Pinnacle Incorporated

# Nabla Copilot: Privacy impact assessment report

13 MAY 2024

# Contents

# Project summary

Nabla Copilot is an AI assistant product designed to ease the workload on medical practitioners developed by Nabla, a French digital health startup. Nabla Copilot captures an audio recording of a medical consultation between a practitioner and a patient and uses a combination of speech-to-text transcription and large language models (LLMs) to produce a structured medical note that can be exported directly into a patient's medical records.

There has been a significant level of interest by general practice in using Nabla Copilot (and similar AI tools) to reduce the amount of time GPs and other practitioners spend on writing up consultations, and it has recently been announced a Nabla Copilot plugin will be integrated into several major practice management systems (PMSs). This privacy impact assessment (PIA) seeks to investigate the potential privacy risks of implementing Nabla Copilot in this form from the perspective of Pinnacle, on behalf of its member practices.

# Scope of the PIA

## Scope

This project covers the implementation of Nabla's "Nabla Copilot" tool within member practices of Pinnacle Midlands Health Network, when used as an integrated part of the practice's practice management system (PMS). This PIA covers:

- what additional information may be collected by Nabla Copilot as part of its use in patient consultations
- how Nabla Copilot may capture, store, and process the information it collects from the patient and practitioner.

Issues considered to be out of scope for this PIA include:

- existing procedures and processes around the storage of patient notes within the PMS.

## The process

This PIA was completed by Pinnacle's Privacy Officer and members of the Data Team, based on public information on the features and technical specifications provided by Nabla as of May 2024.

# Personal information

Nabla Copilot functions as an ambient AI assistant that summarises medical consultations into formatted, medical notes. To do so, it needs to follow the data flow outlined in *Figure 1: Overview of the Nabla Copilot data flow*.

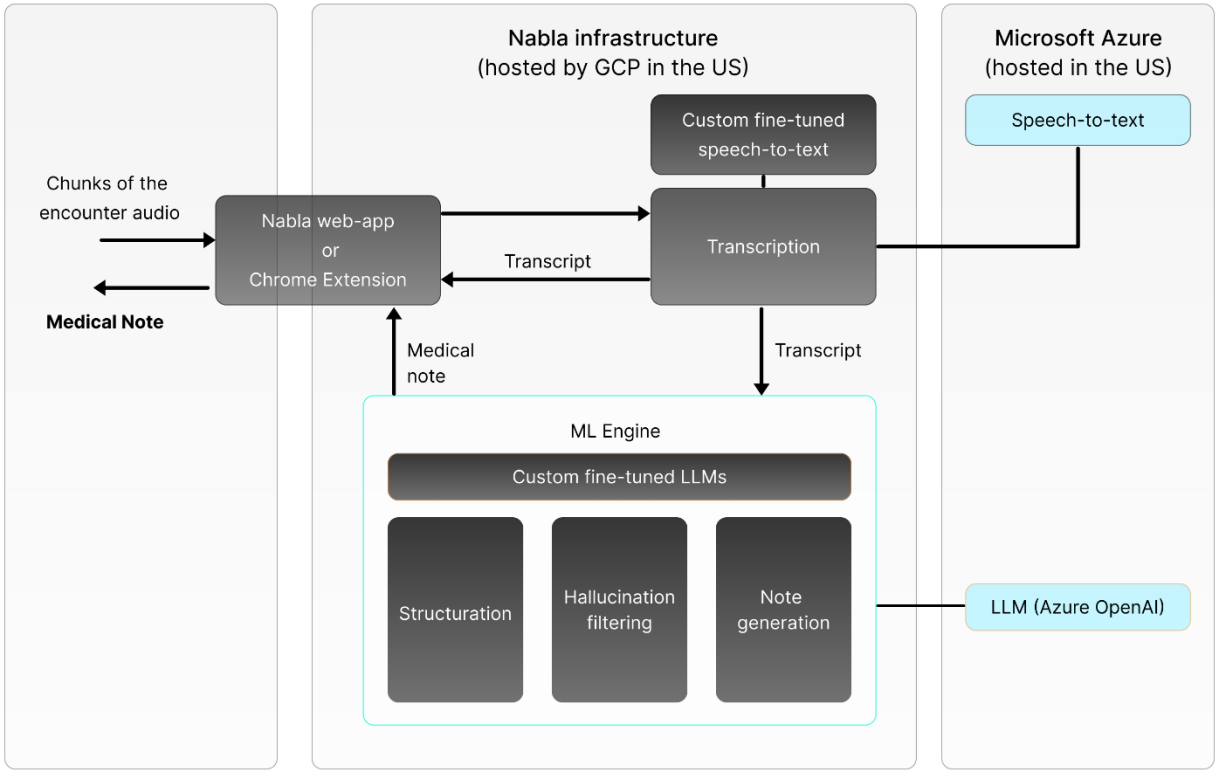


*Figure 1: Overview of the Nabla Copilot data flow [1]*

Before the beginning of the consultation, the practitioner runs Nabla Copilot on a microphone-equipped computer or phone, using either the Nabla web-app, chrome extension, or a plugin integrated directly into the PMS, and clicks “Start encounter” to begin capturing audio.

During this medical consultation, a wide range of personal information may be provided by the patient about themselves, much of which may be particularly sensitive.

As the consultation proceeds, Nabla Copilot uploads portions of the recorded audio to their transcription service, which uses a combination of Microsoft Azure’s commercial speech-to-text API and a speech-to-text model fine-tuned on medical terms to produce a written transcription of the ongoing consultation [2].

Once the consultation is completed and the practitioner clicks “Stop encounter”, the audio capture is stopped. The transcript is then processed by a combination of Microsoft Azure’s commercial LLM and LLMs fine-tuned on medical note generation to produce a structured medical note that summarises the consultation [3].

The consultation transcript and structured medical note are stored in Nabla’s systems temporarily, for a configurable period (by default, 14 days) so that the practitioner can review the note and export it into the patient’s PMS record. Once this period expires, the data is removed from Nabla’s systems;

seven days later after the data is removed from their systems, it is removed from Nabla's backups and patient data is entirely gone from their system.

At this point, data collected by Nabla may be stored in two places: either in the PMS record, or optionally, practitioners may choose to submit a transcription and generated medical note to Nabla for feedback. Since this data is stored permanently in Nabla's systems, this transcription and note are algorithmically de-identified to remove personally identifiable information.

# Privacy assessment

This section outlines the principles set out in the Privacy Act, which form the legal framework Pinnacle needs to consider when deciding to implement Nabla Copilot in practices.

Each row in the following table summarises the key requirements of each of the privacy principles, and outlines how the implementation of Nabla Copilot may impact each principle. Based on this summary, Nabla Copilot has been assessed as being compliant or non-compliant with each principle, with any additional information linked to a numbered risk.

| # | Description of the privacy principle | Summary of personal information involved, use and process to manage | Assessment of compliance |
|---|---|---|---|
| 1 | **Principle 1 - Purpose of the collection of personal information**<br>Only collect personal information if you really need it. | Patient data is collected for the purpose of summarising consultations into medical notes.<br>Unnecessary data may be collected as part of the recording if brought up by either the patient or practitioner during the consultation, but may be edited out of the medical notes by the practitioner after the completion of the consultation. | *Compliant* |
| 2 | **Principle 2 – Source of personal information**<br>Get it directly from the people concerned wherever possible. | Nabla Copilot uses only patient data collected directly from the patient during consultation. | *Compliant* |
| 3 | **Principle 3 – Collection of information from subject**<br>Tell them what information you are collecting, what you're going to do with it, whether it's voluntary, and the consequences if they don't provide it. | To comply with this principle, patients must be aware of and have consented to the recording and processing of their medical consultation by a third-party.<br>This could include, but is not limited to, explicitly informing the patient at the time of consultation, informative signage, or any other appropriate mechanisms of obtaining patient consent. | *Potentially non-compliant*<br>*(R-001)* |

| # | Description of the privacy principle | Summary of personal information involved, use and process to manage | Assessment of compliance |
|---|---|---|---|
| 4 | **Principle 4 – Manner of collection of personal information**<br>Be fair and not overly intrusive in how you collect the information. | Nabla Copilot is a non-intrusive ambient AI assistant, which should result in no change to the way in which information is collected during a consultation. | *Compliant* |
| 5 | **Principle 5 – Storage and security of personal information**<br>Take care of it once you've got it and protect it against loss, unauthorised access, use, modification or disclosure and other misuse. | Nabla lists several safeguards that they have taken in order to secure the limited data that they hold temporarily, including "employee training, third-party audits and penetration testing, strict management of roles and permissions, strong authentication processes, encryption at rest and in transit, continuous vulnerability scanning, logging, monitoring, alerting and much more" [1].<br>Nabla state they are SOC 2 Type 2 compliant and ISO 27001 certified. | *Compliant* |
| 6 | **Principle 6 – Access to personal information**<br>People can see their personal information if they want to. | Implementation of Nabla Copilot would not result in any change to this principle; existing procedures around access to a patient's own notes and medical record would apply, which are out of scope of this PIA. | *Compliant* |
| 7 | **Principle 7 – Correction of personal information**<br>They can correct it if it's wrong or have a statement of correction attached. | Implementation of Nabla Copilot would not result in any change to this principle; existing procedures around corrections to a patient's own notes and medical record would apply, which are out of scope of this PIA. | *Compliant* |
| 8 | **Principle 8 – Accuracy etc. of personal information to be checked before use**<br>Make sure personal information is correct, relevant and up to date before you use it. | The AI models used by Nabla Copilot to summarise a consultation may have biases that affect the accuracy or comprehensiveness of the notes. Inaccurate or incomplete notes could lead to clinical errors or misjudgements, potentially impacting patient care.<br>To minimise this risk, practitioners should thoroughly review generated medical notes to ensure the information in them is correct, complete, and unbiased. | *Potentially non-compliant*<br>*(R-002)* |
| 9 | **Principle 9 – Not to keep personal information for longer than necessary**<br>Get rid of it once you're done with it. | Nabla states they keep transcripts and structured medical notes temporarily, storing them on their systems for a period of 14 days for practitioners to review and export.<br>Storage of these structured medical notes within the patient's PMS record is out of scope for this PIA. | *Compliant* |

| # | Description of the privacy principle | Summary of personal information involved, use and process to manage | Assessment of compliance |
|---|---|---|---|
| 10 | **Principle 10 – Limits on use of personal information**<br>Use it for the purpose you collected it for, unless one of the exceptions applies. | Nabla Copilot uses patient information provided during the consultation for generating medical notes, which falls under the original purpose of collecting that information.<br>Nabla Copilot allows practitioners to submit de-identified transcriptions and generated medical notes for feedback and model training purposes. Although this data is de-identified, there is the risk of re-identification if anonymisation isn't thorough or if residual data patterns persist. This risk can be mitigated by opting out of the ability to send feedback to Nabla. | *Potentially non-compliant*<br>*(R-003)* |
| 11 | **Principle 11 – Limits on disclosure of personal information**<br>Only disclose it if you've got a good reason, unless one of the exceptions applies. | Implementing Nabla Copilot would result in disclosing patient information to Nabla for the purpose of them transcribing the consultation and generating medical notes, which is directly related to the purpose for which the information was obtained (11 (1)(a)). | *Compliant* |
| 12 | **Principle 12 – Disclosure of personal information outside New Zealand**<br>Only send it to someone overseas if the information will be adequately protected. | While Nabla Copilot and the tools it uses to process patient consultations are hosted on global cloud platforms (Google Cloud Platform and Microsoft Azure), the Office of the Privacy Commissioner has stated that using "[u]sing offshore technology providers to store or process your data is not treated as a disclosure under IPP12, so long as they are not using that information for their own purposes" [4]. Opting out of sending feedback, as noted under principle 10, would remove this potential risk.<br>Nabla has indicated their data processing is done in "strict compliance with both HIPAA and GDPR", and that they have agreements to opt out of data retention for all services used to process it. | *Potentially non-compliant*<br>*(R-003)* |
| 13 | **Principle 13 – Unique identifiers**<br>Only assign unique identifiers where permitted. | Implementing Nabla Copilot would not result in the creation of new unique identifiers for patients. | *Compliant* |

# Risk management

Three primary sources of potential risk were identified above.

- Patient consent.
- Accuracy of generated medical notes.
- Sharing of patient information for model training purposes.

These risks are further detailed below.

## R-001: Patient consent

Information privacy principle 3 states in part that if an agency collects personal information about a person, the agency must take any reasonable steps to ensure that the individual is aware of what information is collected, the purpose of that information, and the consequences if that information isn't provided.

As AI transcription of Nabla Copilot's nature is a new and emerging technology, it's especially important that patients are fully informed of its use and how it may impact on their privacy. Failure to do so may result in patient dissatisfaction, ethical issues, or privacy complaints.

Several methods could be used to ensure that patients are aware that Nabla Copilot is being used to record and summarise their consult.

- Seeking explicit consent before activating Nabla Copilot. Nabla suggests using the following copy in jurisdictions that require patient consent [5]:
  *"I'm using a software tool called Nabla to make it easier to write down my medical notes. It helps me spend more time focused on your care, and less time on the computer. This tool will listen to our conversation and convert it into a summary, which I will review and edit for accuracy. The tool will have access to the audio of our conversation during our visit, but it will not be saved and will not be used further by the tool. Your medical information will stay private and only shared with those you allow.*
  *Do you consent to the use of this tool during our visit and future visits?"*
- Signage informing the patient that Nabla Copilot or similar tools may be used during a consultation, as well providing a way to find out more information.

## R-002: Accuracy of generated medical notes

Information privacy principle 8 states that agencies that hold an individual's personal information must not use or disclose that information with taking reasonable steps to ensure that the information is accurate, up to date, complete, relevant, and not misleading.

It is possible for generative AI tools, including summarisation tools like Nabla Copilot, to produce confident errors of fact and logic ("hallucinations"). Additionally, depending on the training data used, AI model may have inherent biases that could affect the accuracy or comprehensiveness of the notes. In the case of Nabla Copilot, Nabla is a French company not primarily focused on the New Zealand market, and so the training data used may not be relevant to New Zealand; as such, any potential biases in the model may disproportionately affect groups of people more commonly found in New Zealand (i.e., Māori or Pacific peoples).

Any issues with the generated medical notes caused by hallucinations or inherent biases, if not identified and corrected, could result in clinical errors or misjudgements, potentially impacting patient care. As such, steps should be taken when using these tools to ensure that the practitioner

carefully evaluates the generated data to ensure correctness and accuracy before entering it into the patient's record.

## R-003: Sharing of patient information for model training purposes

Information privacy principle 10 states that agencies need to limit use of the information they collect to the stated purposes for collection.

Nabla Copilot includes an optional feature for practitioners to share de-identified transcriptions and generated medical notes with Nabla for the purposes of providing feedback and model training. While this data is de-identified using an algorithm which "masks the 18 HIPAA identifiers" [1], it's possible that this may still result in patient information being used.

For example, identifiable data may slip through the de-identification algorithm, or a combination of data points that may not be identifying on their own could be used to identify a patient.

Additionally, information privacy principle 12 states that personal information must not be disclosed overseas unless certain conditions are met. While principle 12 includes an exception for using offshore technology providers to store or process data if they are not using that information for their own purposes, providing patient data to Nabla for the purposes of improving their model would not be covered by this exception.

As such, to avoid the unintentional use of patient data which breaches principles 10 and 12 in this way, practitioners should choose to not provide de-identified data to Nabla for feedback.

# Recommendations to minimise impact on privacy

Summarise the recommendations to minimise the impact on privacy based on your risk assessment.

| Ref | Recommendation | Agreed Y/N |
|---|---|---|
| R-001 | Ensure that patients are made aware of, and consent to, the use of Nabla Copilot in their consultations | |
| R-002 | Ensure that practitioners review medical notes generated by Nabla Copilot for accuracy and potential bias | |
| R-003 | Opt out of providing de-identified feedback to Nabla Copilot | |

# Action plan

This section of the report should describe what actions are being taken (whether short or long term) and how they'll be monitored. There may also be links to other processes in the organisation. For example, a proposed action might relate to security controls (such as restricting access to a system). This will then link in with security processes in the organisation.

Reporting on the outcome of the mitigation may be necessary. If the PIA is being performed as part of a project, then the project is likely to require some reporting on their implementation as part of governance arrangements. Once the project is completed, any on-going privacy monitoring should be incorporated into normal business operations.

In the case of a particularly long or complex programme of work, the PIA may need to be reviewed a number of times to ensure that it continues to be relevant. This section should describe how this will be achieved.

| Ref | Agreed action | Who is responsible | Completion date |
|---|---|---|---|
| R-001 | | | |
| | | | |

# References


[1] C. Baudelaire, “All you need to know about Nabla's privacy and security features,” Nabla, 10 March 2023. [Online]. Available: https://www.nabla.com/blog/privacy-security/. [Accessed 23 April 2024].

[2] R. Stern and S. Humeau, “How we built the best speech-to-text engine for medical encounters,” 3 July 2023. [Online]. Available: https://www.nabla.com/blog/speech-to-text/. [Accessed 23 April 2024].

[3] R. Stern, M. de Vriendt, G. Retourné and S. Humeau, “Evaluating medical note generation,” 28 November 2023. [Online]. Available: https://www.nabla.com/blog/evaluating-note-generation/. [Accessed 23 April 2024].

[4] Office of the Privacy Commissioner, “Artificial intelligence and the Information Privacy Principles,” September 2023. [Online]. Available: https://www.privacy.org.nz/assets/New-order/Resources-/Publications/Guidance-resources/AI-Guidance-Resources-/AI-and-the-Information-Privacy-Principles.pdf. [Accessed 23 April 2024].

[5] Nabla, “Nabla · Patient Consent,” [Online]. Available: https://www.nabla.com/copilot-patient-consent/. [Accessed 23 April 2024].