The current state of assurance in establishing trust in PETs

By Georgia Iacovou and Alice Thwaite, Hattusia.

In this chapter, we will explore the current state of assurance schemes and governing standards for PETs. Our focus is on how these do (or don't) establish trust in those using the PETs (such as data controllers, researchers, or engineers), with a secondary focus on how this trust might filter down to those which the data is about (data subjects).

We will first outline definitions of key terms:

- **What a PET is, and the circumstances for using them**
- **The difference between privacy and security:** we will briefly demonstrate how these two concepts are often conflated, and how it's not clear which umbrella PETs sit under.
- **Assurance**: we also must understand what assurance is, and what kinds of assurances are necessary to support different trust relationships.
    - **Trust relationships:** as part of understanding assurance, we will also detail the kinds of trust relationships at play.

After outlining these definitions, we will identify the different actors involved in the use of PETs, what kinds of trust relationships they have, and which standards or assurances already exist to support these, and which are lacking.

Where there are gaps in assurances or standards, we will explore emerging data governance models which have the potential to provide assurances around not only the PETs themselves, but the contexts in which PETs are used.

A summary of the methods used throughout this analysis can be found at the end of the chapter.

Defining PETs: what are PETs, and what are the circumstances for using them?

For the purposes of this chapter, we understand that a PET is a name for a collection of techniques which currently lack a standardised conceptual boundary. In April 2021, the Ada Lovelace Institute described PETs in this way:

> "There is no single definition or standard for what constitutes a PET, though the term is typically used to refer to technologies or approaches that can help mitigate privacy and security risks. Some popular examples of PETs include forms of encryption such as format-preserving and homomorphic encryption, cryptographic protocols like secure multi-party computation and secret sharing, differential privacy and obfuscation techniques, and various means of anonymisation or pseudonymisation."[1]

This quote demonstrates that while all techniques that sit under the term 'PET' are meant to preserve privacy, they do so using different methods. They are also employed in different use-cases, for a range of different outcomes: secure multi-party computation may not be used in the same circumstances as differential privacy, for instance. In addition to this, PETs are used in combination with each other, which again will be for different outcomes and circumstances.

PETs are not designed for a specific purpose, or for use by a specific actor. So, as long as they are available, PETs can be used by banks, universities, or governments alike.

As such, each PET will support a different kind of trust relationship among its users. This means that the role of assurance will be different in each case. We will go into how different kinds of trust relationships require different kinds of assurances on page [BOOKMARK].

For the purposes of this chapter, we were instructed by the Royal Society to focus on the following types of PETs: homomorphic encryption, trusted execution environments, secure multi-party computation and differential privacy. These methodologies follow on from the previous Royal Society report on PETs, and were agreed by their roundtable of experts.

What's the difference between privacy and security?

Given that 'privacy' is a core part of the language of 'privacy-enhancing technologies', it is appropriate to define what privacy means.

What is privacy? The definition of privacy changes with the introduction of new cultures and technologies[2]. For the purposes of this chapter, we will take Privacy International's definition which states that "privacy enables us to create barriers and manage boundaries to **protect**

[1] Reneiris, E. (2021) Why PETs (privacy-enhancing technologies) may not always be our friends. Access via https://www.adalovelaceinstitute.org/blog/privacy-enhancing-technologies-not-always-our-friends/

[2] For a historical perspective, please see Arendt, H. (2018). *The Human Condition* (2nd ed.). University of Chicago Press.

ourselves from unwarranted interference in our lives": we need privacy for personal autonomy, for being ourselves without judgement and to think freely without discrimination.[3]

The concept of 'data privacy' often refers to the idea that we can protect the privacy of individuals as it's described above, while still using their data across a range of applications. For instance, The Electronic Frontier Foundation advocate for greater interoperability between online services, but data privacy needs to be taken into consideration because: "Policies designed to increase interoperability may weaken the tools that companies currently use to protect their users"[4].

What is security? Security is about **protecting information from outside actors**. The National Institute of Standards and Technology (NIST) defines it as "The protection of information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide confidentiality, integrity, and availability."[5]

So, securing data means that no other actor can access it without permission. **This is not the same as privacy:** e.g. institution X holds the data, and may suffer a security breach, meaning institution Y can access it without permission. This breach in security has nothing to do with privacy; institution X may have been using the data in a way that compromises personal autonomy, increases surveillance, and leads to unwarranted interference in our lives.

While security is needed for privacy, not all security leads to privacy. This is an important distinction, because security and privacy are often conflated, even among data professionals: this happens in discussions around PETs, because just as mentioned in The Society's original PETs report, PETs are useful because an organisation can run analysis on and retrieve insights from data, *without revealing data to unauthorised actors*.[6]

So PETs protect data from unauthorised access, **which is a function of information security**. But, depending on how they're used, they may also preserve privacy as part of a wider data governance framework.

We should also point out that defining privacy as 'unwarranted interference in our lives' will necessitate different actions in different contexts. However, there will be some overarching principles. To flesh this out, we refer to the Royal Society and the British Academy report which took a look at principles for **data governance.**

Data governance can: "mean everything designed to inform the extent of confidence in data management, data uses and the technologies derived from it"[7]. The Royal Society make recommendations on what good data governance looks like in *Data Management and Use: Governance in the 21st Century.* A key take-away is that **the promotion of human**

[3] https://privacyinternational.org/explainer/56/what-privacy

[4] https://www.eff.org/wp/interoperability-and-privacy

[5] https://csrc.nist.gov/glossary/term/infosec

[6] (2017) *Data Management and Use: Governance in the 21st Century,* the British Academy and the Royal Society. Access via https://royalsociety.org/-/media/policy/projects/data-governance/data-management-governance.pdf

[7] Ibid

flourishing is the overarching principle that should guide the development of data governance systems[8].

There are four principles for managing data in a way that promotes human flourishing[9]:

- Protect individual and collective rights and interests
- Ensure that trade-offs affected by data management and data use are made transparently, accountably and inclusively
- Seek out good practices and learn from success and failure
- Enhance existing democratic governance

However, if we examine the function of each PET we are concerned with in this report, it becomes very clear that they all keep data protected from unauthorised access (which is **security**), but not all of them necessarily protect privacy.

We have laid this out in Table I:

Table I: PETs, privacy and security

	Homomorphic encryption	**TEEs**	**SMPC**	**Differential Privacy**
What does this PET do?	Allows you to analyse encrypted data without decrypting it.	This is a secure, isolated environment in which you can run computations, such as a cloud platform.	This allows multiple parties to work together on analysing their respective pieces of data, without revealing the contents of the data to each other.	Mostly for use with large data sets, DP allows institutions to reveal data or derived information to others *without* revealing any sensitive information about the groups or individuals represented in the data set.
In what circumstances would it be used?	To create meaningful insights in research without revealing the contents of a dataset to those running the analysis (which could be done by a trusted third-party).	When data needs to be stored securely, but local machines and operating systems lack the features necessary to do this. Trusted execution environments also allow for running analysis, if again local systems are not equipped to do so.	Removes the need for a trusted central authority that would have access to everyone's data. Rather, multiple organisations can keep their data sets private from each other, but still run joint analysis on the combined data.	An institution may want to share key information that they have derived from their data with another group or with the public, but their data set contains sensitive information which should be kept private.
Who's data is being protected and	The data held by the institution running the research is being	Storing all data with a trusted third party in a highly secure	Each collaborating organisation holds data about individuals, and	Sensitive information about the groups or individuals present in

[8] Ibid
[9] Ibid

from who?	protected from whoever runs the analysis, whether a third-party or the institution themselves. If the third-party were to act in bad faith, they would not have access to the data in question.	environment protects it from any malicious actors who might target the research institution with a cyber attack. The data is also protected from any misconduct or incompetence coming from within the institution itself.	that data is protected from those collaborating on analysis. The data also is protected from any potential misconduct or incompetence from any of the parties.	the dataset is being protected from whoever the data is being shared with or analysed by, whether that's a trusted third-party, the general public, or the institution themselves.
Whose interests are being protected and what are they?	**The data controller**: they have an interest to carry out their research in the safest and most effective way possible. **The data subjects:** those who the data is about have an interest in making sure their data is not accessed by bad actors.	**The data controller**: they have an interest to carry out their research in the safest and most effective way possible. **The data subjects:** those who the data is about have an interest in making sure their data is not accessed by bad actors.	**The collaborating organisations**: they have an interest to carry out their research in the safest and most effective way possible. **The data subjects:** those who the data is about have an interest in making sure their data is not accessed by bad actors.	**The data controller**: they have an interest to carry out their research and share data in the safest and most effective way possible. **The data subjects:** those who the data is about have an interest in making sure their data is not accessed by bad actors.
Does this function fall under security or privacy?	**Security** – because the data is being protected from unauthorised access	**Security** – because the data is being protected from unauthorised access	**Security** – because the data is being protected from unauthorised access. **There is a privacy element** here, because data is being 'shared' so that multiple parties can work with it, but the data remains safe from unwarranted interference.	**Security** – because the data is being protected from unauthorised access. **Privacy** – because this technology provides the potential to give open access to data without infringing on the privacy rights of those the data is about

Table I: PETs, privacy and security

Understanding assurance through a trust lens

The understanding of assurance hinges on the understanding of **trust**: in December 2021, The Centre for Data Ethics and Innovation published a report entitled: "The Roadmap to an Effective AI Assurance Ecosystem". In this, they state that assurance systems are mechanisms for building trust:

> "Assurance is about building confidence or trust in something, for example a system or process, documentation, a product or an organisation. [...] These assurance services provide the basis for consumers to trust that the products they buy are safe, and the confidence for industry to invest in new products and services."[10]

The above paragraph illustrates that trust is an integral part of assurance; in order to develop effective assurances for something, we must understand **what kinds of trust** those assurances are meant to support, and **who the trust is between (i.e. the trust relationship).**

In order to understand more about trust, let's discuss how the most respected thinkers in this field define trust, and how trust is developed.

In *Who Can You Trust*,[11] Rachel Botsman defines trust as **a confident relationship with the unknown.**

There are a number of ways of interpreting this depending on who the actors are and the 'unknown' in question; I may trust the same actor to do one thing, but not do another thing. For example, I may trust a cab driver to take me to my destination safely, but I wouldn't trust them to manage my passwords.

Trust relationships involve two actors:

1. **Trustors:** the person or party who gives trust. They will hold their own individual attitudes, trusting beliefs, and "generalized faith in humanity".[12]
2. **Trustees:** the person or party who receives trust. This could be an individual, an institution, the government, or even a profession.

So the trustor will be confident in the trustee based on a range of factors, some of which are dependent on context. **If you do not have complete certainty over a particular outcome, that's where trust comes in.**

Trust also involves human agency; we make conscious decisions to trust, and therefore understand that disappointment *may* follow. One of the most critical and established players in this space is Onora O'Neill, who writes about the ethics of communication. Her 2002 Reith Lectures have inspired a lot of thinking around trust, trustworthiness and transparency.

---

[10] (2021) *The roadmap to an effective AI assurance ecosystem.* Centre for Data Ethics and Innovation. Access via: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1039146/The_roadmap_to_an_effective_AI_assurance_ecosystem.pdf

[11] Rachel Botsman. (2017). *Who Can You Trust? How Technology Brought Us Together and Why It Might Drive Us Apart*. Perseus Books, USA. p.257

[12] Becker, M. & Bodó, B. (2021). *Trust in blockchain-based systems*. Internet Policy Review, 10(2). https://doi.org/10.14763/2021.2.1555

> "And when we place trust we don't simply assume that others are reliable and predictable, as we assume that the sun rises reliably, and the milk goes off predictably. When we trust we know – at least when we are no longer small children – that we could be disappointed."[13]

There is a difference between **knowing** that the sun will rise tomorrow, and **trusting** that the tomatoes sold in your local supermarket are not poisonous.

Secondly, we shouldn't trust every person and institution blindly – instead they should be **trustworthy**. One way that an institution can demonstrate their trustworthiness is to have governance and processes in place which ensure they are accountable, and they are working to a high standard.

The 2020 Edelman Trust Barometer identified two types of trust[14]:

1. **Moral[15]**: the trustor believes the trustee can articulate and act on the best interests of the trustor.
2. **Competence:** the trustor believes the trustee has the ability to deliver on what has been agreed.

As we will explore further on in the chapter, **trust in competence** will also apply when the PET – the technology itself – is the trustee. It's important to highlight this, because the user of a PET needs to trust that it will work for its intended purpose. Therefore the PET itself (or a system of PETs) is a key part of the trust relationship at play.

The Edelman Trust Barometer report also explains the importance of moral trust in institutions:

> "Trust is undeniably linked to doing what is right. After tracking 40 global companies over the past year through our Edelman Trust Management framework, we've learned that ethical drivers such as integrity, dependability and purpose drive 76 percent of the trust capital of business, while competence accounts for only 24 percent."[16]

In both these cases, there needs to be an established idea of what 'good' looks like. What does it mean for someone to be a competent professional? How can you know what someone's best interests are?

It's often argued that transparency is a key prerequisite for accountability, and that then leads to trust. But really, transparency is not a silver bullet for trust: it is neither necessary for achieving trust, nor sufficient just on it's own. However, it might be a good stepping-stone to demonstrate institutional trustworthiness. It's important to remember that full transparency

---

[13] http://downloads.bbc.co.uk/rmhttp/radio4/transcripts/20020410_reith.pdf
[14] https://www.edelman.com/trust/2020-trust-barometer
[15] The Edelman Trust Barometer uses the term 'ethical' here, instead of 'moral'. We opt to use the word moral because their definition of ethics does not match how Hattusia defines ethics (which is 'the study of how humans should live'). We are happy if the Royal Society wants to adopt the word 'ethical' instead for this chapter - and we've discussed this with June.
[16] https://www.edelman.com/trust/2020-trust-barometer

makes trust irrelevant[17]; if everyone knows anything about the inner-workings of an institution, there's no room for building a confident relationship with the unknown.

This brings us to what good governance is: For example, you can trust that the electrical wiring in your house is in good working order, because it lives up to certain standards. You're sure that when you leave your phone charging overnight, it won't start a fire. You are *aware* that there is a governance model for this in place, but you don't have full knowledge of how it works — you trust the model.

So, if standards are the method by which we maintain 'good' practice, we must identify what 'good' looks like, and work from there. For electricity, we all already know that 'good' means safeguarding against fire and electrocution. But this might be a little less clean cut in other areas. So, we should ask: what does 'good' look like in data privacy?

Key questions that should be answered in order to build trust in PETs:

1. What does 'good' look like in data privacy?
2. What does it mean to be a competent professional in data privacy?
3. What does it mean to have a 'competent' technology or technological system?
4. How can trustees demonstrate to trustors they have their best interests at heart?

[17] O'Neill, Onora (2002). *A Question of Trust: The Bbc Reith Lectures 2002*. Cambridge University Press. Retrieved from: http://downloads.bbc.co.uk/rmhttp/radio4/transcripts/20020427_reith.pdf

Assurance and PETs

When defining trust, we quoted a CDEI report called “The Roadmap to an Effective AI Assurance Ecosystem”. Here they state that assurance systems are mechanisms for building trust, and that assurance services can allow consumers to trust the products they buy, and the confidence for industry to invest in new products and services[18].

So in their definition of assurance, they have identified two ‘trustors’: one is the consumer (or the general public) and the other is the industry itself. Later on in the report they separate out the ‘industry’ into two distinct actors: ‘executives deploying AI systems’ and ‘front line users’.

To help us understand more about how assurance works with PETs, lets apply this basic framework of trustors to data governance:

- The ‘executives deploying AI systems’ become executives deploying data-led strategies and research
- The ‘front line users’ become the users of the PETs themselves (such as engineers or data scientists).
- The ‘consumers’ become data subjects (i.e. those that the data is about, including organisational data such as data referring to revenue or productivity).

These trustors will have different kinds of trust relationships with each other, and therefore require different kinds of assurances. For example, the user of the PET must be assured that the PET will work for its intended purpose, which is a trust in competency. Whereas the user’s trust in the executive is moral, because the user of PETs needs to be assured that what they are being asked to do is ‘right’.

Diagram I illustrates the flow of assurance between these three trustors, and what kinds of trust relationships sit between them.

[18] (2021) *The roadmap to an effective AI assurance ecosystem.* Centre for Data Ethics and Innovation. Access via: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1039146/The_roadmap_to_an_effective_AI_assurance_ecosystem.pdf

Diagram I: The flow of assurance for PETs between key trustors

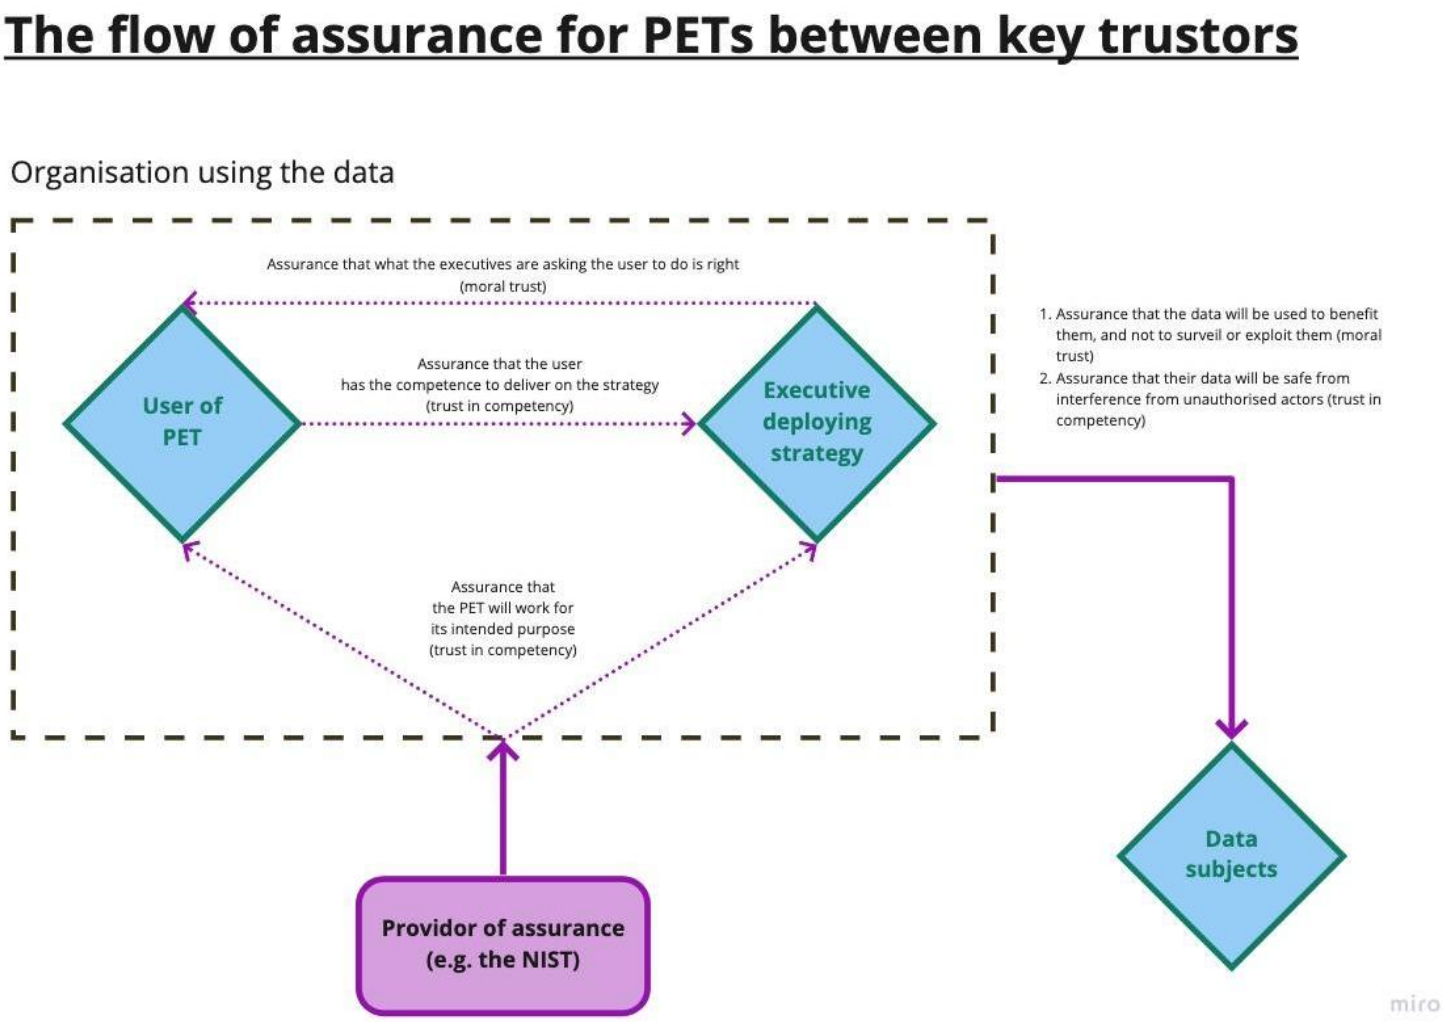

Diagram I - the flow of assurance for PETs between different trustors[19]

In this diagram we can see that both moral trust and trust in competency are required in the assurance of PETs overall. This means that the kinds of assurances needed around PETs are more than just technical ones, which simply assure the user that the PET will work. If moral trust is at play, **there need to be assurances on the applications of PETs too,** e.g. standards which dictate when it is or isn't appropriate to use a certain PET. So, to be clear:

- Trusting that the technology will work is a trust in competency, which requires **technical assurance**
- Trusting that the technology is being used in a way that is 'right' is a trust in morals, which requires **assurance (or standards) in the application** of the technology.

Another critical part of Diagram I is the 'provider of the assurance'. While conducting pilot research for this project, we developed the following taxonomy of assurance providers.

[19] See https://miro.com/app/board/o9J_ljPbmnM=/?moveToWidget=3458764515194701407&cot=14

1. **Legal standards**: these are standards which are enforced by law and where there is legal accountability to them. An example might be UK Health and Safety legislation[20], or the legal standards which exist to ensure the safe electrical wiring inside private properties in the UK[21]. For the application of PETs, legal standards can be exemplified in regulations such as the GDPR.
2. **Non-legal standards**: these are non legally binding standards, but are dictated by a third party organisation, which occasionally awards a kitemark or a certificate to products and services which adhere to the ethical standard. For example, a coffee distributor may be awarded the 'Fair Trade' kitemark when they pay a fair and minimum price to the coffee producers.[22] In technology, a non-legal standard could be one provided by a body such as the ISO.
3. **Educational and competency standards**: these standards demonstrate that an organisation, or a practitioner within an organisation, has a certain level of skill which mandates they can have a certain level of responsibility. For example, a surgeon will have passed many exams in order to complete an operation. In other cases, a person may require a certification which states they have completed a course, such as in first-aid standards. Some educational standards are required legally, and some are enforced culturally. A culturally enforced standard may be the desirability of an undergraduate degree for certain jobs. In cybersecurity, it's widely accepted that being a Certified Information Systems Security Professional qualifies you to be a chief information manager.
4. **Reputational assurance:** this is not a standard as such, because one does not acquire a certification which states that they have a good reputation, but instead this is earned through undocumented social means. Nevertheless, it is an important part of assurance. There is no clean-cut example for this in technology, because there are many theories as to why we favour some tools over others. E.g. a social platform may not be useful unless 'everyone' is using it.

N.B. This taxonomy was sense-checked with members of the Royal Society, as well as informal conversations with Hattusia's network.

[20] https://www.hse.gov.uk/legislation/index.htm
[21] https://www.gov.uk/government/publications/electrical-safety-standards-in-the-private-rented-sector-guidance-for-landlords-tenants-and-local-authorities/guide-for-landlords-electrical-safety-standards-in-the-private-rented-sector
[22] https://www.fairtrade.net/standard/aims

What types of assurances are needed to support the trust relationships present in the use of PETs?

Now that we've defined assurance, trust, PETs, and the difference between privacy and security, we can look at the kinds of trust relationships required for using PETs, and then the assurances needed to support those relationships. The trustors in Table II are the actors we identified above when defining assurance (taking inspiration from the CDEI[23]):[24]

Table II: Trust, trustors, trustees

Trustors	**Trustees**	**Moral trustworthiness**	**Trust in competence**	**Assurances needed**
The 'user' of PETs e.g. engineers or data scientists	The technology itself; collaborators; other actors who may try and get some confidential information; the executives.	Do they trust the morality of the executives in terms of whether what they are being asked to do is the right thing?	Will the PET fulfil its expected function? Will the data remain secure from outside actors who want access to it? Does the PET user have knowledge about whether the PET is appropriate for the task at hand?	**Technical assurance** in the technological systems, which would come from institutions such as the NIST. **Assurance in the application** of the PETs, which is assurance that the executives are using good participatory data governance models.
Executives deploying data-led strategies and research	The user of PETs; the technology itself	N/A	Do the developers have the competence to deliver on their strategy? Will the PET fulfil its expected function?	**Technical assurance:** Professional qualifications or certifications which prove the user knows how to deploy PETs correctly. **Technical assurance** in the technological systems, which would come from institutions such as the NIST.

[23] (2021) *The roadmap to an effective AI assurance ecosystem.* Centre for Data Ethics and Innovation. Access via: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1039146/The_roadmap_to_an_effective_AI_assurance_ecosystem.pdf

[24] Please note that this table is based on Hattusia's experience in this field and does not constitute robust research in the area. We recommend that more work is done to understand the best assurance schemes for data privacy in the context of trustors, trustees, moral trustworthiness and trust in competence.

The people whom the data is about (data subjects)	The ecosystem of the people and organisations that collect, sell, and use the data	Will their data be used in a way that benefits them, and not lead to increased surveillance and exploitation.	Will their data remain safe from interference from unauthorised actors?	**Assurance in the application** of the PETs, which is assurance that their data is being used in good participatory data governance models.

Table II: Trust, trustors, trustees

The user of PETs as the trustor

As outlined in our table, **the user of the PET must trust that the technology itself will work for its intended purpose,** which requires a **technical assurance**. The user of the PET also trusts that executives directing the work are asking them to do the right thing. This requires **standards in the application** of the PET to be followed.

At the moment, the only assurances or standards that exist for our chosen PETs are technical ones. There are very few of these, and they are very new -- they detail how to apply specific PETs across a range of possible use-cases, and the bulk of them have only been published in the last three to six months.

For an overview of our methodological process to gather this information please see the end of this chapter.

The majority of specific standards we found came from the following bodies:

- The Institute of Electrical and Electronics Engineers (IEEE)
- The International Organisation for Standardisation (ISO)
- National Institute of Standards and Technology (NIST)

These have been presented in Table III

Table III: Non-legal standards in PETs

	Homomorphic Encryption	**Trusted Execution Environments**	**Differential Privacy**	**Secure Multi-Party Computation**
IEEE	There is not a standard on homomorphic encryption but a standard on Biometric Privacy includes info about it[25]. *May 2021*	IEEE 2830-2021: IEEE Standard for Technical Framework and Requirements of Trusted Execution Environment based Shared Machine Learning[26]. *October 2021*		IEEE 2842-2021: IEEE Recommended Practice for Secure Multi-Party Computation[28]. *November 2021*
		P2952 - Standard for Secure Computing Based		

[25] https://standards.ieee.org/standard/2410-2021.html
[26] https://standards.ieee.org/standard/2830-2021.html
[28] https://standards.ieee.org/standard/2842-2021.html

		on Trusted Execution Environment[27]. *Not yet published*		
ISO	ISO/IEC AWI 18033-8 Information security — Encryption algorithms — Part 8: Fully Homomorphic Encryption[29]. *Under development*			ISO/IEC WD 4922-2.3 Information security — Secure multiparty computation — Part 2: Mechanisms based on secret sharing[31]. *Under development*
	ISO/IEC 18033-6:2019 IT Security techniques — Encryption algorithms — Part 6: Homomorphic encryption[30]. *May 2019*			ISO/IEC CD 4922-1.2 Information security — Secure multiparty computation — Part 1: General[32]. *Under development. .*
NIST	Blog series published on this technology but there is (as yet) no standard[33]	Hardware-Enabled Security: Enabling a Layered Approach to Platform Security for Cloud and Edge Computing Use Cases (2nd Draft)[34]. *October 2021*		

Table III Non-legal standards in PETs

Other relevant standardisation projects include:

- **Homomorphic Encryption Standardisation**: an open consortium set up specifically to set standards for homomorphic encryption[35]
- **The PEC Project:** the NIST are working on setting general standards for cryptography[36]
- **PETs Adoption Guide**: the Centre for Data Ethics and Innovation have a general guide for best practices in adopting PETs[37]

It's clear that these standards and assurances only play a part in assuring the user of the PET that the PET will fulfil its intended function, which is to protect data from access from unauthorised actors. None of these address the privacy concerns outlined in our definitions, nor are there any standards around the application of PETs. There are however emerging

[27] https://standards.ieee.org/project/2952.html
[29] https://www.iso.org/standard/83139.html
[30] https://www.iso.org/standard/67740.html
[31] https://www.iso.org/standard/80514.html
[32] https://www.iso.org/standard/80508.html
[33] https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/collaboration-space/focus-areas/de-id/dp-blog
[34] https://csrc.nist.gov/publications/detail/nistir/8320/draft
[35] https://homomorphicencryption.org/
[36] https://csrc.nist.gov/projects/pec
[37] https://cdeiuk.github.io/pets-adoption-guide/

standards and governance models for data in this regard, which we will cover when examining the assurances required for data subjects.

The executive as the trustor

In our table of trustors, we outline that the executives deploying data-led strategies and research will have to put their trust into the users of PETs to use the PETs correctly. This is a **trust in competence** and therefore requires **technical assurance**.

As discussed throughout this chapter, PETs fulfil a function that is much closer to information security, than to privacy. So to understand the kind of technical assurance that the executive might need from the user of PETs, we can look at some existing cybersecurity professional qualifications.

Cybersecurity qualifications, certifications, or accreditations for individuals offer assurance that those individuals are able to carry out work in cybersecurity. E.g. understanding how to set up an office network so that all communications, such as emails or video conferences, are properly encrypted.

Here are a few examples of cybersecurity qualifications. Earning these demonstrate varying degrees of competence in the individual:

- Certified Information Systems Security Professional (CISSP): this is an internationally recognised qualification designed for positions such as Chief Information Officer[38]. In the UK, CISSP is considered to be equivalent to a masters degree[39].
- Certified Information Security Manager (CISM): CISM differs slightly from CISSP in that it will validate expertise, but offer less support in training to achieve this[40].
- Certified Cloud Security Professional (CCSP): This certification is offered by the same body as CISSP[41], and is notable in that it provides training specifically for those working with cloud security[42]. This certification is therefore somewhat relevant to trusted execution environments.

From these we can see that in cybersecurity, there are well-established schemes to certify that individuals are qualified to use the relevant tools and methods required for their work. In the assurances and certifications found for PETs, there were no such assurances. While there are assurances that the PETs will fulfil their function, there are no certifications that an individual can earn to prove that they are qualified to use PETs.

The data subject as the trustor

The table of trustors shows that the data subjects, which could be individuals or organisations, need to trust that their data is safe from access by unauthorised actors, which

[38] https://www.isc2.org/Certifications/CISSP#
[39] https://www.infosecurity-magazine.com/news/cissp-equal-masters-degree/
[40] https://www.isaca.org/credentialing/cism
[41] https://www.isc2.org/
[42] https://www.isc2.org/Certifications/CCSP#

is a **trust in competency**, and therefore requires technical assurance from data controllers. This is very much a function of information security.

Crucially however, data subjects also need to have a **moral trust** that their data will be used for their benefit, and not for surveillance or exploitation. This is therefore a function of data privacy, so the public require assurance PETs will be used in appropriate applications – this assurance is more nuanced, because the general public will not necessarily be aware of PETs or how they work.

In our definitions, we quoted Rachel Botsman who says trust is a confident relationship with the unknown; this means that we can put our trust in complex systems without fully understanding how they work. We are somewhat aware that there are **governance models** at play in our food supply chains, which mean that the tomatoes we buy in the supermarket will be safe to eat. We do not know how these models work – we **trust** that they work.

This same principle can potentially be applied to data governance: can data subjects be aware of a system that is governing their data, and what makes this system ‘good’?

Diagram II: The assurances needed to support the types of trust at play in PETs

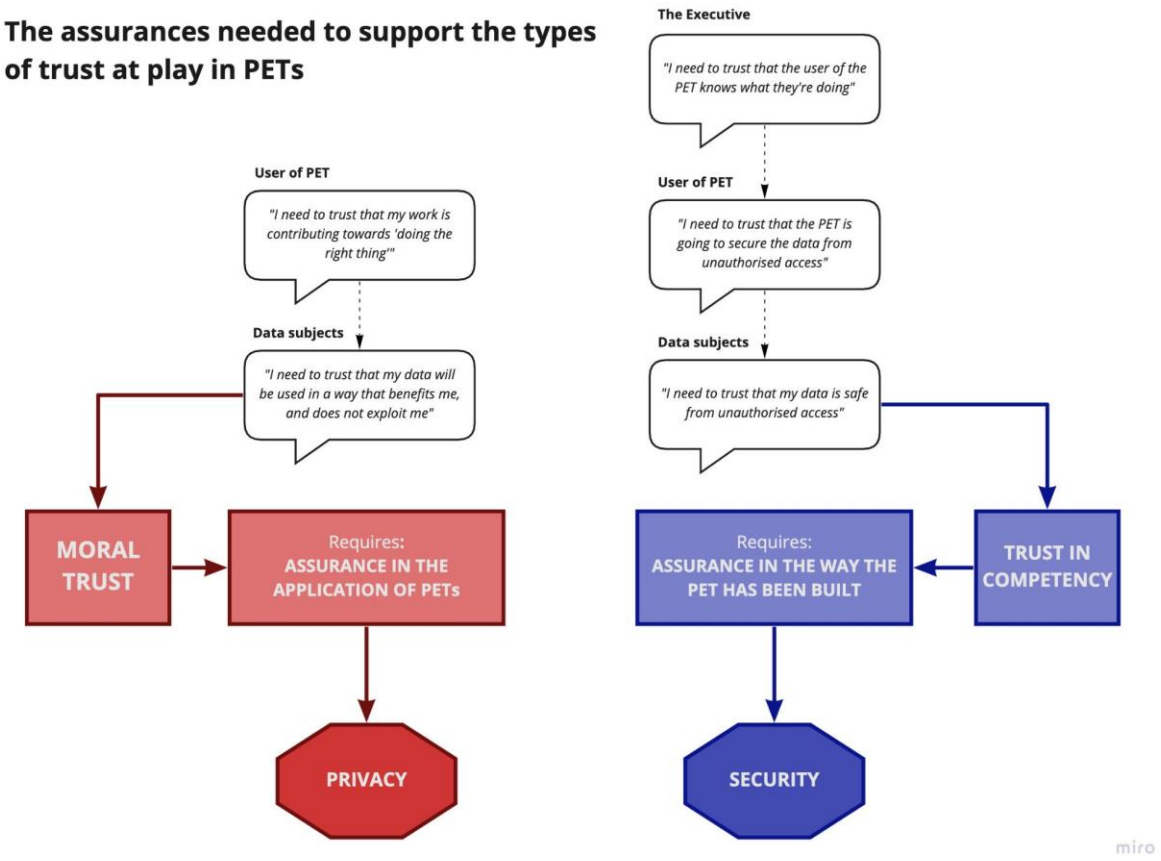

Diagram II The assurances needed to support the types of trust at play in PETs[43]

New data governance models can provide assurance for PETs

As we've demonstrated, the trust relationships present in the use of PETs are not solely based on trust in competence, but also trust in morality. Moral trust requires assurance that goes beyond just the technical. Here we will discuss ideas around data governance which can potentially provide assurance for all of our identified trustors, and the kinds of trust relationships they have with each other.

An emerging facet of data governance which holds a lot of potential is the **data trust** model. A data trust is a legal model which provides an alternative way to process data: data trusts are made up of **trustees** who have a legal duty to process data in the best interest of the **beneficiaries** – the people who the data is about.

The trustees duties are as follows:

[43] See https://miro.com/app/board/o9J_ljPbmnM=/?moveToWidget=3458764515586915869&cot=14

- **Duty of loyalty**: they can only act in the best interest of the beneficiaries.
- **Duty of independence**: they not only act in the best interest of the beneficiaries, but in the *sole* interest -- so they cannot have an interest in anything else, including the data.
- **Duty of care**: they must keep the data secure. E.g. do not send the data through an encrypted network.

These duties are put in place because there is an assumption that there is information asymmetry between the trustees and beneficiaries. This dynamic is similar with patients and doctors -- we visit doctors because they have knowledge that we do not have. The legal duty of a trustee also means that they will suffer consequences if they act in bad faith or make serious mistakes.

The data trust must also have a strong purpose which is well-defined and has measurable outcomes. For instance, a trust that says it will use data to 'improve health in the UK' is too vague, and has too many potential metrics. A more specific purpose such as 'to cure lung cancer' would be more appropriate for a data trust. The idea that data trusts, and the gathering of data in general should have a strong purpose, is one put forward by Anouk Ruhaak, a fellow at the Mozilla Foundation. Ruhaak explains:

> "As it turns out, getting data governance right is hard and highly context dependent. Governing financial data in the context of banking is a different thing from governing data pertaining to the milk production of cows, which is altogether different from the governance of health related data by our medical practitioners. And in each of these contexts, the specific use of the data and problem it attempts to solve further impacts how it needs to be governed."[44]

There are also arguments in data governance for **data minimisation**; that models should be built around what problem they are trying to solve, and not what data should be collected. Data governance expert Sean MacDonald has argued this point by using the NHS as an example:

> "The NHS explained, for example, that its data grab is to "save lives," which is an impossibly broad and endlessly reusable purpose. Contrast the breadth of purpose with something like "develop treatments and cures for COVID-19," which is a more specific pursuit, with clearer boundaries, indicators for success, and – to an extent – logic for prioritizing resources."[45]

This demonstrates that a trustworthy governance model must have a well defined purpose, and that **the purpose must come before the data;** preventing credit card fraud and ensuring traffic flows around a city efficiently are two very distinct challenges that require two different kinds of data.

---

[44] Ruhaak, A. (2021) Getting started with data governance. Access via: https://foundation.mozilla.org/en/blog/getting-started-with-data-governance/

[45] McDonald, S. (2021) Data governances' new clothes. Access via: https://www.cigionline.org/articles/data-governances-new-clothes/

Currently, all the PETs we have examined in this report, and the few standards and assurances that follow them, **are not concerned with enhancing privacy at the stage of collecting or creating data**. The PETs in question are techniques which are applied at the analysis stage, *after* data has already been collected.

The data governance models we have discussed here therefore **have the potential to provide assurance that supports moral trust** because if a governance model is built around a specific problem or purpose, the users of the PETs, executives, and data subjects will all have the awareness that the only data being collected, is the only data that is needed.

Summary and further considerations

Assurance in the realm of PETs is at a very early stage. There are technical standards, which are still being developed, which assure us that PETs will fulfil their intended purpose. While these standards assure data controllers that they can complete their work effectively, without exposing data to unauthorised actors, there are not yet any assurances around the applications of PETs. Take for example homomorphic encryption: there's nothing stopping a bad actor from using this technique on an encrypted dataset that they acquired by questionable means.

The adoption of PETs themselves is also at an early stage; as we've argued, there's no consensus on whether PETs currently sit, or indeed whether they should sit, under the umbrella of security or privacy. PETs lack a conceptual boundary – they are a set of techniques that can offer particular protections over data, but they are not designed for any specific use-cases. They can be used on any data, and by anyone. As such, **it's challenging to apply standards or assurances which support moral trust relationships**: the trust not just in that the PETs will work, but also in that the PETs are being used in a way that promotes human flourishing, and within a data governance framework that improves well-being, rather than increase surveillance on the public.

This moral trust relationship is required between many different parties. It is required between:

- the data subject and the ecosystem,
- the PET user and the executives they work for,
- and also between each and every organisation which uses and shares data to gain insights about the world we live in.

Therefore, assurances and standards going forward should address the context in which a PET is being used, not just the PET itself.

Assurances and standards on PETs should also be developed in a future-facing manner. The use of PETs at scale will have an effect on the environment because of the computational power needed to facilitate these techniques. For example, trusted execution environments rely on cloud computing, and cloud computing platforms contribute to higher carbon emissions, and these are rising annually[46].

Furthermore, future technologies should also be taken into consideration. Current encryption methods are designed to protect against the computing power we have at our disposal today, but would be ineffective against quantum computing, for instance. A recent study by quantum physicists has shown that using a quantum computer would make it possible to break encryption within the next 25 years[47]. This indicates that there is a need to consider new forms of encryption, and therefore new standards for PETs.

---

[46] *Climate Risk Analysis* by The Shift Project: https://theshiftproject.org/en/en-notation-climat/

[47] *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits* by Craig Gidney, Martin Ekerå; revised April 2021. https://arxiv.org/abs/1905.09749

Ultimately, our main recommendation is that more research should be conducted to understand the wider context in which PETs sit. Can assurances that PETs are being used by data controllers be part of a wider assurance to all actors that their data is being used in a way that benefits them? This much larger question on the role of assurance should be examined in greater detail.

The development and use of PETs are very much in their infancy, and therefore so are the assurances and standards surrounding these; when developing standards, we must be sure to understand not only what function PETs are meant to fulfil, but the contexts in which we use PETs, and at what scale.

Information on methods

In order to construct this chapter, we primarily adopted the Western analytic philosophical tradition to create normative arguments. These arguments stem from premises that are either analytic, or synthetic.

Analytic statements are logical definitions of a word or a concept; we backed these up by positions taken by credible organisations. In some cases we brought together many expert positions, such as in the section ‘understanding assurance through a trust lens’. This may be shown to be wrong if we have missed out a vital component to the terminology.

In contrast, synthetic statements are descriptions of the world based on experience. These might come from social science methods, computer science or natural sciences. In some of our work in this chapter, we’ve had to rely on our pre-existing expertise to make assumptions where there is a lack of research available. In order to make this more robust, we would advise further research using qualitative surveys or other methods. For example, Table II on trust, trustors and trustees is based on

In other parts, we’ve conducted our own desk research - the method for Table III can be seen below.

Table III: Non-legal standards in PETs

In order to uncover these assurance schemes we conducted desk research using Google search, the English language and an IP address based in London, UK.

We based our search criteria on a pilot piece where we looked into what kinds of assurance schemes were available in general in Anglo-American societies, as well as assurance schemes in adjacent fields such as cyber security and open data policies. This pilot project started with assumptions about types of assurance service, and we developed this taxonomy:

We were advised to concentrate on standards which would be of most use to the user’s of PETs, which meant that we chose to look at legal, non-legal and educational standards as opposed to reputational or educational standards. In our pilot, we found no legal nor educational standards that involved PETs. We identified a number of organisations which were most likely to have developed, or be developing, non-legal technical standards. These were:British Standards Institution, CREST, Cyber Essentials, ENISA, IAPP, IEEE, ISO, MCSS, NESAS scheme, NIST, PCI.

We also agreed with the Royal Society the versions of PETs that we would seek assurance schemes for. These were: homomorphic encryption, trusted execution environments, secure multi-party computation and differential privacy. These methodologies follow on from the previous Royal Society report on PETs, and were agreed by their roundtable of experts.

For each of these organisations we used the ‘in’ function on Google search to look within the body of their website, and searched for the *name of the PET*.

For example:

“in:nist.gov secure multi party computation”

Table III of our search criteria:

By PET	Differential Privacy, Homomorphic encryption, SMPC, TEEs
By Standards / Assurance body	British Standards Institution, CREST, Cyber Essentials, ENISA, IAPP, IEEE, ISO, MCSS, NESAS scheme, NIST, PCI

In our pilot we developed a series of criteria through which we were going to evaluate the standards. However, given the standards were in such an embryonic state, this part of the research was discarded. These evaluations are shown in Table IV:

By data type	Electoral data, Health data, Transport data
Who's privacy is primary?	Data controller, Data subject, Third party researcher
Which actor do we want privacy from?	A government, A malevant actor, A technology, Commercial organisation
Privacy rights and protections	Protected sensitive characteristics, Right to change/delete data, The right to anonymity

Please note, a key limitation in this work regarded the paywalled status of the standards. We made the choice, in conjunction with the Royal Society, not to pay for access to the standards which meant that we couldn’t look into the technical details they proposed.