APPENDIX 10: Safety Assessment

Safety obligations when sharing data

Legal requirements to protect data

The SynD organisations have legal obligations to maintain data quality, manage data security, ensure data is disposed of appropriately, and take accountability for the data shared internally or externally.

These legal obligations come from the Commonwealth and State and Territory privacy laws, as well as the range of State Records and Archives laws.

If personal information is leaving a state or territory

If data about identifiable individuals is to be disclosed to a recipient outside of the state or territory where it is currently held, the organisation must take reasonable steps to ensure that the information will not be held, used or disclosed by the recipient of the information inconsistently with the organisation's legal privacy obligations.[58] This is typically satisfied by way of a contract / agreement or confidentiality deed.

Data about identifiable individuals (such as source data, or synthetic health data that is still 'health data') should be held by the organisation who collected it as a rule (including when stored on secure cloud platforms) unless the SynD organisations have completed a risk assessment and agreed on an alternate storage location.

Keeping data secure in transit

There are multiple techniques for 'sharing' data: data might be 'shared' by uploading data using a secure file transfer protocol or a system such as Kiteworks, building a dashboard report via an API, providing a Data Requestor with direct access to a data warehouse, or by providing a Data Requestor with a data extract.

However, whether or not a particular technique can be considered 'safe' will depend on the context. Some techniques, such as emailing records without further security safeguards such as encryption or password protection, will not be considered safe. The security of the proposed method or technique for sharing data should be assessed prior to the sharing taking place to ensure it is suitable and meets any additional organisational requirements.

Disposing of data appropriately

The rules for data retention and disposal are set by the State Records and Archives laws and applicable privacy laws working together. Organisations must set retention periods for any synthetic health datasets generated under this Framework. As soon as possible after that period has expired, the data must be deleted or securely destroyed from all production and non-production environments (unless going into permanent government archives).

[58] IPP 12 (s.19(2)(g) of the PPIP Act), and HPP 14

When sharing synthetic health data amongst the SynD organisations, the Data Provider will need to work with the Data Requestor and specify how long the Data Requestor is allowed to keep the data, and what their obligations are in relation to either returning or destroying the data at the end of that period.

Assurance

Once synthetic health data is approved for sharing, steps should be taken to ensure that the conditions required under the Data Sharing Agreement are being met, and to address any deficiencies. Types of assurance activities that could be planned and carried out include:

- Reviewing relevant current security certifications or system specifications
- Confirming that data is stored and maintained in the approved system and that IT controls are in place and effective
- Verifying any required training has been conducted (e.g. privacy and security training for End Users, or data analytics training)
- Ensuring that access to the data is being managed appropriately, and that End User access is being revoked in a timely manner when no longer need
- Checking data access logs periodically for any unusual behaviour or unauthorised accesses
- Checking any outputs arising from the analysis to confirm they are aligned with the approved use case
- Confirming that any required assessments (e.g. technical security assessments) have been completed and the results are acceptable
- Confirming that any third-party service or supplier arrangements are appropriately managed
- An audit of an organisation's compliance with the Data Sharing Agreement

Organisations with formal assurance policies and processes should apply these to the data sharing arrangement.

Taking a 'five safes' approach to generating, using and sharing synthetic health data

Every instance of generating and handling synthetic health data carries some privacy risks, so the benefits of each use case need to substantially outweigh those risks.

The Five Safes Framework[59] offers a useful way of thinking about how to control for *two* particular types of privacy risks:

[59] Australian Computer Society, *Data Sharing Frameworks: Technical White Paper*, September 2017; available at https://www.acs.org.au/content/dam/acs/acs-publications/ACS_Data-Sharing-Frameworks_FINAL_FA_SINGLE_LR.pdf

- Inadvertent disclosure (also known as 'statistical disclosure'), such as *re-identification risk*, and
- Misuse of data by authorised users (data recipients) – in other words, *misuse risk*.

The Five Safes Framework was designed to manage these two types of privacy risk *in a particular context*: the sharing of data in a controlled environment, in which a party will perform analytical operations on the data, and then share the results of the analysis.

The Five Safes Framework is not a legal requirement, and it does not override our legal obligations. It is not a way of measuring the level of 'identifiability' of data. The Five Safes Framework is not an assessment of whether it is lawful, or ethical, to engage in generating and handling synthetic health data. That is why working through *all* steps and assessments in this Framework is critical. However, the Five Safes Framework is useful as a way of thinking about risk management when handling data.

Each 'safe' refers to an independent but related aspect of managing these two types of privacy risk, in the context of handling data in a controlled environment:

- Safe data – Has appropriate and sufficient protection been applied to the data?
- Safe projects – Is the data to be used for an appropriate purpose?
- Safe settings – Does the access environment prevent unauthorised use?
- Safe people – Is the requestor or end user appropriately authorised to access and use the data?
- Safe outputs – Are the statistical results non-disclosive?

These five elements are intended to be viewed wholistically to create an *overall* level of safety, in which the different elements may balance each other out. In other words, if one type of control has been 'dialled up', it may be safe to 'dial down' another control.

The type of synthetic health data project and its expected outputs will have an impact on which risk management controls can be dialled up or down. For example, when releasing data to the public at large, we must assume there is zero safety in terms of the 'people' or 'settings'. In such cases, to achieve an overall level of risk management to enable safe sharing, it will be critical to 'dial up' other elements such as the safety of the data, through extremely stringent de-identification techniques.

The Safety Assessment Checklist below reflects the different contexts in which data sharing might take place, while still using the broad 'five safes' concept.

Safety Assessment Checklist

The accountable decision-maker must assess the request in consideration of the following:

Considerations:	
Safe data	Has the data been presented so that it can be clearly understood, appropriately footnoted, with data sources acknowledged? Will the integrity and quality of data be maintained by the Data Requestor? For example, will the data remain segregated from other data held by the requestor, and data lineage (including that the data originated at the Data Provider) maintained? Has re-identification risk been tested? Are there any data fields which could raise 'safe data' concerns, due to their identifiability? Examples could include: • Date of birth • Combinations of demographic fields such as age, gender, postcode, Aboriginality, ethnicity or health status • Attribute data that could lead to re-identification by rendering an individual unique in the dataset, e.g. geolocation data, or longitudinal data • Text-based fields / unstructured data
Safe projects	Has a Privacy Impact Assessment (and any other relevant assessment such as a data linkage assessment) been completed, where applicable? Is the data to be used for an appropriate purpose? Reasons for concern could include if: • There is a Data Owner outside the organisations (e.g. data provided by another organisation) • The data was indirectly collected, generated or inferred • A new metric or indicator needs to be built for this proposal • The synthetic health data request does not have a clear objective or methodology
Safe settings	What controls will be in place to prevent unauthorised access or unauthorised use? For example, the accountable decision-maker must be satisfied that: • The Data Requestor can meet all of the Data Provider organisation's requirements including security requirements

	• The data will be transferred, stored, managed and disposed of securely and appropriately, to prevent unauthorised or accidental access, modification, loss, and damage or copying. Systems should allow for user access to be controlled, monitored and audited. • The system within which the data will be used and stored, and any transfer mechanisms, must be subject to a technical security assessment to ensure it is sufficiently secure in the circumstances. What is considered 'sufficiently secure' should take into account the type and format of data, the level of privacy risk associated with the data, and the potential impacts (both legal and non-legal) in the event of a data breach or misuse. For example, unit record data with a 'more than very low' re-identification risk may need to be stored and accessed via a trusted research environment or secure data enclave. Other types of secure storage options may be suitable where data only has a very low re-identification risk, and other reasonable safeguards are in place. • The data has been labelled according to any relevant data classification policies • Obligations relating to the return or disposal of the data have been agreed • If approved for sharing, there will be a legally binding Data Sharing Agreement (DSA), Data Use Agreement (DUA), or other appropriate written agreement in place between the organisations, and with End Users, with appropriate confidentiality and privacy provisions. (Organisations should seek legal support on questions regarding DSAs and DUAs. See discussion on DSAs and DUAs below.) • Where contracted service providers are involved in the synthetic health data project (for example, a cloud platform provider), the contracting organisation has a service contract in place with appropriate data and privacy protection clauses, and carries out appropriate onboarding and oversight of the contracted service provider's performance. • The data must be held in Australia, and by an organisation that is subject to legal privacy obligations and oversight by a body (e.g. a statutory regulator such as a Privacy Commissioner or Information Commissioner) who can enforce these obligations. As an example, private sector organisations (with an annual turnover of less than $3 million) and South Australian public sector agencies are not subject to legal privacy obligations. Any proposed exception to this requirement must first undergo a legal risk

	assessment, and appropriate clauses in the Agreement must be included.
Safe people	Has the Data Requestor provided a list of people who will be authorised to access / use the data, and those people have been approved? • Consider whether an end user or recipient of the data possesses specialised skills or technology, or has access to relevant data, which may increase the risk of re-identification. How can these risks be controlled?
Safe outputs	How will the Data Requestor use or publish the data later, or outputs from analysis? What guarantee is there that their actions will not lead to re-identification of individuals, or other forms of harm to cohorts or the community?

Data Sharing Agreements and Data Use Agreements

Following the completion of the steps and assessments required under this Framework, and where the accountable decision-maker has approved the synthetic health data request, the participating organisations (i.e. the Data Provider(s) and Data Requestor(s)) will need to enter into a **Data Sharing Agreement (DSA) prior any data being shared**. Entering into a DSA has the benefit of documenting what the organisations agreed to regarding the synthetic health data request and setting out requirements for data handling and security. DSAs can also contain clauses around ongoing assurance and enforcement rights once the data has been shared (e.g. whether the Data Provider has a right to inspect or audit compliance by the Data Requestor).

End Users at the Data Requestor organisation who will access the synthetic health dataset for the approved use case will also be required to complete a **Data Use Agreement (DUA) before they are granted access to the data**. A DUA can be used to notify End Users of their responsibilities and obligations when accessing and using synthetic health datasets.

While a legal expert should be involved in drafting DSAs and DUAs, organisations should consider incorporating the following:

- Identifying the parties involved in the data sharing
- Agreement expiry date
- A description of the synthetic health data to be shared under the agreement, including the nature of the data and the source dataset from which it is drawn. This should include a list of specific data fields (and may be contained in data specification document attached as an appendix)
- Confidentiality clauses

- Requirements that parties will comply with Australian and any relevant state or territory privacy legislation to the extent the shared data includes personal information
- Approved purposes for which the data may be used and shared by the parties
- Details of the sharing mechanism to be used, including frequency
- Details of the security requirements for transfer and storage of data (including any relevant standards or policies)
- Data storage location – expected locations may be the ACT, NSW, Northern Territory, Queensland, Tasmania, Victoria and Western Australia (Western Australia will only be a suitable location after the substantive privacy provisions in the PRIS Act become effective in July 2026, unless the source data originated in Western Australia), unless parties agree on another location (which must be subject to a legal risk assessment first)
- Restrictions on who may access the data (i.e. limited to authorised and identified individuals based on their roles or functions within the organisation)
- Conditions on use, release, or publication of the data
- Restrictions on using synthetic health data, or combining synthetic health data with other data, in a manner that could reasonably re-identify an individual or that generates personal information about an individual or otherwise increases the risk of re-identification
- Prohibitions on attempts to re-identify the data
- Conditions on the use or release of outputs (for example, outputs from data analysis, results, insights, statistics, or other information or data generated from the synthetic health data), including any review or approvals from the participating organisations, and whether the data should be identified as originating from the Data Provider. The Agreement should specify who is responsible for the management, storage and the destruction of the outputs. Outputs should not include any personal information, or other information that could reasonably be used to identify an individual (either alone or in combination with other information or knowledge)
- Requirements around sharing the data with (including providing access to) a third party, including notification and/or approval by the parties.
    - For example: if a third-party service provider (e.g. a cloud platform provider) will have access to the shared data, requiring service contracts between a party to the Agreement and the service provider to first be in place which contain appropriate privacy protections (including restrictions on any use of the data, security obligations, and data breach and incident reporting responsibilities).
    - Also, where a party is required by law to disclose the data to a third party (for example, to comply with a court order), they may need to (where lawful) first notify the other parties of the disclosure.

- Obligations relating to the quality of the data to ensure it is fit-for-purpose. For example, that the data is provided in the agreed format; that it is accessible; and that it meets the agreed description in terms of accuracy, completeness, reliability and currency. Data being shared should be limited to what is necessary to achieve parties' objectives in line with approved purposes. Parties may wish to include a requirement that records be maintained about what data was provided and/or combined for traceability and verification purposes – such as source systems, description, date of extraction, etc. Parties must notify the Data Custodian and/or Data Providers of any data quality issues.
- Obligations relating to the return or disposal of the data, including by what date. A Data Retention Schedule can be useful for setting out retention timeframes, and obligations for data retention and destruction (for example, if certain datasets must be retained for certain timeframes to meet legal requirements).
- Obligations relating to reporting / escalation pathways in relation to any privacy incident, inadvertent inclusion of personal information in the dataset, privacy complaint, access or correction request, or suspected or actual data breach. This includes incidents or breaches involving the model (or aspects of the model) used to generate the synthetic health dataset. The Agreement should specify which party is responsible for managing these events.

A DUA for End Users will not typically need to be as detailed as the overarching DSA, and will be a one-way agreement or acceptance of certain terms and conditions completed by the End User. DUAs may incorporate:

- A statement / declaration agreed to by the End User about accepting certain responsibilities and obligations regarding the data
- Relevant elements from a DSA, such as:
    - Confidentiality clauses
    - Restrictions on accessing and using the data only for an approved purpose and during an approved timeframe
    - Restrictions or conditions on how outputs from analysis can be handled (e.g. releasing or publishing outputs)
    - Compliance with any relevant policies or guidelines when accessing the data, such as security policies that apply to devices and / or relevant ethical guidelines
    - Requirements to escalate an incident to an appropriate designated person
    - Prohibitions on activities which may impact re-identification risks

If the SynD organisations wish to develop a template DSA and DUA, the above suggested content can be removed and the Framework can instead point to these templates.

Further resources

Relevant organisational policies that may need to form part of an organisation's safety assessment can be set out / linked to here

E.g. information security policies; data governance policies; data breach and incident management policies, data release policies, acceptable IT use policies, etc.