The Alan Turing Institute

FRIDGE
A DARE UK PROJECT

FRIDGE: Federated Research Infrastructure by Data Governance Extension

Martin O'Reilly | Director of Research Engineering, The Alan Turing Institute

05 March 2025 | NFCS Network+ Launch event

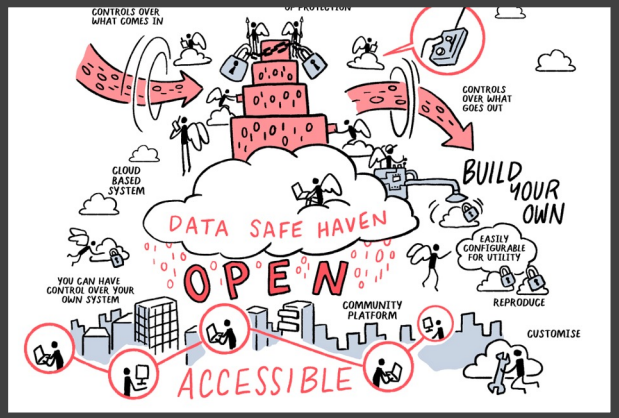

Turing Data Safe Haven

Image credits: Scriberia, created for The Turing Way, Turing, DARE UK

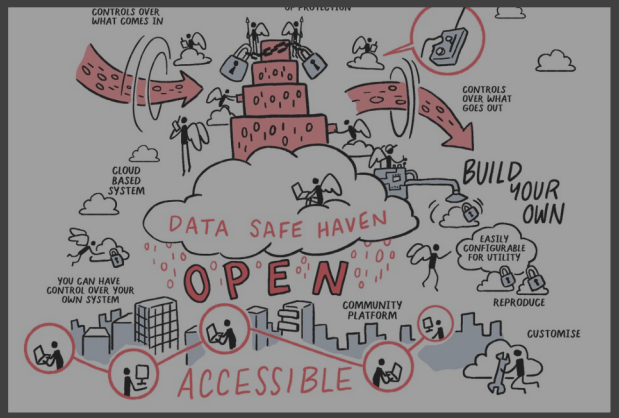

Turing Data Safe Haven

UK TRE

trusted research environment

UK TRE Community

Image credits: Scriberia, created for The Turing Way, Turing, DARE UK

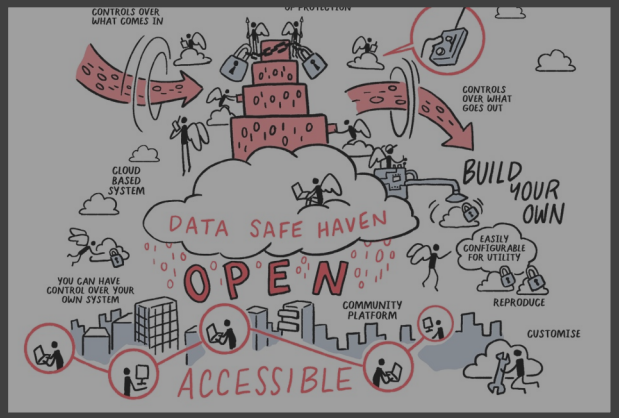

Turing Data Safe Haven

UK TRE

trusted research environment

UK TRE Community

FRIDGE

FRIDGE HPC TRE

Image credits: Scriberia, created for The Turing Way, Turing, DARE UK

What is FRIDGE?

- A Trusted Research Environment capability for High Performance Compute clusters…

What is FRIDGE?

- A Trusted Research Environment capability for High Performance Compute clusters…
- …supporting the extension of a client organisation's information governance model to the HPC cluster

What is FRIDGE?

Why TREs?

Why TREs?

What are TREs?

Why TREs?

What are TREs?

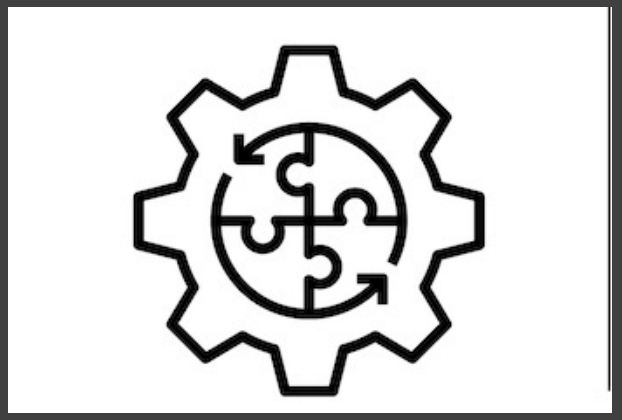

How TREs work

Why Trusted Research Environments?

Why Trusted Research Environments?

To do effective research on sensitive data

Why Trusted Research Environments?

To do effective research on sensitive data

While keeping data safe

Why Trusted Research Environments?

To do effective research on sensitive data

While keeping data safe

And maintaining public confidence

What are TREs?

What are TREs?

TREs are highly secure computational analysis environments …

What are TREs?

TREs are highly secure computational analysis environments …

… where approved researchers …

What are TREs?

TREs are highly secure computational analysis environments …

… where approved researchers …

… can carry out approved research …

What are TREs?

TREs are highly secure computational analysis environments …

… where approved researchers …

… can carry out approved research …

… using sensitive data

How TREs work

How TREs work

- **Safe projects**: research is approved by data custodians

Ritchie 2017, The 'Five Safes': a framework for planning, designing and evaluating data access solutions. http://dx.doi.org/10.5281/zenodo.897821

How TREs work

- 📎 **Safe projects**: research is approved by data custodians
- 👩 **Safe people**: researchers are trained and authorised

Ritchie 2017, The 'Five Safes': a framework for planning, designing and evaluating data access solutions. http://dx.doi.org/10.5281/zenodo.897821

How TREs work

- 📎 **Safe projects**: research is approved by data custodians
- 👩 **Safe people**: researchers are trained and authorised
- 🏢 **Safe settings**: a secure environment controls data access

Ritchie 2017, The 'Five Safes': a framework for planning, designing and evaluating data access solutions. http://dx.doi.org/10.5281/zenodo.897821

How TREs work

- 📎 **Safe projects**: research is approved by data custodians
- 👩 **Safe people**: researchers are trained and authorised
- 🏢 **Safe settings**: a secure environment controls data access
- 🔒 **Safe data:** data used is processed to protect confidentiality

Ritchie 2017, The 'Five Safes': a framework for planning, designing and evaluating data access solutions. http://dx.doi.org/10.5281/zenodo.897821

How TREs work

- **Safe projects**: research is approved by data custodians
- **Safe people**: researchers are trained and authorised
- **Safe settings**: a secure environment controls data access
- **Safe data:** data used is processed to protect confidentiality
- **Safe outputs:** outputs are processed to be non-disclosive

Ritchie 2017, The 'Five Safes': a framework for planning, designing and evaluating data access solutions. http://dx.doi.org/10.5281/zenodo.897821

How TREs work

- 📎 **Safe projects**: research is approved by data custodians
- 👩 **Safe people**: researchers are trained and authorised
- 🏢 **Safe settings**: a secure environment controls data access ← TREs?
- 🔒 **Safe data:** data used is processed to protect confidentiality
- 📊 **Safe outputs:** outputs are processed to be non-disclosive

How TREs work

- **Safe projects**: research is approved by data custodians
- **Safe people**: researchers are trained and authorised
- **Safe settings**: a secure environment controls data access
- **Safe data:** data used is processed to protect confidentiality
- **Safe outputs:** outputs are processed to be non-disclosive

← TREs

How TREs work

📎 **Safe projects**: research is approved by data custodians	Governance
👩 **Safe people**: researchers are trained and authorised	
🏢 **Safe settings**: a secure environment controls data access	Technical
🔒 **Safe data:** data used is processed to protect confidentiality	Statistical
📊 **Safe outputs:** outputs are processed to be non-disclosive	

Why FRIDGE?

Why FRIDGE?

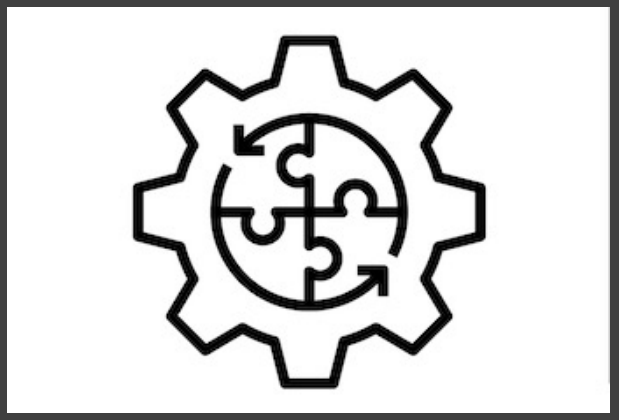

How FRIDGE works

Why FRIDGE?

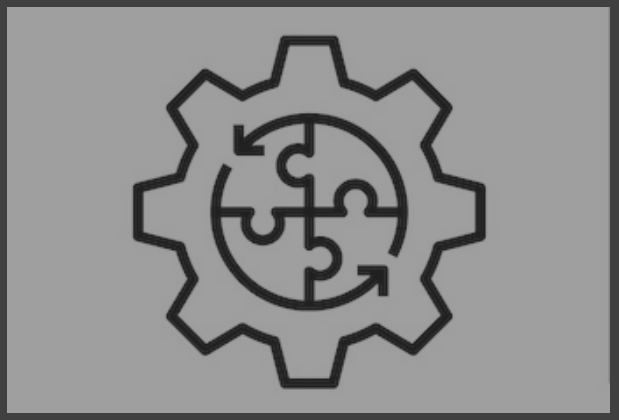

How FRIDGE works

How to get involved

Why is FRIDGE needed?

Why is FRIDGE needed?

- HPC clusters generally don’t support TRE level isolation between projects and their data

Why is FRIDGE needed?

- HPC clusters generally don’t support TRE level isolation between projects and their data
- TREs generally don’t include large scale HPC capabilities

Why is FRIDGE needed?

- HPC clusters generally don’t support TRE level isolation between projects and their data
- TREs generally don’t include large scale HPC capabilities
- Very difficult to establish trust between two systems, even when both are TREs

How FRIDGE works

How FRIDGE works

- Put information governance problem front and centre

How FRIDGE works

- Put information governance problem front and centre
- Extend client organisation's information governance model to HPC cluster

How FRIDGE works

- Put information governance problem front and centre
- Extend client organisation's information governance model to HPC cluster
- Replicate shared responsibility model used for TREs deployed in public cloud

How FRIDGE works

- API for client organisation to configure a strongly isolated HPC system segment as a TRE tenancy

How FRIDGE works

- API for client organisation to configure a strongly isolated HPC system segment as a TRE tenancy
- TRE tenancy supports a common infrastructure layer (Kubernetes-compatible container service)

How FRIDGE works

- API for client organisation to configure a strongly isolated HPC system segment as a TRE tenancy
- TRE tenancy supports a common infrastructure layer (Kubernetes-compatible container service)
- Code to reproducibly deploy TRE components

How FRIDGE works

- API for client organisation to configure a strongly isolated HPC system segment as a TRE tenancy
- TRE tenancy supports a common infrastructure layer (Kubernetes-compatible container service)
- Code to reproducibly deploy TRE components
- Standardised shared responsibility TRE tenancy hosting agreement, leveraging existing standards

How FRIDGE works

- Deploy a standalone TRE into the HPC TRE tenancy

How FRIDGE works

- Deploy a standalone TRE into the HPC TRE tenancy, or
- Extend an existing TRE with a satellite segment deployed into the HPC TRE tenancy

How to get involved

- **Data providers / TRE administrators**: What would make you comfortable extending your existing information governance to a HPC TRE tenancy?

How to get involved

- **Data providers / TRE administrators**: What would make you comfortable extending your existing information governance to a HPC TRE tenancy?
- **HPC administrators:** Can you support the TRE tenancy model and how would you do so?

Learn more

FRIDGE: moreilly@turing.ac.uk (website coming next week)

TRE community: https://www.uktre.org/

Turing Data Safe Haven:
https://www.turing.ac.uk/research/research-projects/data-safe-havens-cloud

moreilly@turing.ac.uk | @martinoreilly.bsky.social | @martinoreilly@mastodon.social

Discussion questions

- **Data providers / TRE administrators**: What would make you comfortable extending your existing information governance to a HPC TRE tenancy?
- **HPC administrators:** Can you support the TRE tenancy model and how would you do so?

moreilly@turing.ac.uk | @martinoreilly.bsky.social | @martinoreilly@mastodon.social