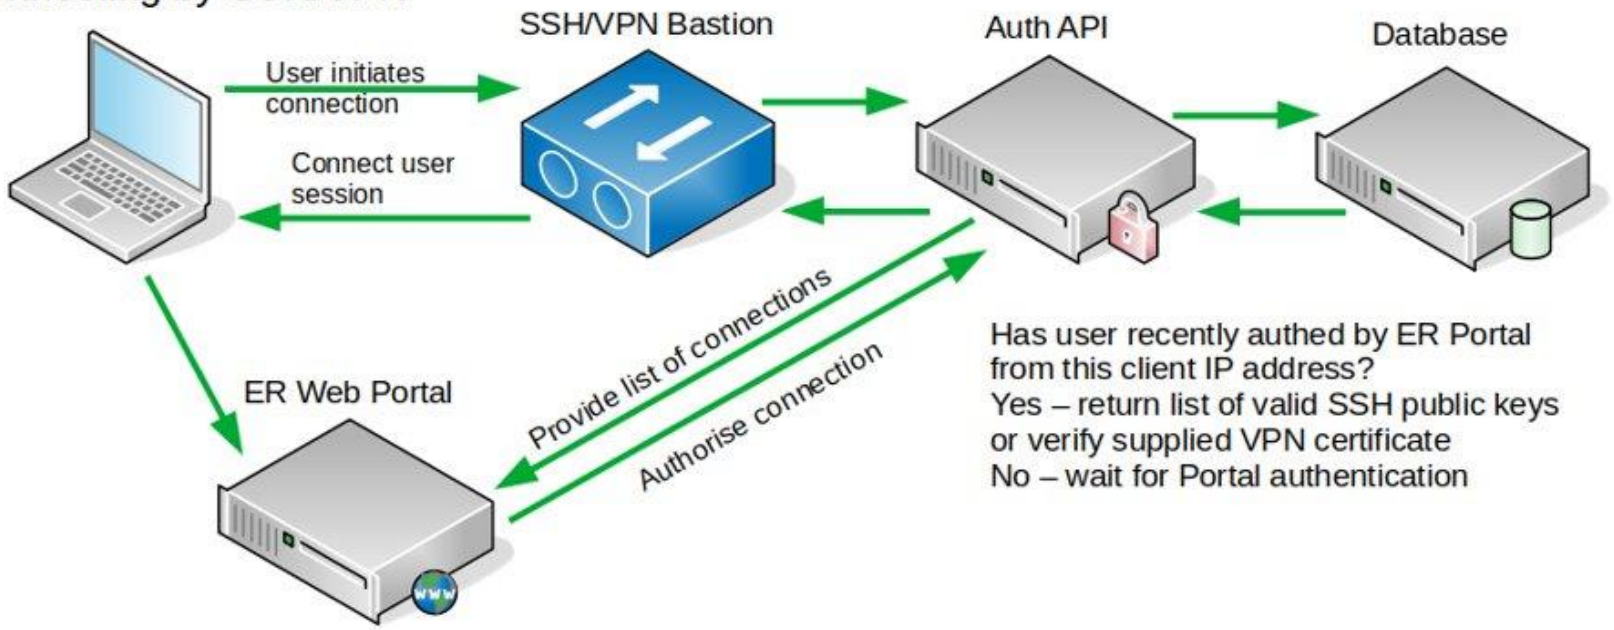

# SSH MFA access to CREATE HPC https://youtu.be/HEVx4CeIv-8

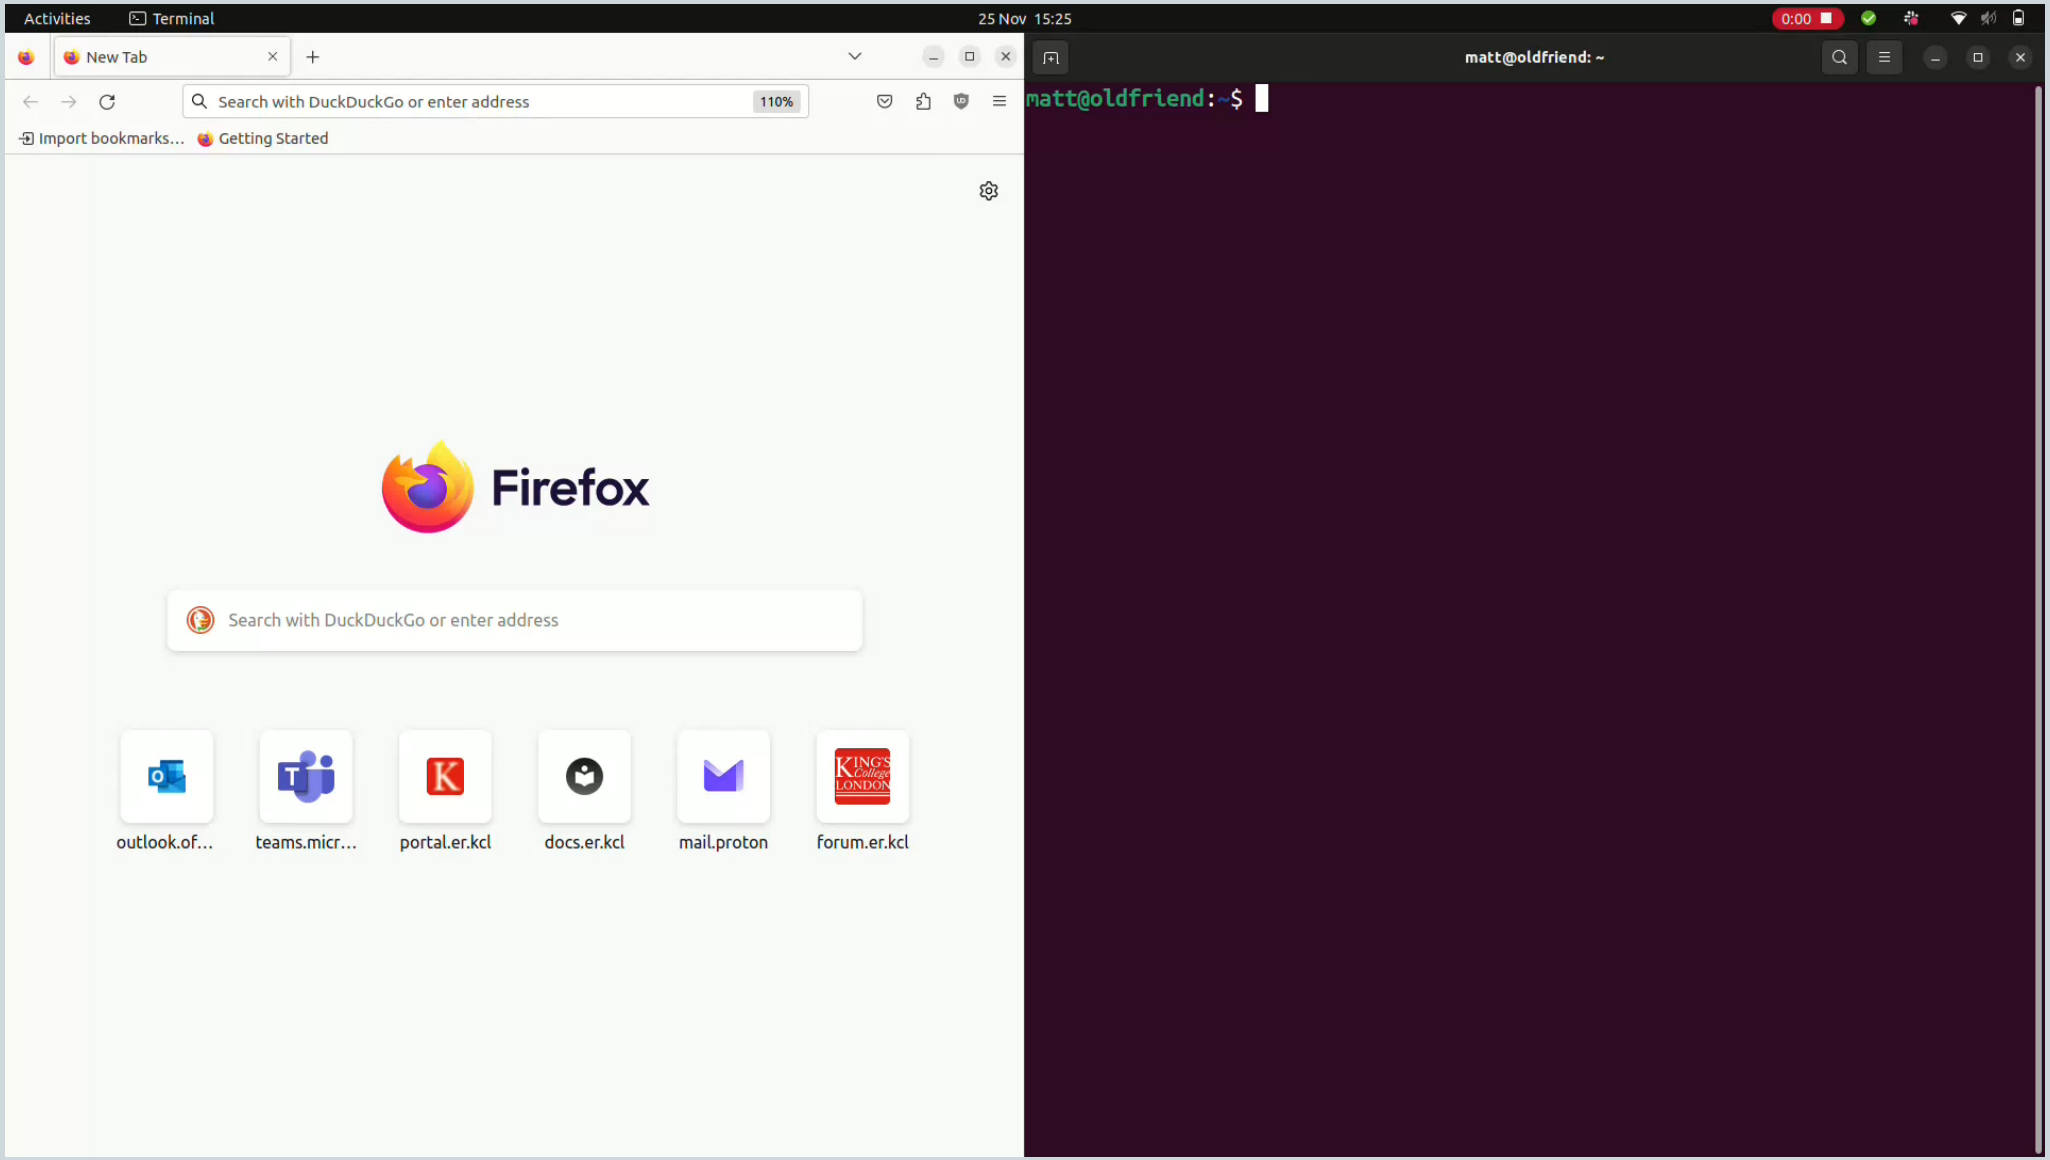

# Usage to date

- 1000+ users since CREATE HPC launched in April 2022
- 458 unique HPC logons in October
- Also used for SSH bastion (with -J / ProxyJump) to:
  - OpenStack VMs
  - On-campus SSH servers

# Group Management

- Most obvious thing to self-service well for efficiency
- Again, lean on institutional AD
- TRE increased motivations for:
  - Audit log
  - Increase context for approvals (department, user type, photo)
- Sync mechanics
  - Per action pushes to from Portal to AD (e.g. add/remove this user to group)
  - Per group sync job from AD to Portal runs every 6 hours

# Group member addition https://youtu.be/-r7AwiWifW0

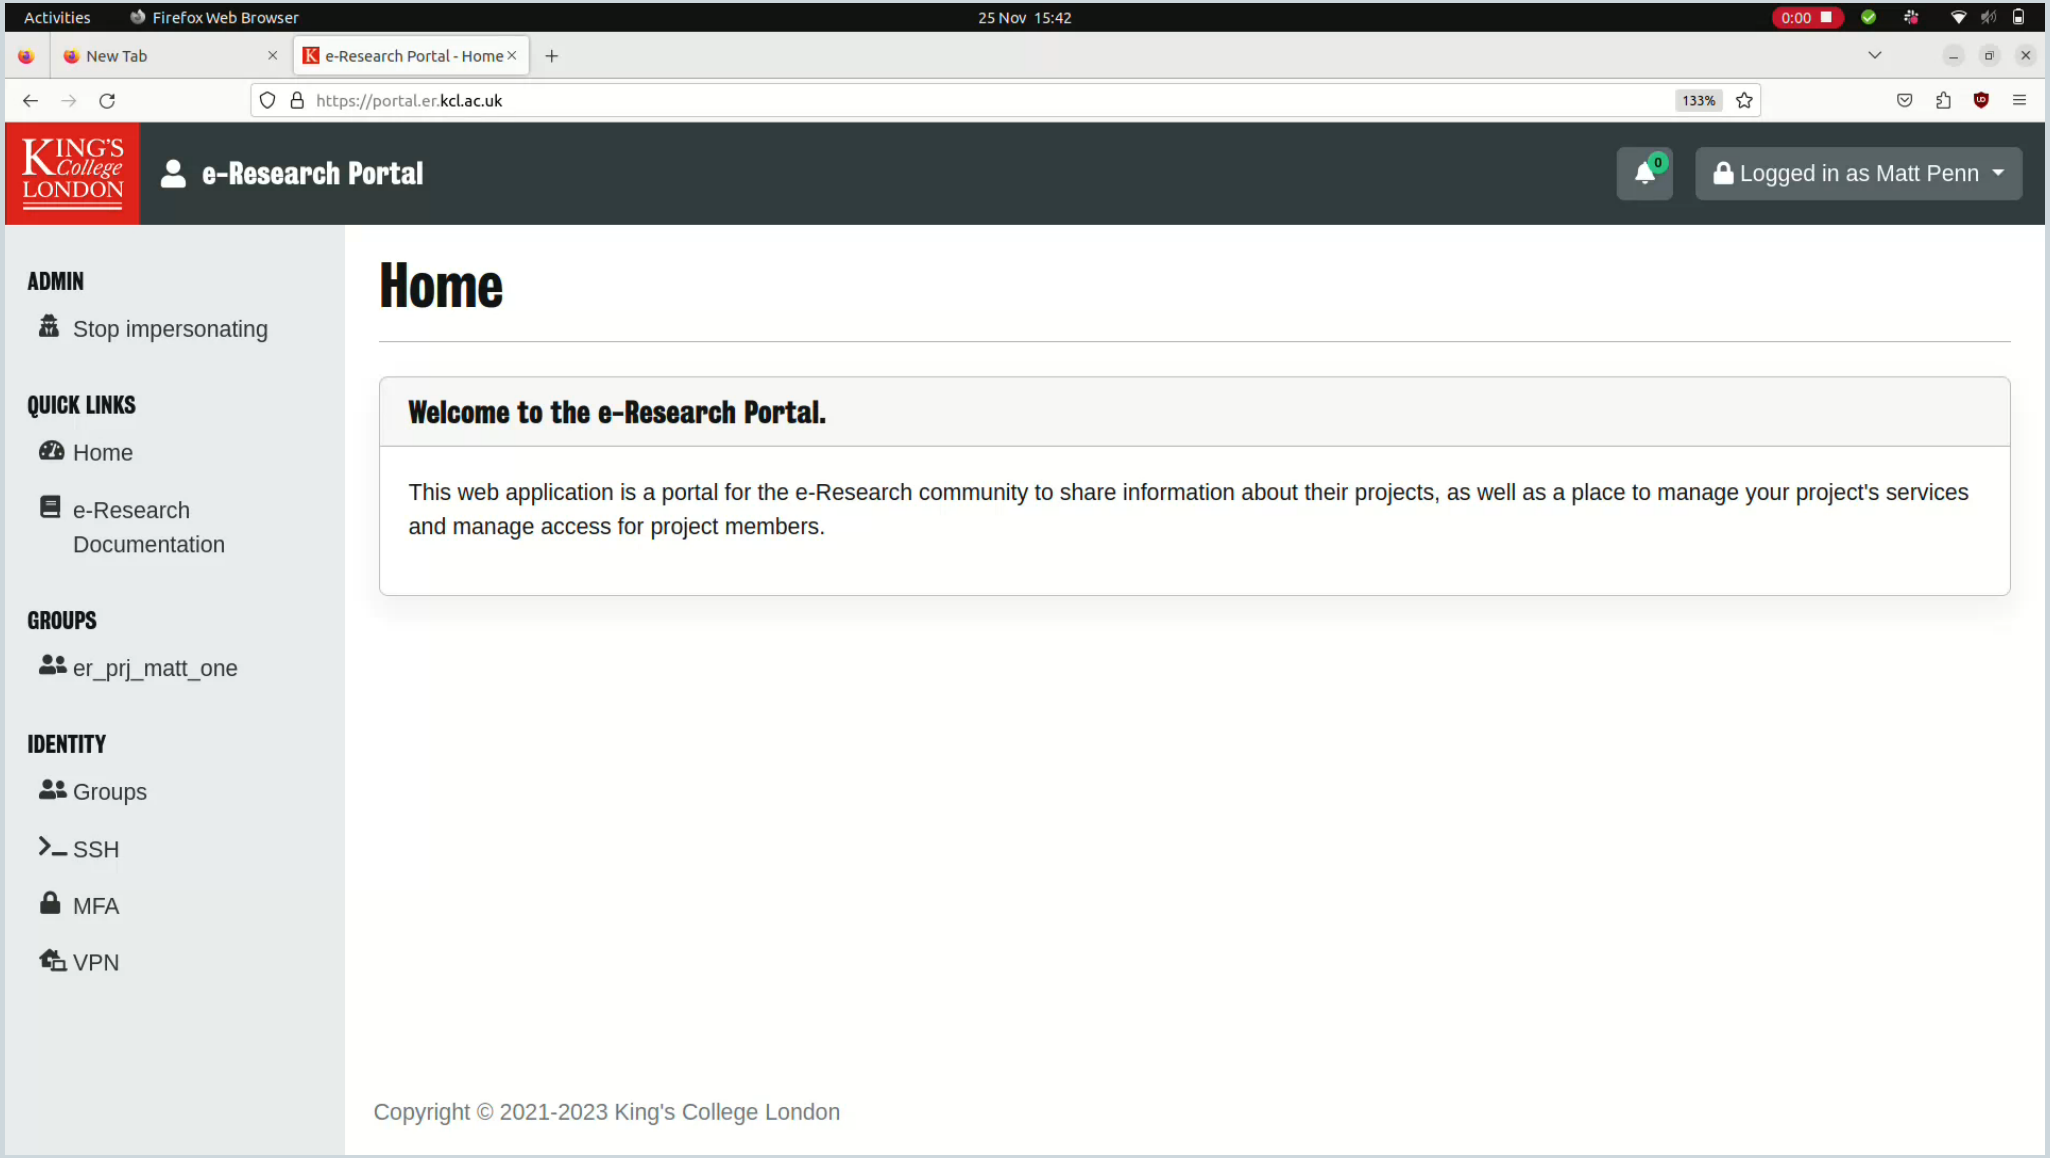

# Web proxy self-service

- Want users of our private cloud to be able to expose web applications to the outside world at ease
- NGINX / Apache plumbing
- Take care of certificate signing (Let's Encrypt)
- Built-in vulnerability scanning (OWASP-ZAP)
- Web application firewall support (ModSecurity)
- Allow authn/authz from mod_mellon upstream of app
  - Allow restricted access to internal applications to save users time on implementing themselves
  - Allow "crusty" things to be run behind appropriate security controls

https://github.com/kcl-eresearch/webfarmd

# Proxying to OpenStack website and enabling institutional auth

https://youtu.be/ja-DbF13myI

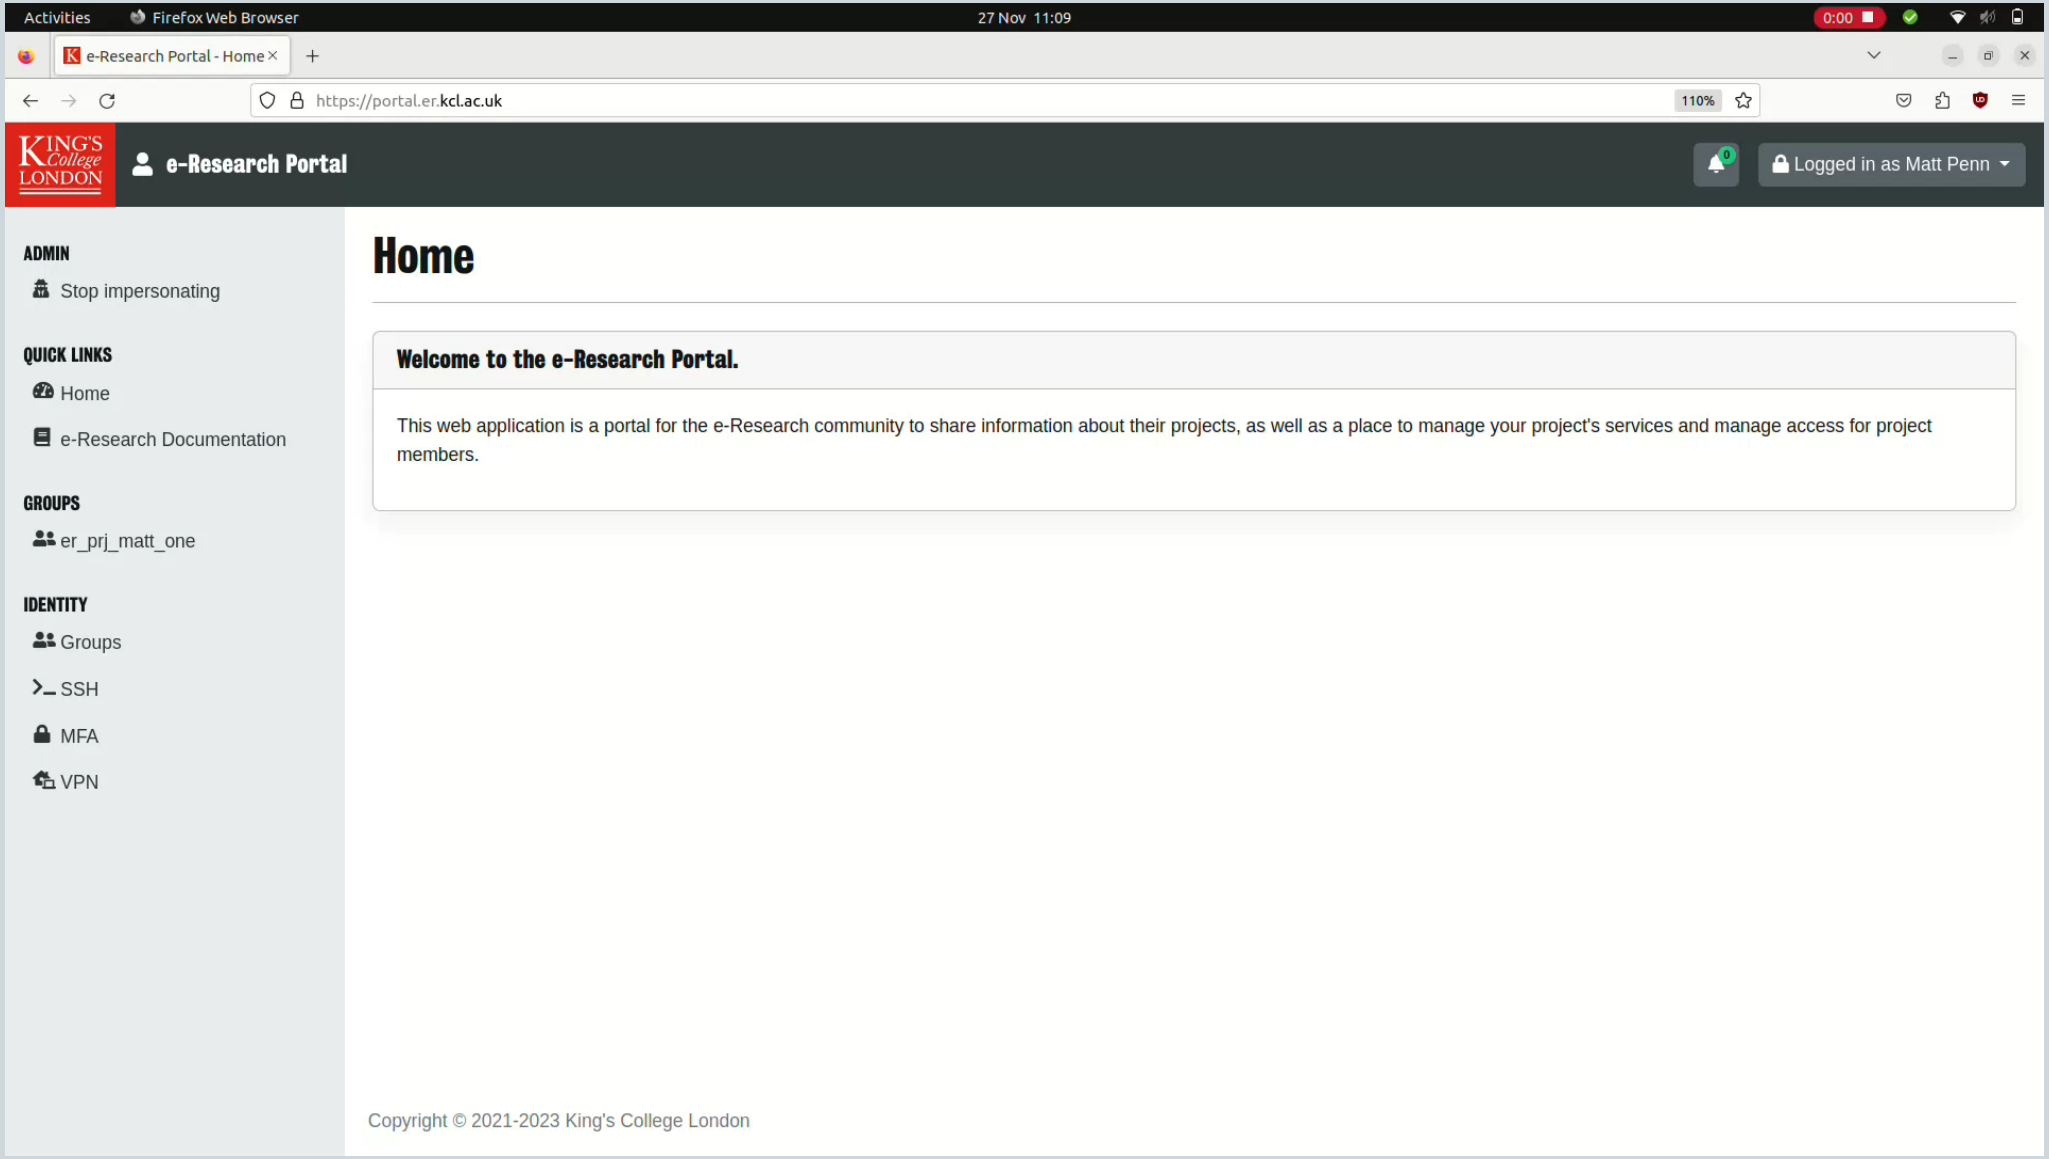

# Use case example 1 – crusty old stuff :)

- We've all been there right...
- … vendor application tied to scientific instrument
- CentOS 6
- HTTP
- Basic auth
- We can layer on top:
  - HTTPS
  - MFA
  - Group based authorisation

# Relevance alarm!!

# Use case 2 – web app for inference

- Recently deployed A30/A40/A100 hypervisors in OpenStack
- Many research groups also have GPUs in private hypervisors
- Proxy through to OpenStack instance with GPU
- Want to present trained models trained on the HPC cluster externally

Disclaimer: early case so we set this up for user (using Portal)

but expecting more use as we roll out training and documentation on this pattern

# Use case 2 – LLMs as predictive model from patient data

- https://foresight.sites.er.kcl.ac.uk/
- LLM trained on records from 1,000,000 KCH patients,
  20,000 SLaM patients
- Use inputted patient timeline to predict what will happen next
- Funded by NIHR Maudsley BRC
- Preprint https://arxiv.org/abs/2212.08072

# You need to work well with your IT department/identity team

- AD OU with delegated access
- Entra/Azure AD SAML registrations and configuration
- GraphAPI (this one also involved information compliance)

# Thinking about software sustainability

- Very custom stuff
- Bus factor is a worry
- Sharing knowledge and effort with RSE team
- Possibility of open source projectification
- Hiring junior posts (inc. industry placement)
- Keep calm, assess risk and trust your engineers

# Thanks!

- Sky and Xand for the incredible work
- The rest of the e-Research team
- King's IT
- King's researchers
- Our partners
  - StackHPC
  - Lenovo

# Questions?

**Contact details/for more information**

**Now**

**Later**

**matt.penn@kcl.ac.uk**

**HPC-SIG Slack**

**https://docs.er.kcl.ac.uk/**