The Alan Turing Institute

FRIDGE
A DARE UK PROJECT

FRIDGE: Trusted research environments supercomputers

UCL | University of Bristol | University of Cambridge | King's College London | UK Biobank

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

NCFS Spring conference | 26 February 2026

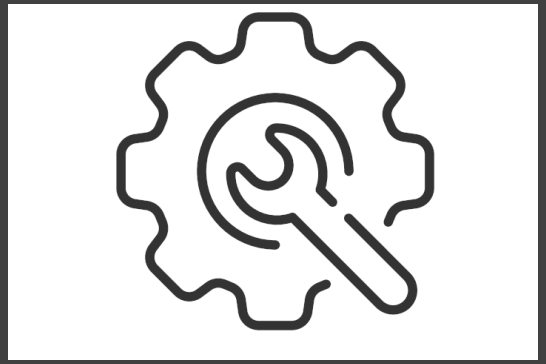

What FRIDGE is

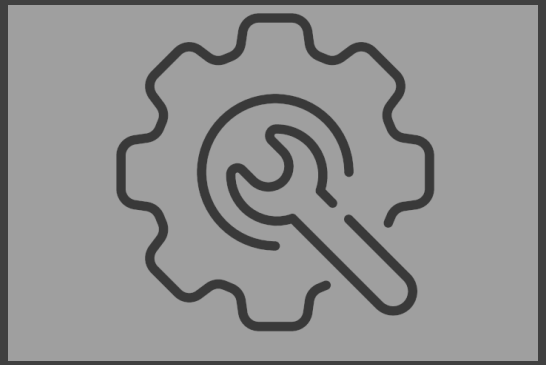

What FRIDGE is

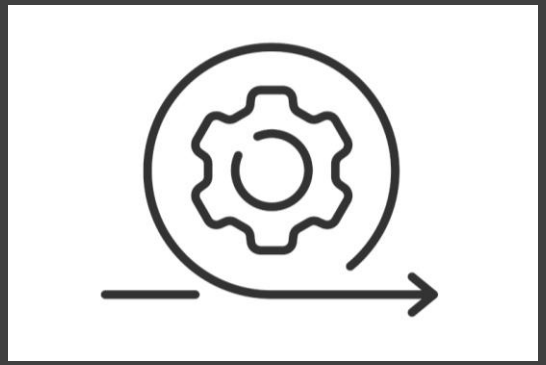

FRIDGE in practice

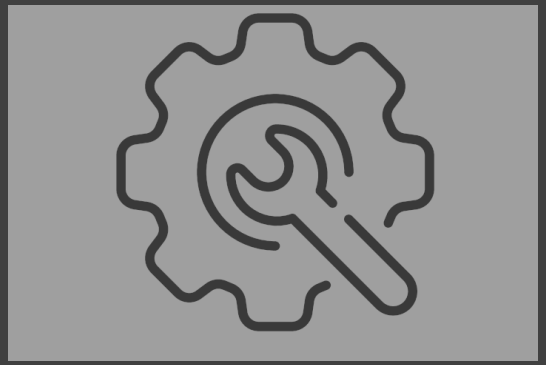

What FRIDGE is

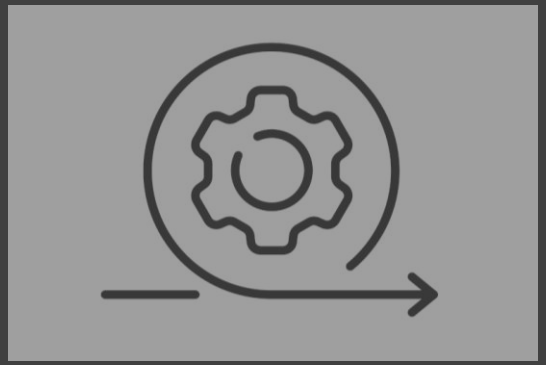

FRIDGE in practice

Beyond FRIDGE

What is FRIDGE?

- A Trusted Research Environment capability for supercomputers …

What is FRIDGE?

- A Trusted Research Environment capability for supercomputers …
- …supporting the extension of an existing TRE's information governance model to the supercomputer

Why is FRIDGE needed?

Why is FRIDGE needed?

- TREs generally aren’t supercomputers

Why is FRIDGE needed?

- TREs generally aren’t supercomputers
- Supercomputers generally aren’t TREs

Why is FRIDGE needed?

- TREs generally aren’t supercomputers
- Supercomputers generally aren’t TREs
- Establishing trust between systems is hard

Why is FRIDGE needed?

- TREs generally aren't supercomputers
- Supercomputers generally aren't TREs
- Establishing trust between systems is hard

  … even when both systems are TREs

How FRIDGE works

How FRIDGE works

Two options:

How FRIDGE works

Two options:

- Deploy a standalone TRE into the supercomputer

How FRIDGE works

Two options:

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

How FRIDGE works

Two options:

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

How FRIDGE works

- Puts information governance problem front and centre

How FRIDGE works

- Puts information governance problem front and centre
- Extends an existing TRE's information governance model to the supercomputer

How FRIDGE works

- Puts information governance problem front and centre
- Extends an existing TRE's information governance model to the supercomputer
- Adapts shared responsibility model used for TREs deployed in public cloud

How FRIDGE works

FRIDGE lets groups of computers within a supercomputer be strongly isolated from each other in a way that complies with the governance rules for different existing TREs

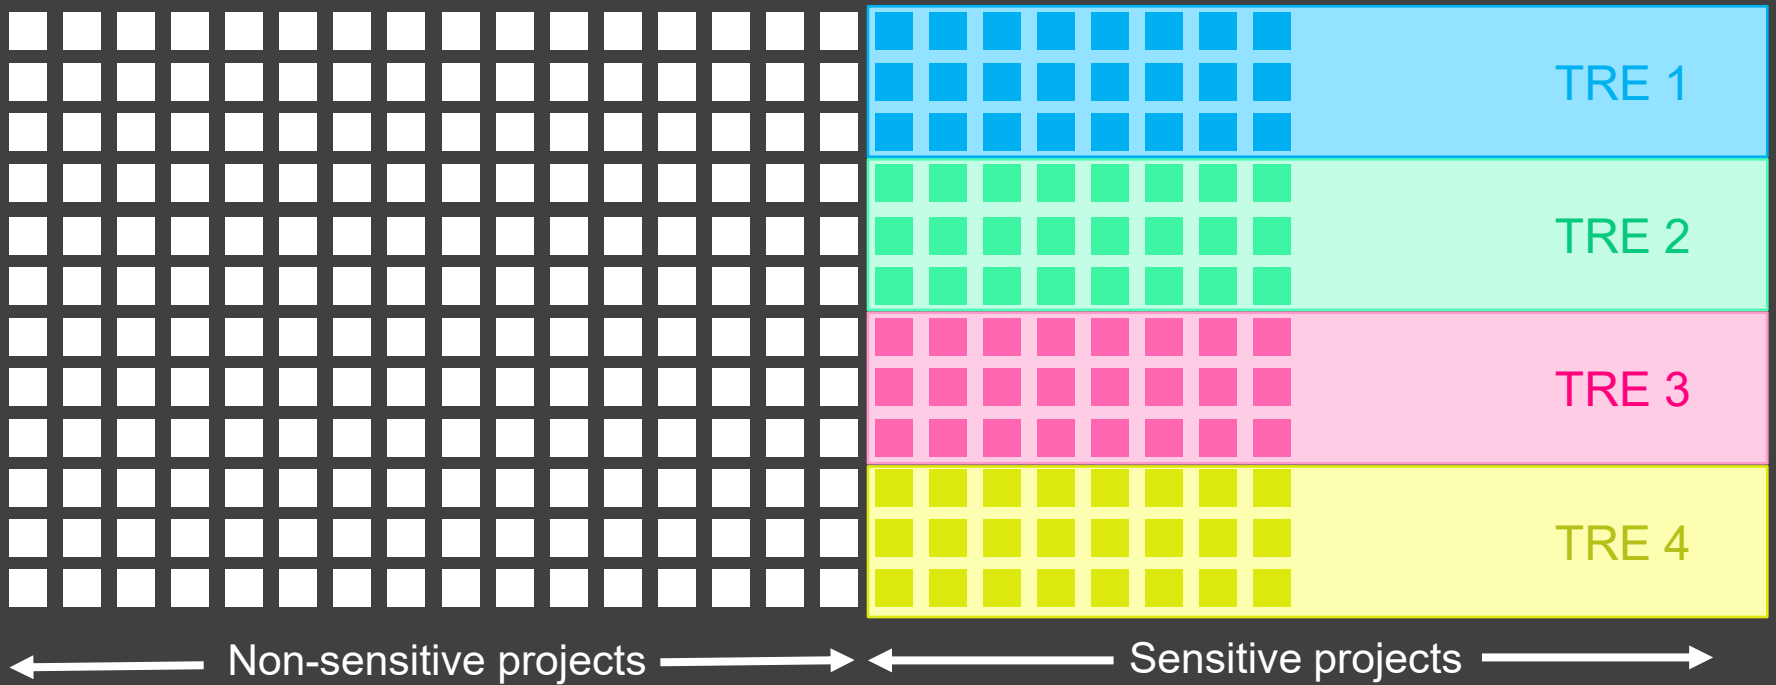

How FRIDGE works

- Provides technical protections to strongly isolate groups of computers within a supercomputer

How FRIDGE works

- Provides technical protections to strongly isolate groups of computers within a supercomputer
- Allows these isolated groups of computers to be controlled and configured by existing TRE owners

How FRIDGE works

- Provides technical protections to strongly isolate groups of computers within a supercomputer
- Allows these isolated groups of computers to be controlled and configured by existing TRE owners
- Tells the supercomputer owner what they need to do to ensure security

How FRIDGE works

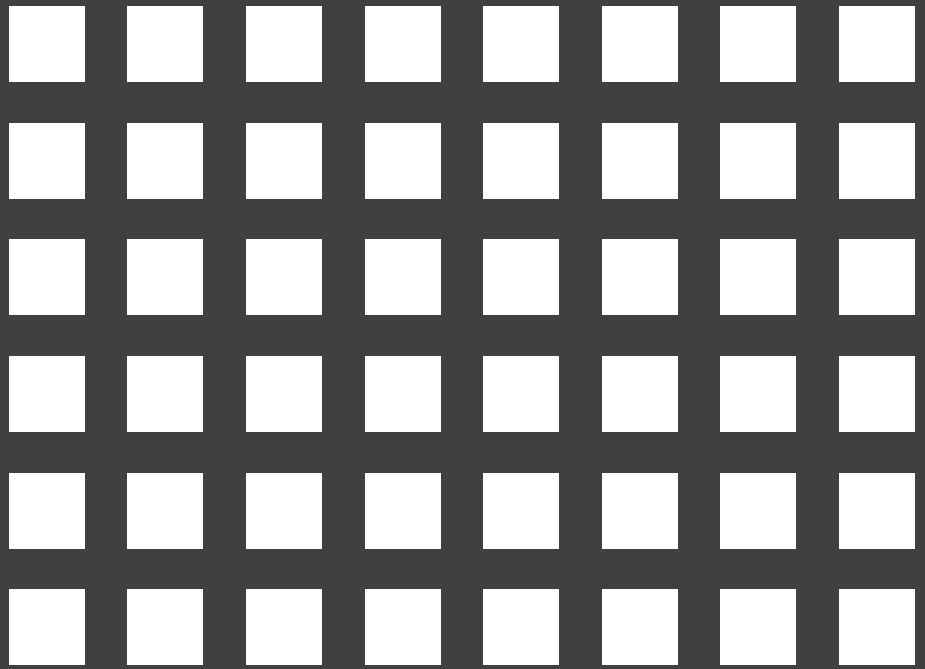

Supercomputer (Cambridge / Bristol)

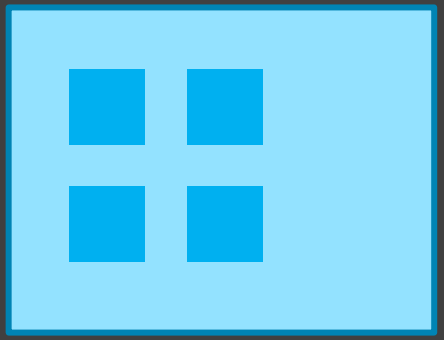

TRE (Turing / KCL)

How FRIDGE works

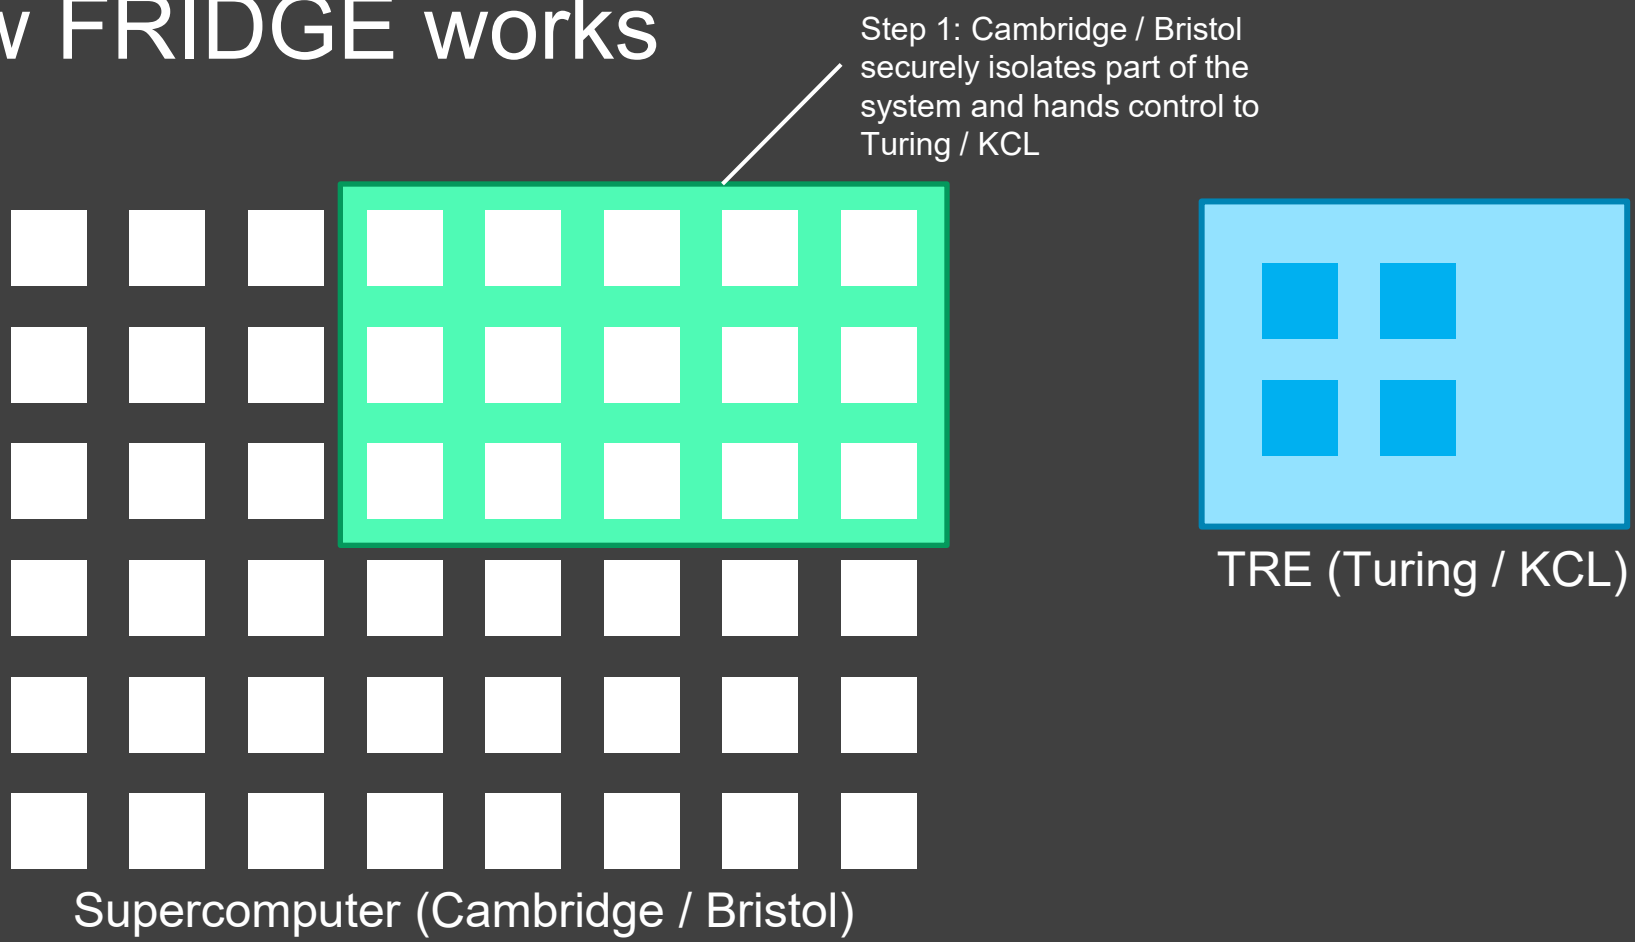

How FRIDGE works

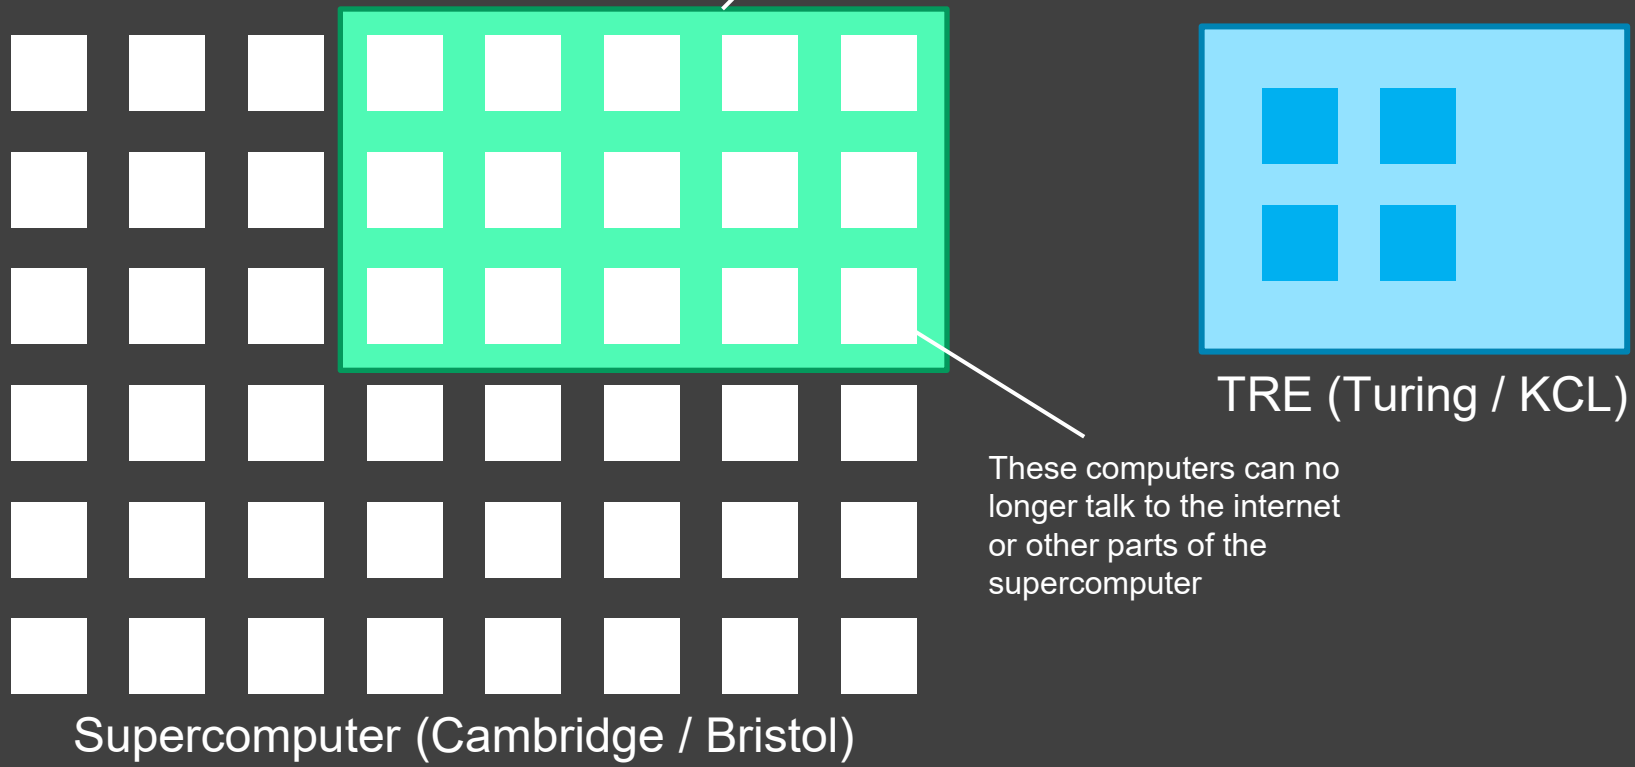

How FRIDGE works

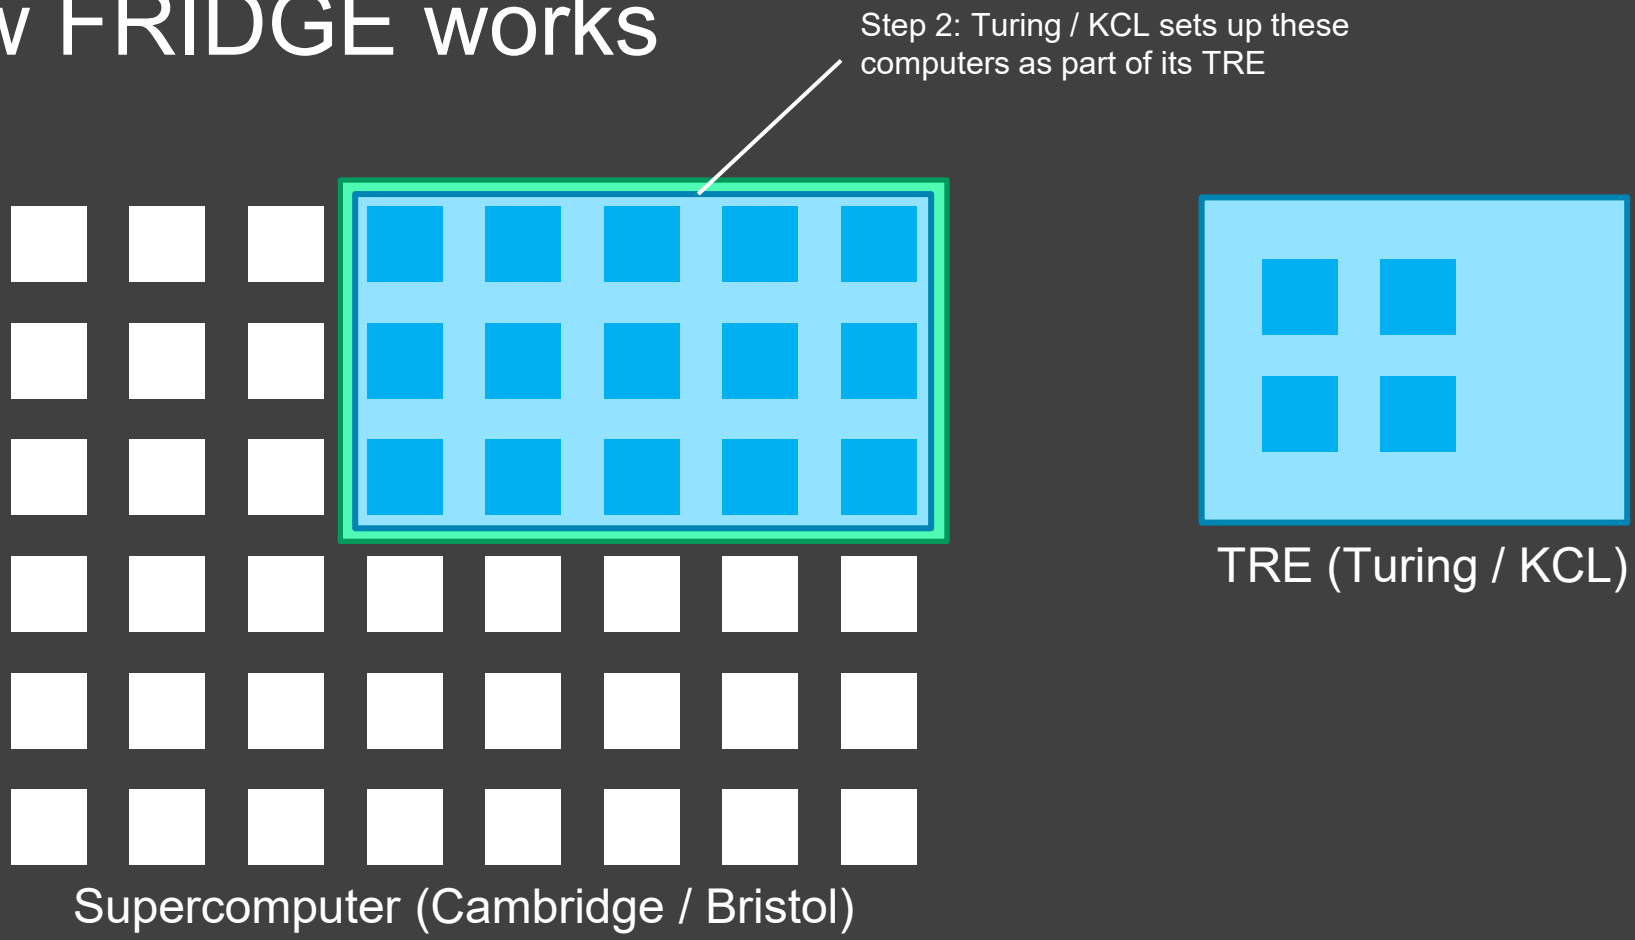

How FRIDGE works

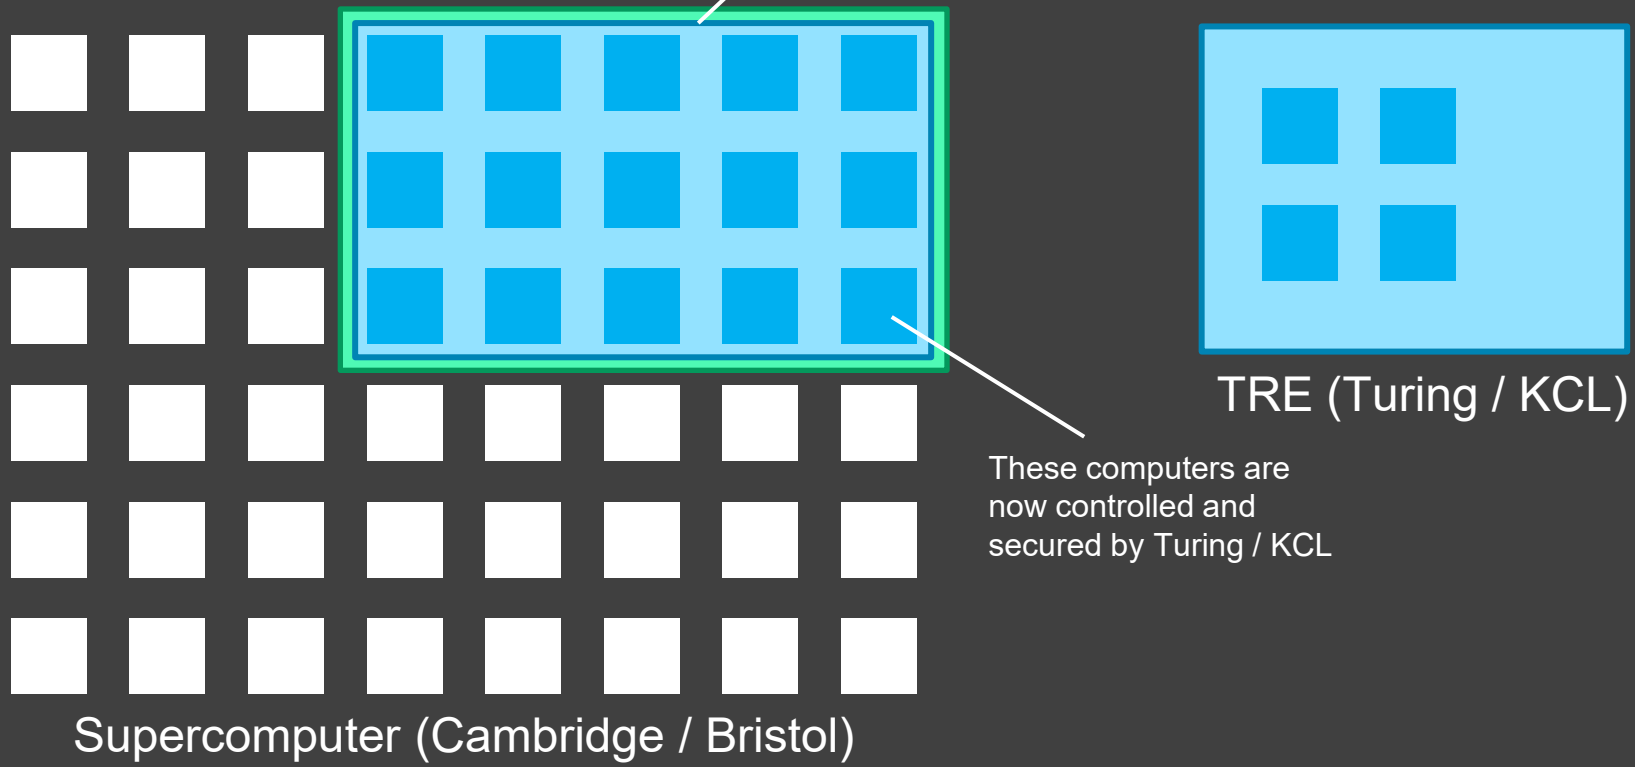

How FRIDGE works

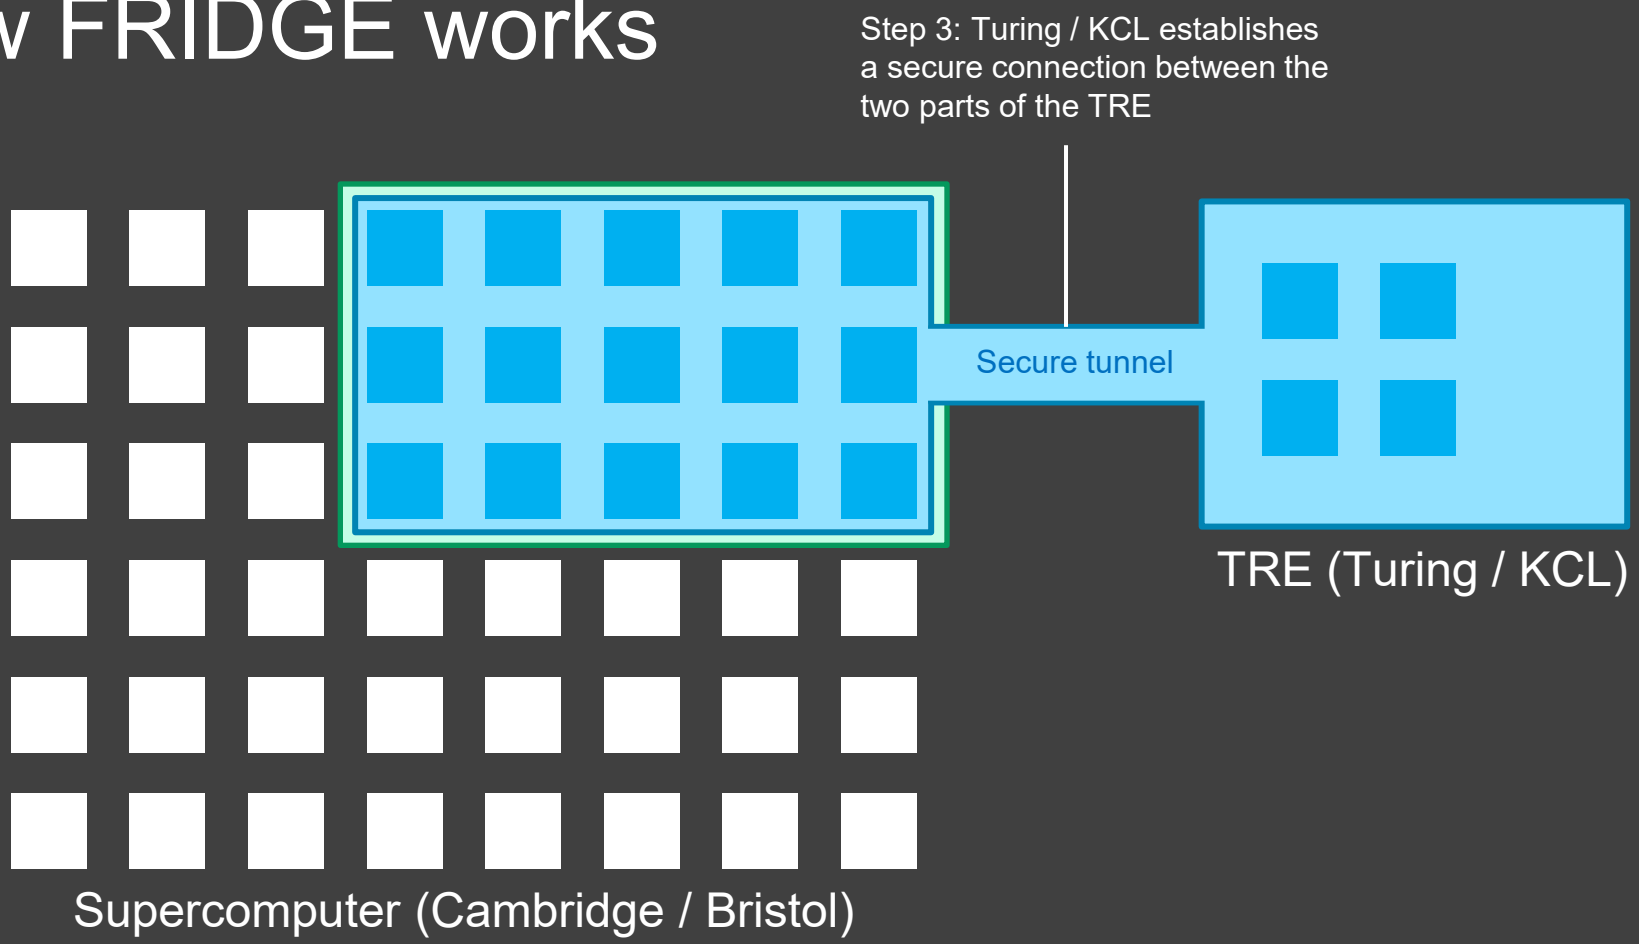

How FRIDGE works

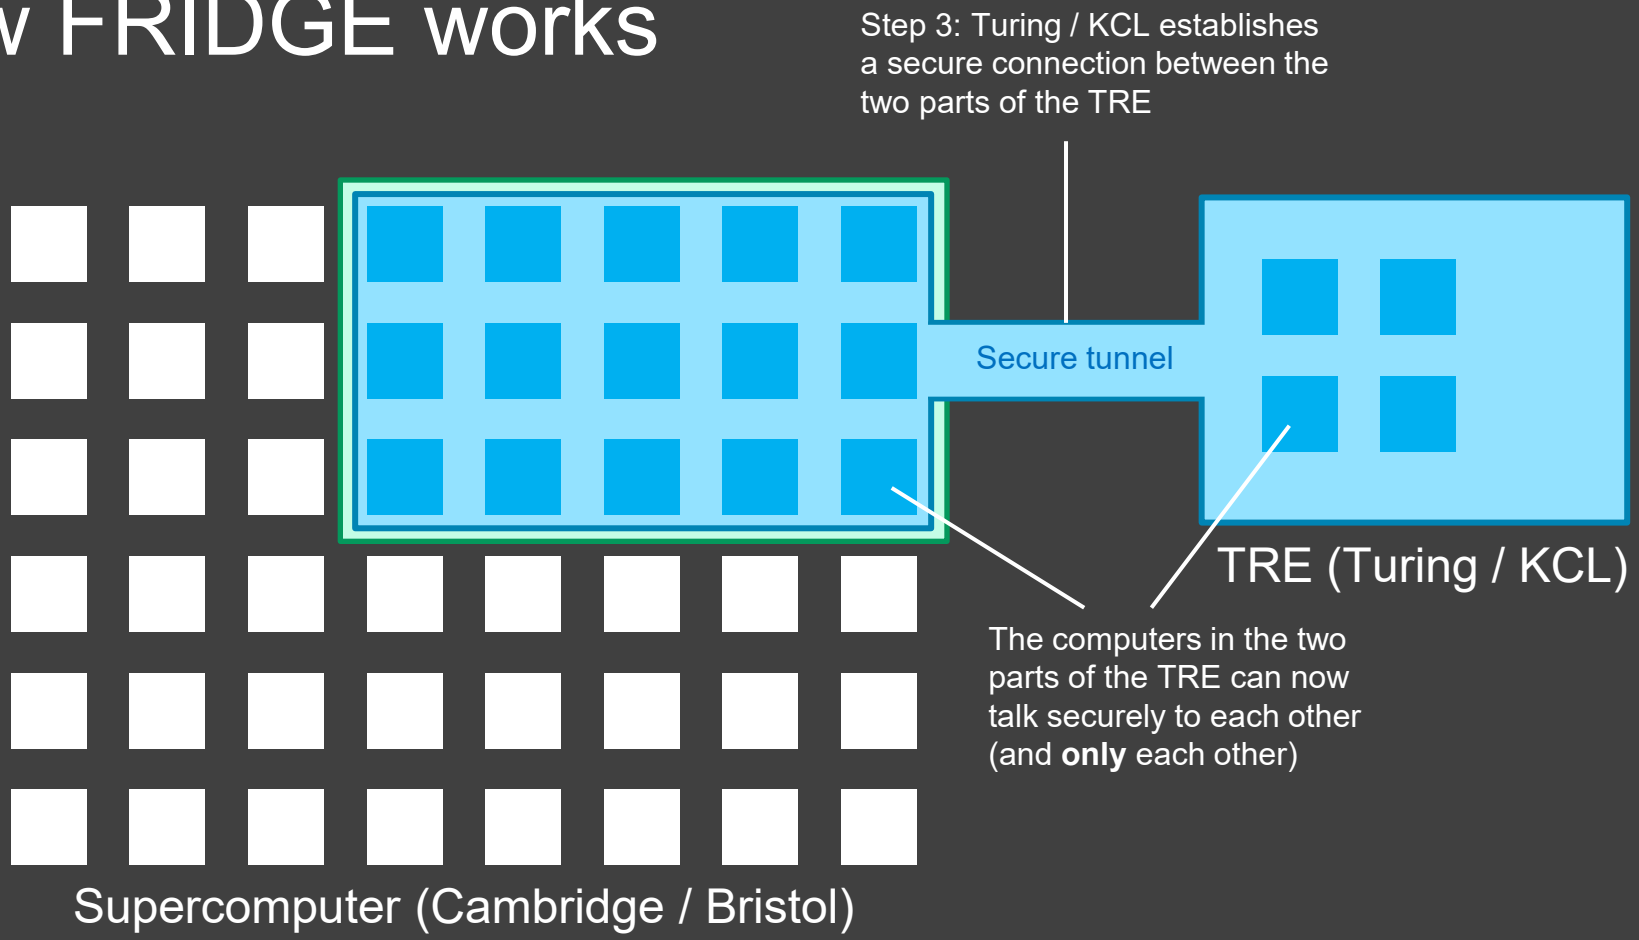

How FRIDGE works

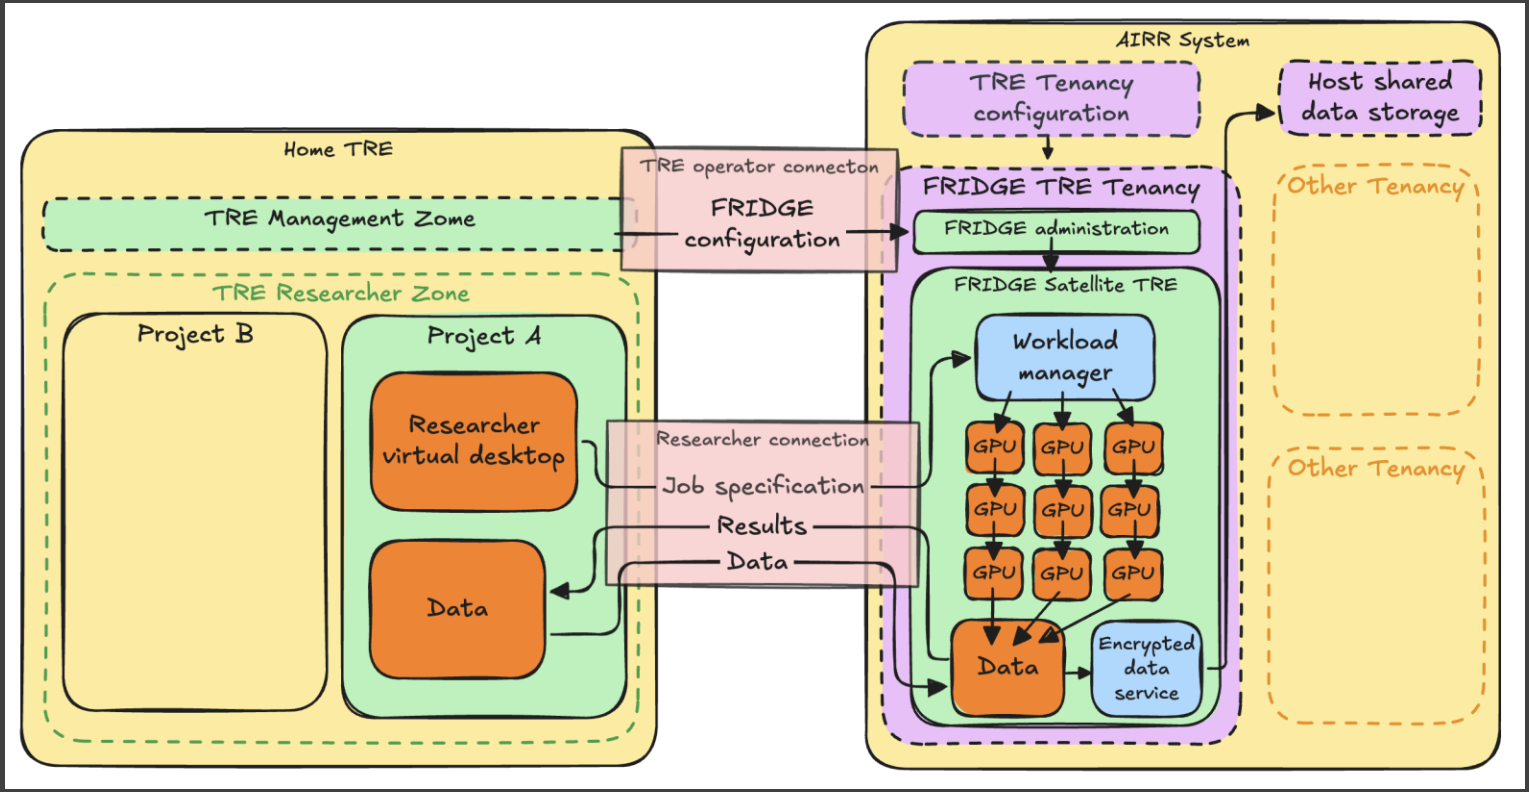

Putting FRIDGE into practice

What are we doing?

- **Goal:** Enabling the safe use of large-scale health datasets on the UK's national AI supercomputers

What are we doing?

- **Goal:** Enabling the safe use of large-scale health datasets on the UK's national AI supercomputers
- **How:** Showing this can be done for three large-scale high-impact datasets using existing real-world research projects

What are we doing?

- **Goal:** Enabling the safe use of large-scale health datasets on the UK's national AI supercomputers
- **How:** Showing this can be done for three large-scale high-impact datasets using existing real-world research projects
- **How:** Showing how this will scale to all users of these datasets and to other sensitive datasets

Datasets

Datasets

- **UK Biobank**: Large 500,000 subject longitudinal research cohort. Long follow-up includes many different health outcomes. Deep information on lifestyle, environmental, biological factors.

Datasets

- **UK Biobank**: Large 500,000 subject longitudinal research cohort. Long follow-up includes many different health outcomes. Deep information on lifestyle, environmental, biological factors.
- **PharosAI**: Large-scale multimodal data (digital pathology, imaging, omics DNA/RNA, longitudinal clinical data). 12,000 patients per year, 100,000 patients linked sample data, 28 million datapoints.

Datasets

- **UK Biobank**: Large 500,000 subject longitudinal research cohort. Long follow-up includes many different health outcomes. Deep information on lifestyle, environmental, biological factors.
- **PharosAI**: Large-scale multimodal data (digital pathology, imaging, omics DNA/RNA, longitudinal clinical data). 12,000 patients per year, 100,000 patients linked sample data, 28 million datapoints.
- **AI Centre for Value Based Healthcare**: 170,000 brain MRI images from 40,000 patients with metadata and clinical radiology reports.

Datasets

- **UK Biobank**: Large 500,000 subject longitudinal research cohort. Long follow-up includes many different health outcomes. Deep information on lifestyle, environmental, biological factors.
- **PharosAI**: Large-scale multimodal data (digital pathology, imaging, omics DNA/RNA, longitudinal clinical data). 12,000 patients per year, 100,000 patients linked sample data, 28 million datapoints.
- **AI Centre for Value Based Healthcare**: 170,000 brain MRI images from 40,000 patients with metadata and clinical radiology reports
- **Commonality**: Large-scale multi-modal health datasets with large imaging component, with lots of projects using AI.

Key challenges

Key challenges

- Securely hosting and updating a single shared copy of each dataset

Key challenges

- Securely hosting and updating a single shared copy of each dataset
- Enforcing data provider’s access controls

Key challenges

- Securely hosting and updating a single shared copy of each dataset
- Enforcing data provider’s access controls
- Integrating into DSIT’s allocation process for AIRR

Key challenges

- Securely hosting and updating a single shared copy of each dataset
- Enforcing data provider’s access controls
- Integrating into DSIT’s allocation process for AIRR
- Providing an easy pathway for other TREs to connect to AIRR

Key challenges

- Securely hosting and updating a single shared copy of each dataset
- Enforcing data provider’s access controls
- Integrating into DSIT’s allocation process for AIRR
- Providing an easy pathway for other TREs to connect to AIRR
- Providing an option for researchers with no “home” TRE

Key challenges

- Securely hosting and updating a single shared copy of each dataset
- Enforcing data provider’s access controls
- Integrating into DSIT’s allocation process for AIRR
- Providing an easy pathway for other TREs to connect to AIRR
- Providing an option for researchers with no “home” TRE
- Identifying scaling path to 100s of projects and 1,000s of researchers

Key challenges

- Securely hosting and updating a single shared copy of each dataset
- Enforcing data provider’s access controls
- Integrating into DSIT’s allocation process for AIRR
- Providing an easy pathway for other TREs to connect to AIRR
- Providing an option for researchers with no “home” TRE
- Identifying scaling path to 100s of projects and 1,000s of researchers
- Supporting adoption by other sets of TREs and supercomputers

Beyond FRIDGE

The limits of FRIDGE

The limits of FRIDGE

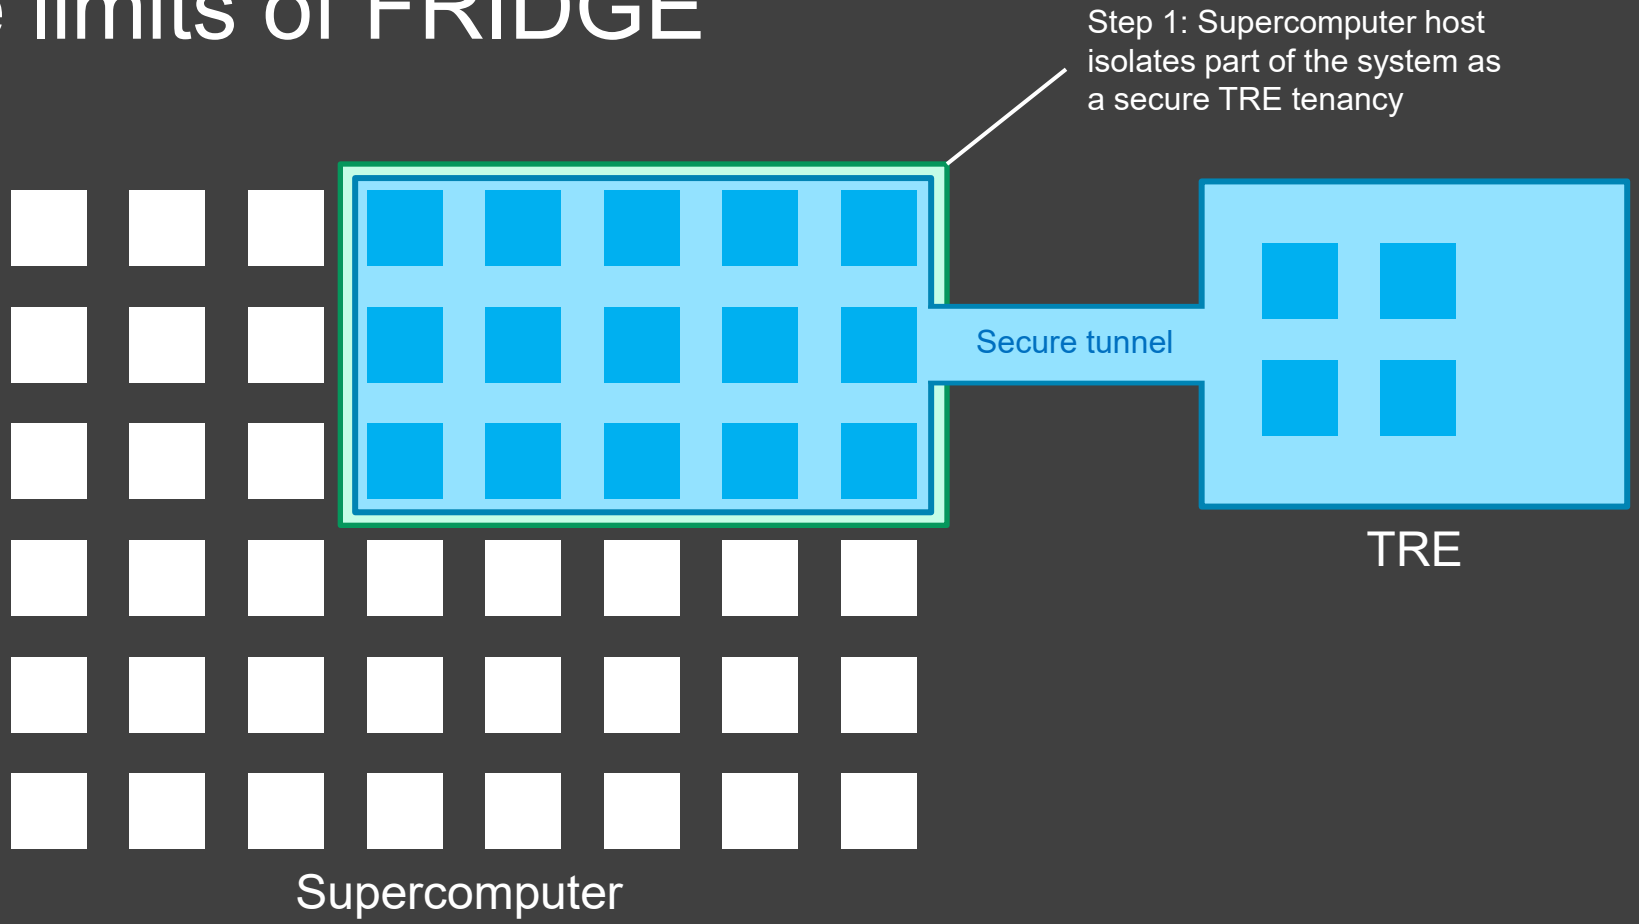

The limits of FRIDGE

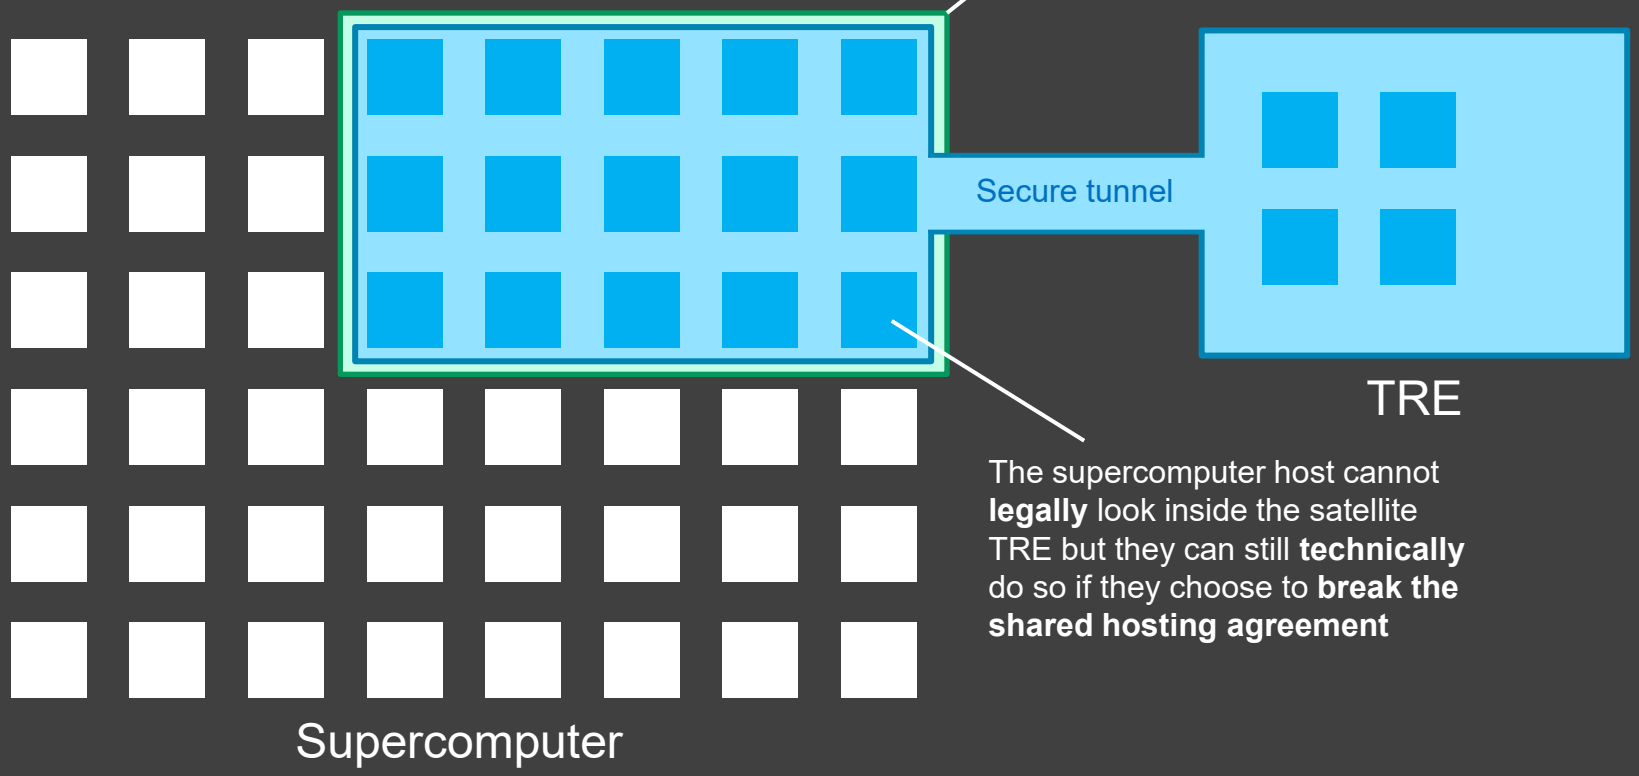

The limits of FRIDGE

The limits of FRIDGE

- Relies on a shared responsibility model

The limits of FRIDGE

- Relies on a shared responsibility model
- May not be sufficient for highly sensitive data

The limits of FRIDGE

- Relies on a shared responsibility model
- May not be sufficient for highly sensitive data
- Doesn’t support the “long tail” of small-scale AI workstations run by individual labs or researchers

The limits of FRIDGE

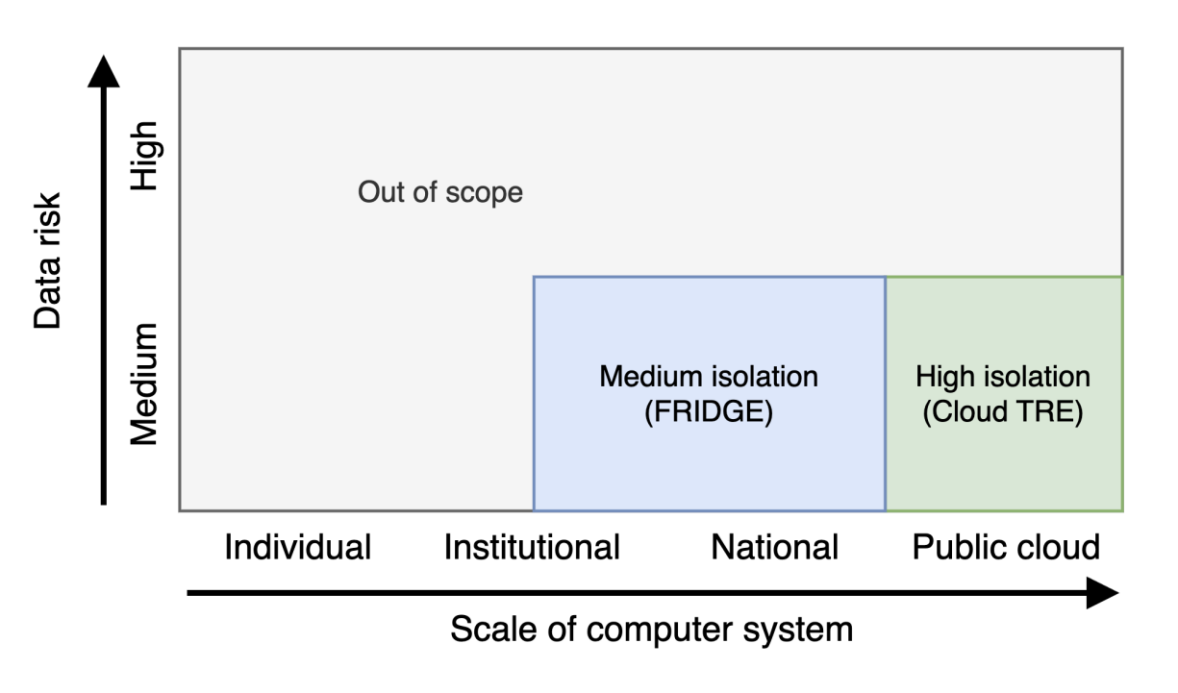

Beyond FRIDGE (TRUSTEE)

Beyond FRIDGE (TRUSTEE)

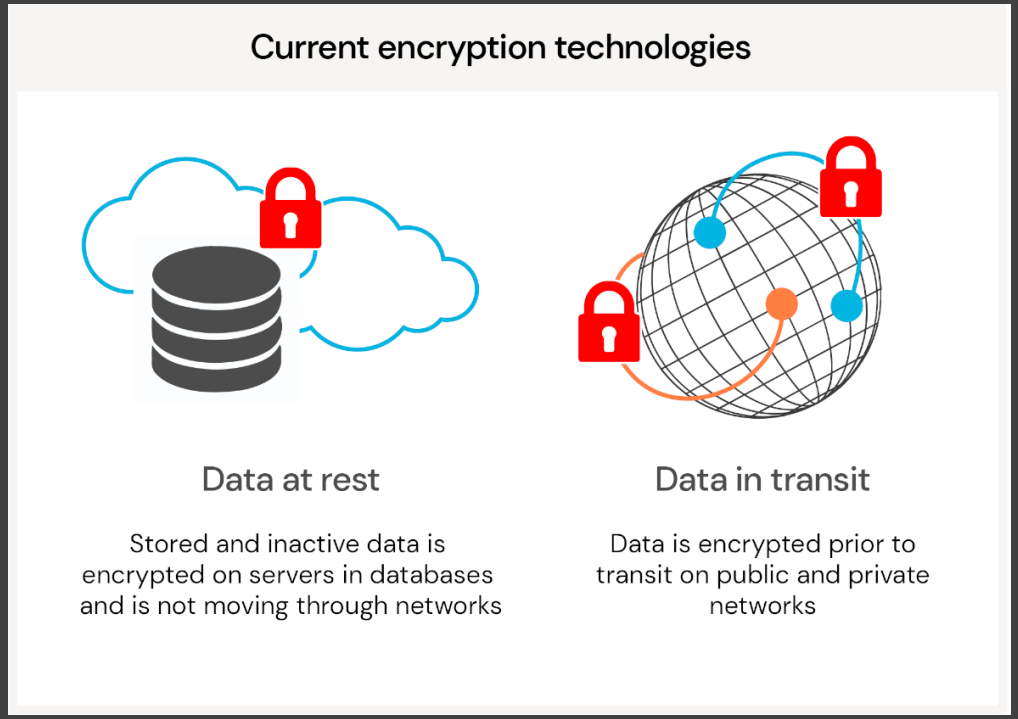

Beyond FRIDGE (TRUSTEE)

Current encryption technologies

Data at rest

Stored and inactive data is encrypted on servers in databases and is not moving through networks

Data in transit

Data is encrypted prior to transit on public and private networks

Confidential computing

Protect data in use

Data is protected in use by RAM encryption and hardware-based technologies that protect data during computation

Beyond FRIDGE (TRUSTEE)

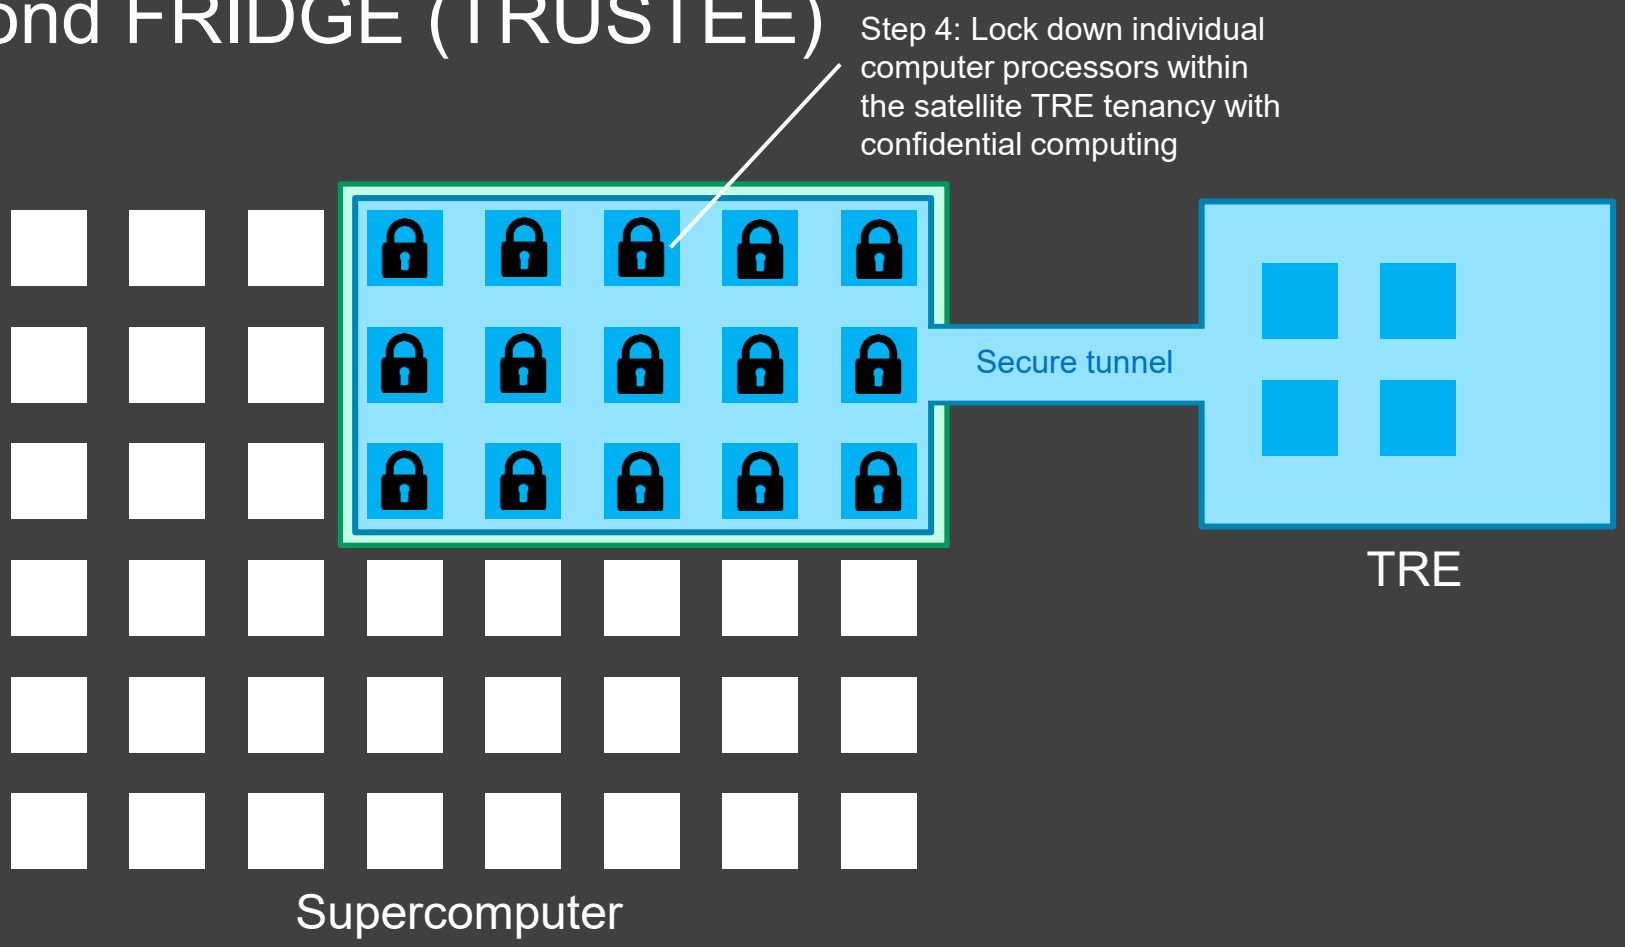

Beyond FRIDGE (TRUSTEE)

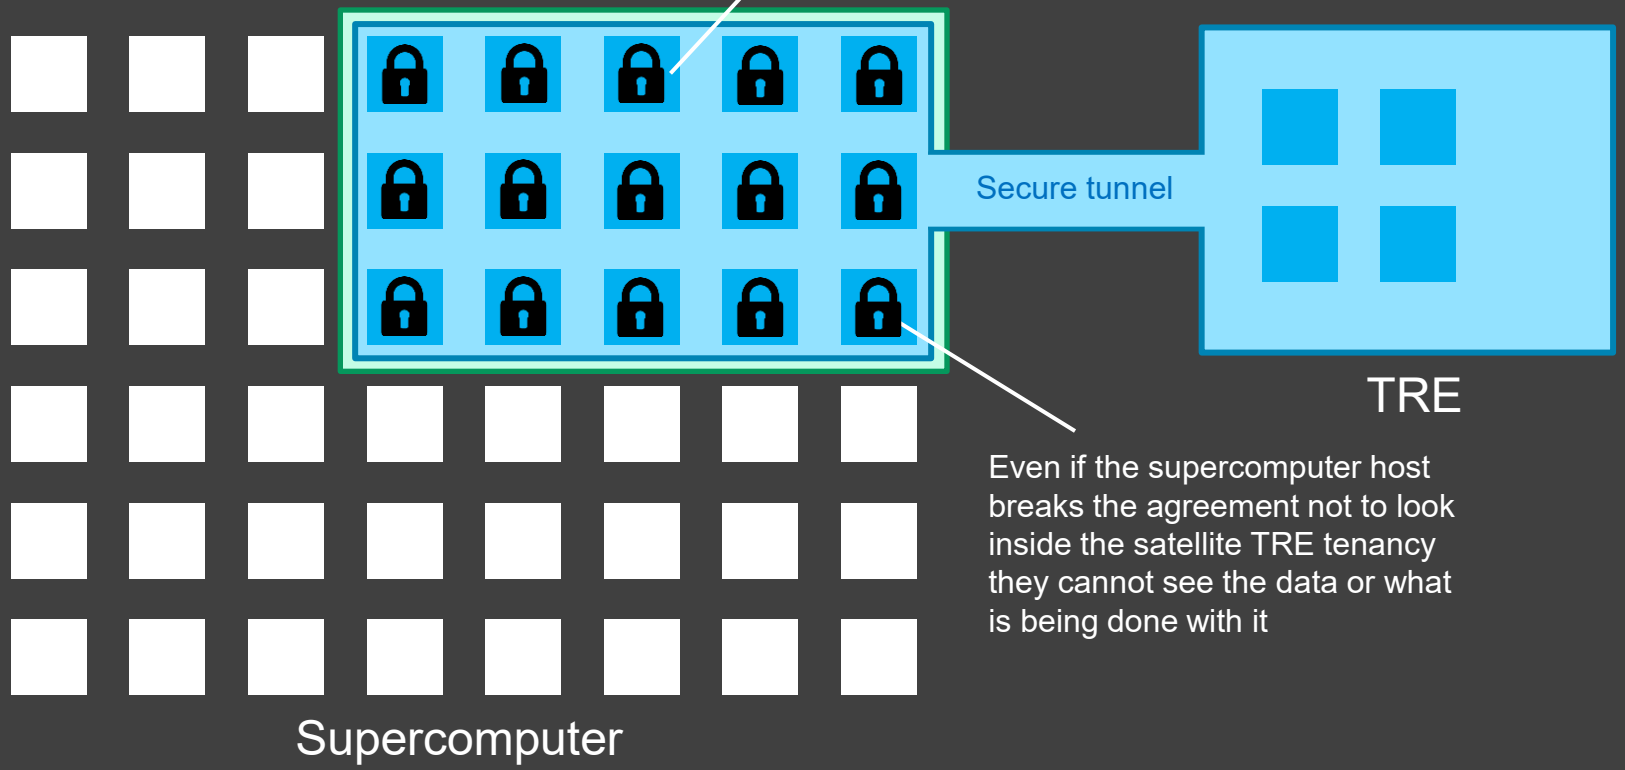

Beyond FRIDGE (TRUSTEE)

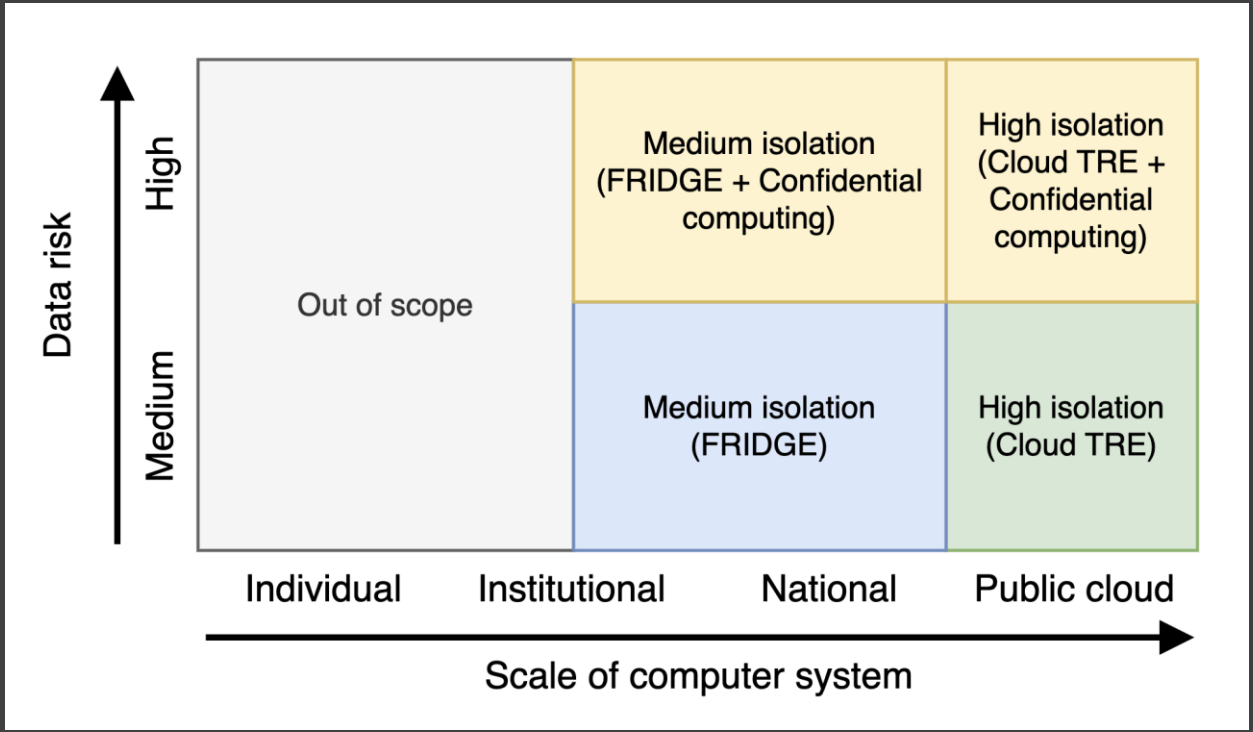

Beyond FRIDGE (TRUSTEE)

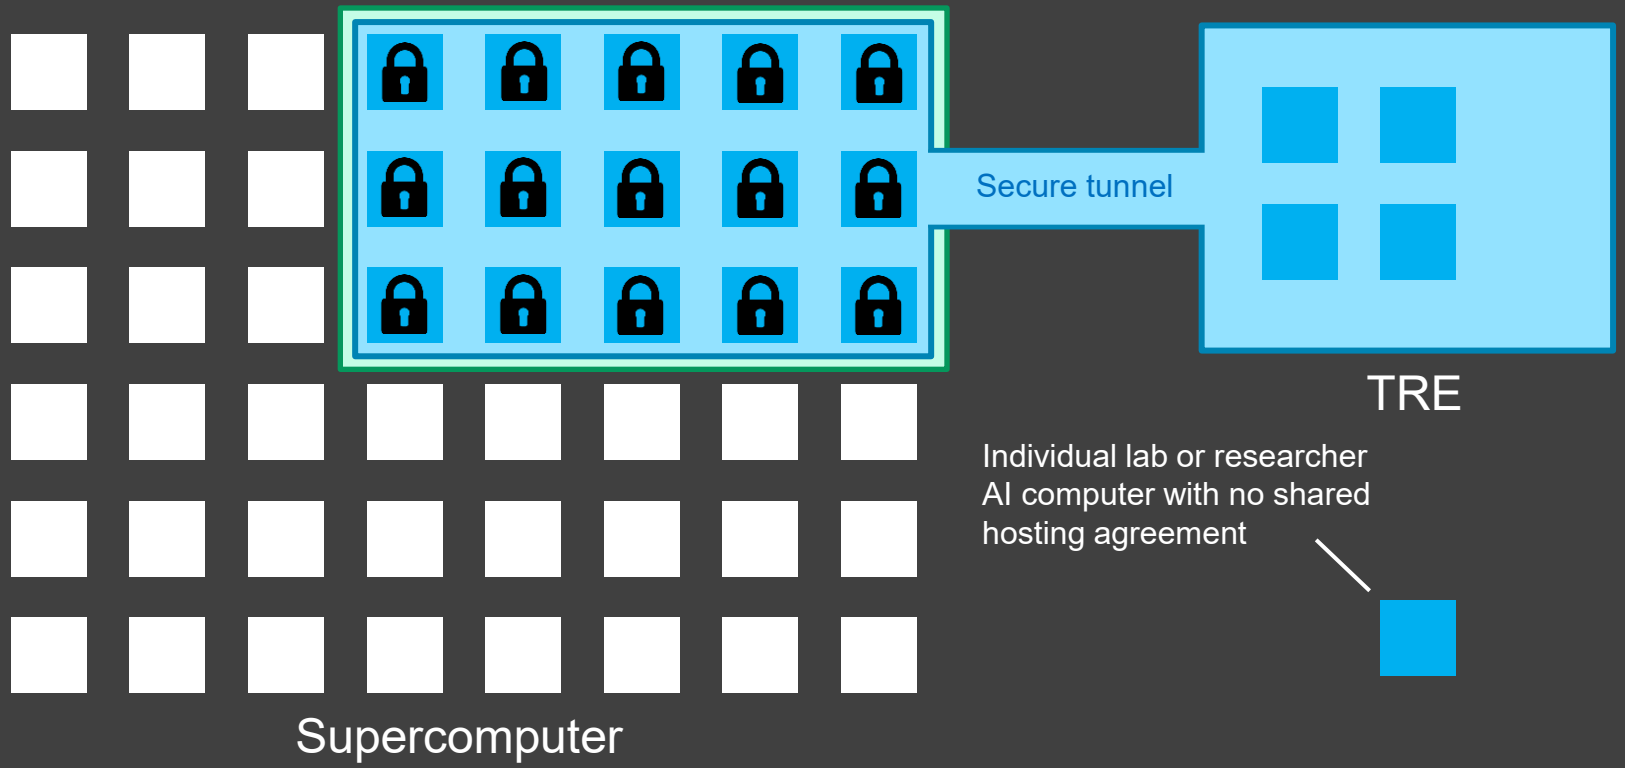

Beyond FRIDGE (TRUSTEE)

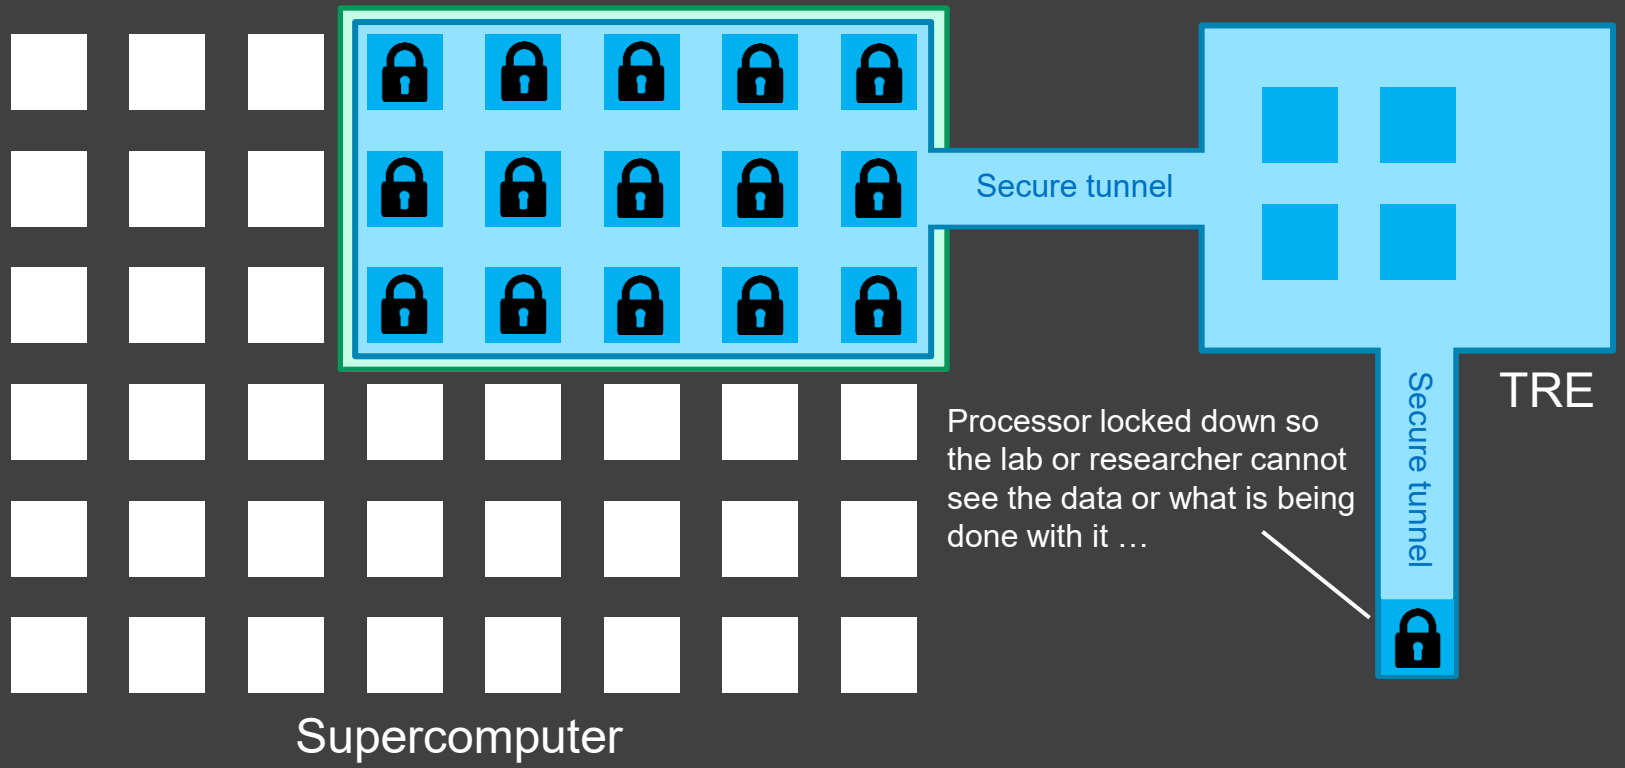

Beyond FRIDGE (TRUSTEE)

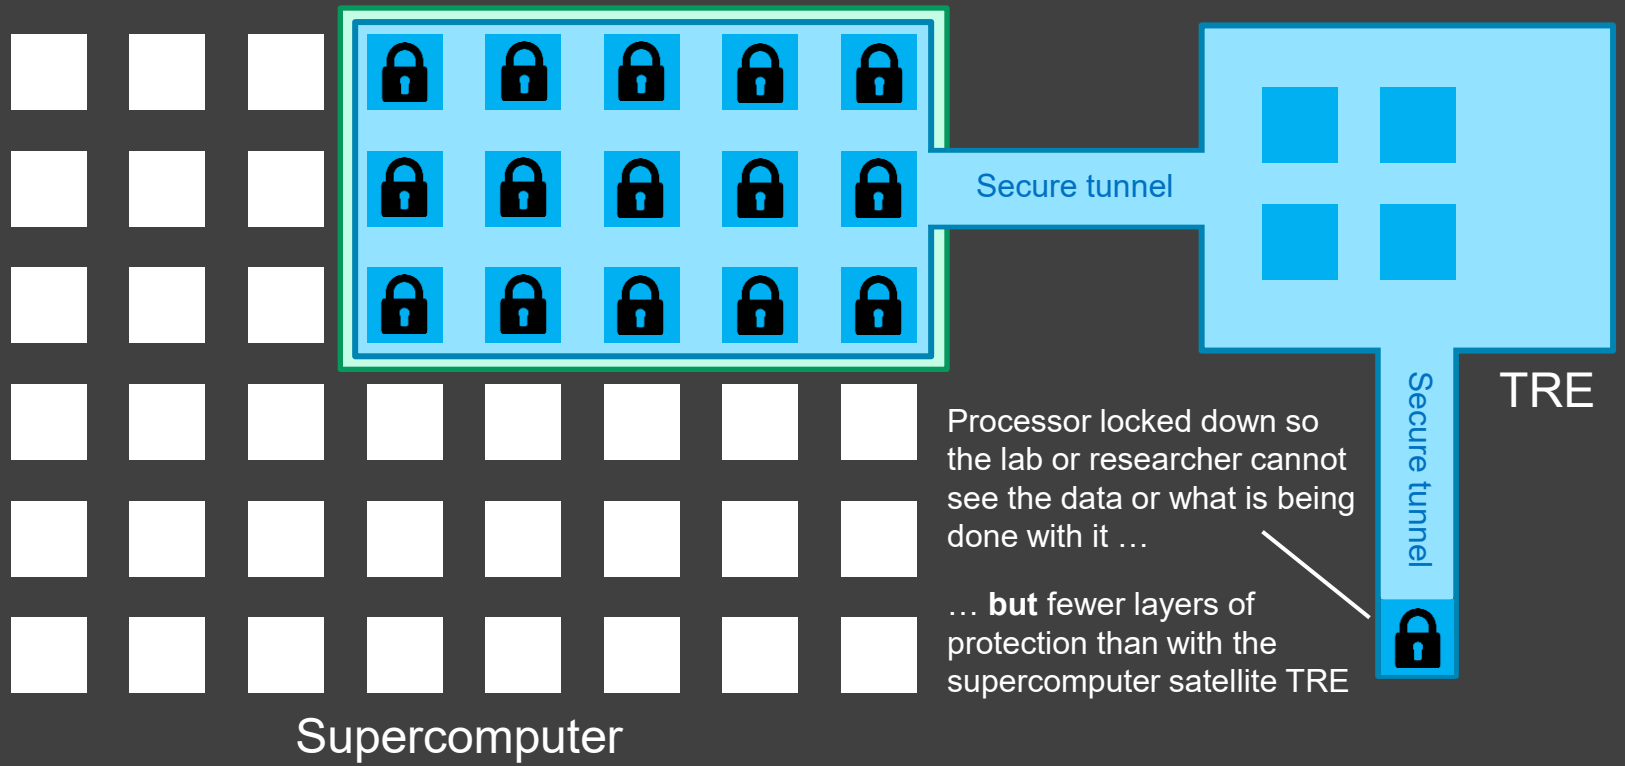

Beyond FRIDGE (TRUSTEE)

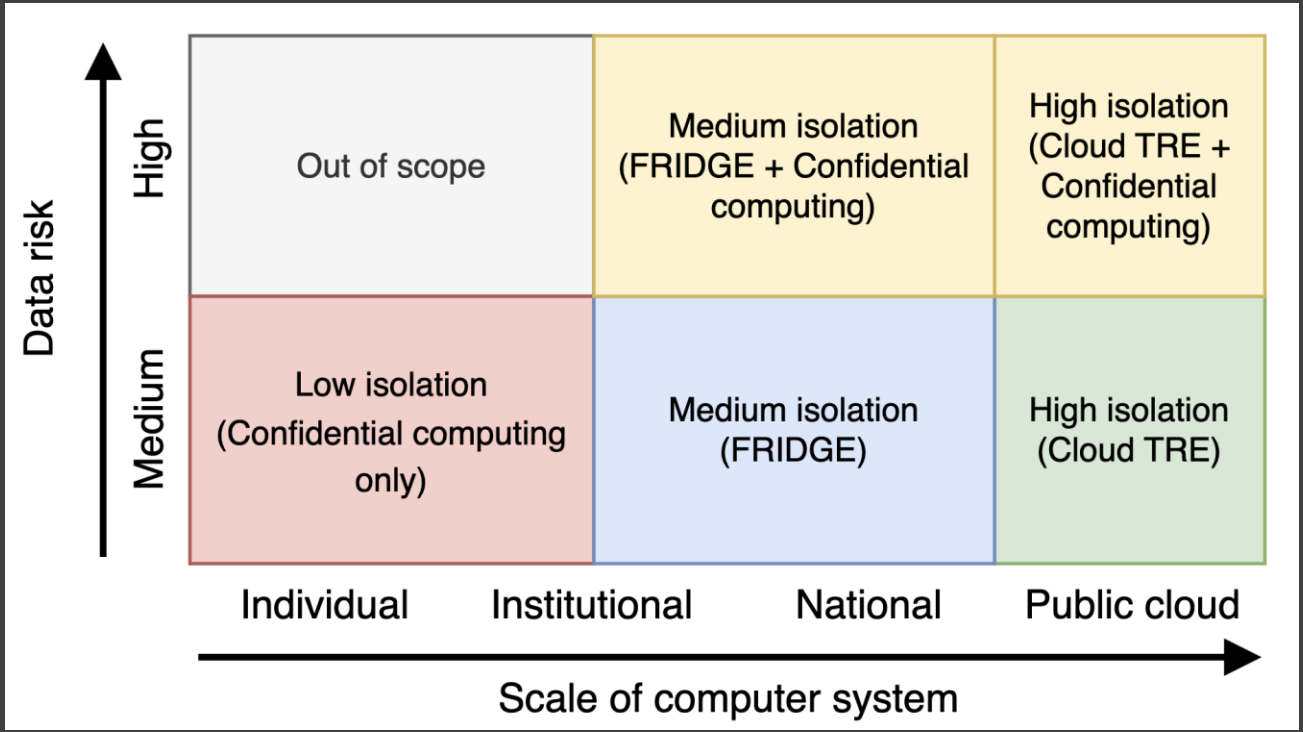

What are we doing?

What are we doing?

- **Goal**: Defining a bring your own compute model for TREs

What are we doing?

- **Goal**: Defining a bring your own compute model for TREs
- Evaluation of confidential computing capabilities of current CPUs and GPUs

What are we doing?

- **Goal**: Defining a bring your own compute model for TREs
- Evaluation of confidential computing capabilities of current CPUs and GPUs
- Validation of confidential computing approaches within TRE and supercomputer contexts

What are we doing?

- **Goal**: Defining a bring your own compute model for TREs
- Evaluation of confidential computing capabilities of current CPUs and GPUs
- Validation of confidential computing approaches within TRE and supercomputer contexts
- Sharing code and guidance for safely configuring confidential computing capabilities for TREs

Getting involved

Getting involved

- **Data provider**: Does the FRIDGE model work for your information governance context?

Getting involved

- **Data provider**: Does the FRIDGE model work for your information governance context?
- **TRE provider**: Can you support the FRIDGE satellite model?

Getting involved

- **Data provider**: Does the FRIDGE model work for your information governance context?
- **TRE provider**: Can you support the FRIDGE satellite model?
- **Supercomputer provider**: Can you support the FRIDGE satellite model? Do you have confidential computing enabled CPUs / GPUs?

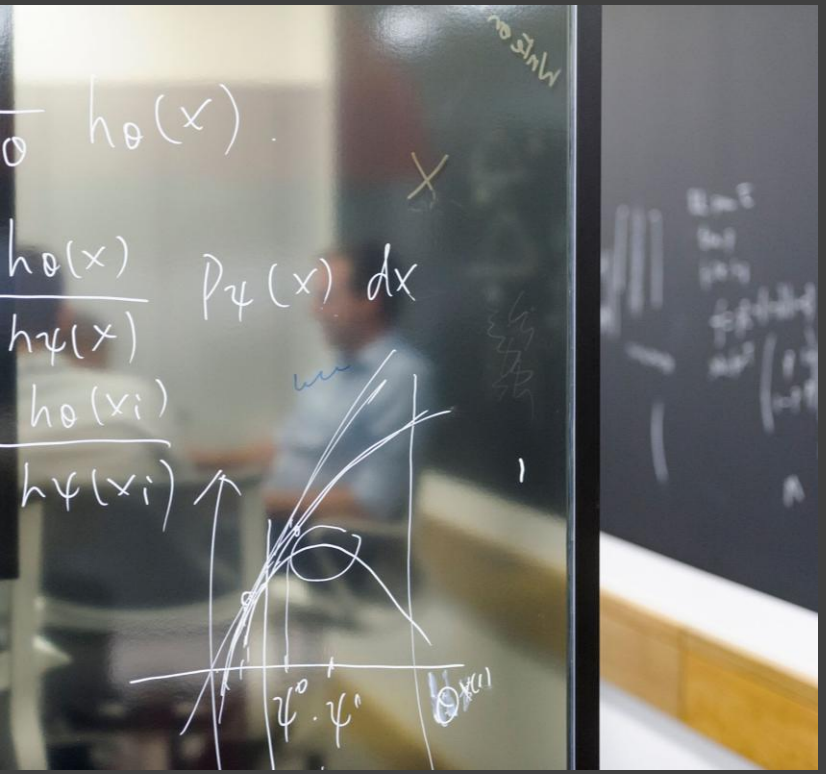

Getting involved

- Have a one-to-one chat
- Come to a workshop
- Try out FRIDGE on your system (from summer)

moreilly@turing.ac.uk | @martinoreilly.bsky.social | https://turing.ac.uk/research-engineering