Privacy Impact Assessment: Scribeberry Ltd. (AI Medical Scribe Service)

Cover Letter

May 8, 2025
Office of the Information and Privacy Commissioner (OIPC) of Alberta
Re: **Privacy Impact Assessment for Scribeberry Ltd.**

Dear Commissioner,

Please find enclosed the comprehensive Privacy Impact Assessment (PIA) for Scribeberry Ltd. Scribeberry is a physician-led provider of an AI medical scribing and documentation service, now operating across Canada and the USA. This PIA has been prepared in accordance with Alberta's Health Information Act (HIA) and follows the Alberta Health **"Completing a Privacy Impact Assessment – Annotated Template (Version 1.1, April 2023)"** structure. It details how Scribeberry, as a custodian of health information, ensures the privacy and security of personal health information (PHI) in its custody. We confirm that Scribeberry's designated Privacy Officer is Dr. Zaahir Moloo. The PIA covers all Scribeberry deployments Canada-wide (including Alberta, Ontario, etc.) and in the United States. Key privacy safeguards – including end-to-end encryption, data residency in Canada for Canadian users, strict access controls, and third-party agreements with Microsoft Azure, Google Cloud, and OpenAI – are described. We have also included an information flow diagram, a legal authority and purpose table, and summaries of relevant policies, contracts, notices, and consent processes. Scribeberry does not have access to identifiable patient data; PHI is encrypted and visible only to the end-user clinician. The attached assessment demonstrates Scribeberry's compliance with the HIA and other applicable privacy legislation, and outlines our ongoing risk mitigation, monitoring, and compliance activities.

We trust that this PIA addresses all requirements. Scribeberry Ltd. is committed to protecting patient privacy and will implement any recommendations from your office. Thank you for your review and consideration.

Sincerely,

Dr. Zaahir Moloo
Privacy Officer & Co-Founder, Scribeberry Ltd.
Email: zaahir@ualberta.ca, Phone: 7806601092

Cover Page

Official Name of System: Scribeberry – AI Medical Scribe and Documentation Service
Custodian Legal Name: Scribeberry Ltd. (Dr. Zaahir Moloo, MD)
PIA Prepared By: Dr. Zaahir Moloo – Privacy Officer & Custodian (with input from privacy consultant)
Delegated HIA Compliance Officer: Dr. Zaahir Moloo – Privacy Officer (Contact: hello@scribeberry.com, 7806601092)
PIA Submission Date: May 8, 2025
Expected Implementation Date: System in production (launched 2023; PIA previously submitted in 2023 under Hello Mental Health (Amended PIA)
Related OIPC Submission: *Hello Mental Health*

Section A: Project Summary

Overview: Scribeberry Ltd. offers an **AI-powered medical scribing and documentation tool** designed for health care professionals across Canada and the US. The system captures patient encounter information (e.g. clinician-patient conversations) and automatically generates clinical documentation (medical notes, summaries, charts) in real time. Scribeberry's objective is to **reduce administrative burden** on clinicians by streamlining note-taking and improving accuracy and completeness of health records. By leveraging advanced speech recognition and language models, Scribeberry allows clinicians to focus on patient care while documentation is handled in the background. The ultimate purpose is to improve workflow efficiency, enhance record quality, and facilitate better communication in health care.

Need for Health Information: To function, the Scribeberry system **collects, uses, and discloses individually identifying health information** only as part of documenting patient care. The service inherently must handle personal health details spoken or entered during a medical encounter – for example, patient demographics, symptoms, medical history, diagnoses, and treatment plans. This information is necessary for Scribeberry to create accurate clinical notes that will eventually form part of the patient's health record. Scribeberry **only collects the minimum information required** to fulfill this documentation purpose, in alignment with HIA's data minimization principle. Health information is processed *only* to assist in providing health services to the individual patient and clinician (e.g. to produce a consultation note for their chart). Without using actual patient information, the system could not generate meaningful documentation; thus, collection and use of PHI are essential for Scribeberry's objectives. However, data is sanitized, anonymized, and used with great care so as to not be used for any third-party purposes.

Participants and Roles: Scribeberry Ltd. itself is the custodian operating the system (for this PIA's scope). **Dr. Zaahir Moloo** serves as Scribeberry's Privacy Officer and lead for HIA compliance, given his role as a physician co-founder. Health professionals (clinicians) across

various clinics are the authorized end-users of Scribeberry; they act as affiliates of Scribeberry when using the system under this model. These clinicians input or capture patient data via the Scribeberry application during patient visits. Scribeberry is also supported by a small internal team (e.g. a technical lead, engineers) who manage the platform, but notably **no Scribeberry staff have routine access to patients' identifiable health information** in the system. Additionally, Scribeberry engages **third-party service providers** to support its technology: Microsoft Azure and Google Cloud (for cloud hosting and storage), and OpenAI (for AI language processing). Each third-party's role is strictly limited via contract to processing or storing data on Scribeberry's behalf (see Section C.6).

System Operations and Architecture: Health information will be accessed and stored in a **secure cloud-based environment**. Scribeberry utilizes **Microsoft Azure and Google Cloud data centers in Canada** to store and process Canadian users' data, ensuring Canadian health information remains within Canadian jurisdiction. For example, a clinician in Alberta or Ontario using Scribeberry will have their data processed and stored on Azure/GCP servers located in Toronto, Montreal, or (for certain components) Edmonton, rather than in the US. If a user in the USA uses Scribeberry, data is routed to US-based servers to comply with data residency preferences for HIPAA. The system primarily operates as a cloud application that clinicians access via secure web or mobile app from their clinic or device.

Servers: Key components include an on-premises (self-hosted) transcription server (managed by Scribeberry in a secure environment) to convert speech to text, and connections to OpenAI's cloud API for natural language processing. The physical server infrastructure for transcription is maintained either in Scribeberry's controlled data center in Canada (potentially in Edmonton) or within its secured cloud tenant, adding an extra layer of control for audio data. All servers (whether on-prem or cloud) have robust security safeguards and encryption.

Access Methods: Clinicians and their authorized staff will access Scribeberry **within clinic offices or remotely via secure devices**. In typical use, a clinician uses Scribeberry on a tablet or computer in the exam room (or a smartphone for dictation). Access is through authenticated user accounts – only the clinician and any designated delegates can log in to view the notes. Mobile access is allowed (e.g. a doctor reviewing notes from home on a secured device) to support flexibility. Scribeberry enforces secure login and session controls to protect data when accessed outside the clinic. No patients or general public users have direct access to the system; however, clinicians may share the generated documentation with patients through normal record release processes outside the application.

Vendor Involvement: Scribeberry (the vendor) is responsible for deploying and maintaining the system, including storing the encrypted data in the cloud. **Azure and Google Cloud** host Scribeberry's application and database – but **Scribeberry retains custody/control of the data** through these services via contractual agreements. Scribeberry's contracts with these providers ensure that they act as information managers and do not use the health information for any purposes other than storage/processing on Scribeberry's instructions. **OpenAI** is engaged for its AI language model service; Scribeberry sends transcribed text to OpenAI's secure endpoint to generate the structured clinical note. OpenAI returns the result to Scribeberry. Importantly,

Scribeberry has executed **Business Associate Agreements (BAAs)/data protection addenda** with OpenAI and all cloud providers to impose strict privacy obligations. Scribeberry's team manages the health information within the system (e.g. ensuring it flows correctly to Azure, OpenAI, etc.), but **does not view or intervene in the content of any patient records**. No third-party (including OpenAI) has independent rights to retain or use the data beyond providing the service. There is **no disclosure of patient information to other health providers or systems** through Scribeberry itself – it is an internal documentation tool. (Clinicians may later transfer the finished notes into their official electronic medical record (EMR) or share with patients or referring doctors as needed, but those actions are outside Scribeberry's direct scope.)

HIA Compliance and Safeguards: In summary, this project involves a new system for handling health information, triggering the need for a PIA under HIA Section 64. Scribeberry recognizes its duty under HIA Section 60 to protect health information with administrative, technical, and physical safeguards. This PIA explains **what the Scribeberry system does and how it operates**, and **how we will meet our obligations under the HIA during and after implementation**. Key privacy features include: end-to-end encryption of data in transit and at rest, role-based access control limiting data visibility to the end-user only, data residency controls, and thorough organizational privacy practices (detailed in Section B). By implementing privacy by design, Scribeberry ensures compliance with HIA Sections 60, 61, and 63 (protection, accuracy, and policies) throughout the system's deployment. The next sections provide detailed analysis of Scribeberry's privacy management program, specific privacy impacts, risk mitigations, and compliance measures.

Section B: Organizational Privacy Management

Scribeberry Ltd. has established a comprehensive privacy management program to fulfill its legislative requirements under the HIA. As a **community-based custodian**, Scribeberry's approach to privacy is grounded in strong organizational structure, clear policies and procedures, staff training, and robust incident response processes. The following subsections describe each aspect of Scribeberry's privacy management framework.

1. Management Structure

Custodian and Privacy Officer: Scribeberry is a physician-led organization. **Dr. Zaahir Moloo** (MD), co-founder of Scribeberry, serves as the designated **Privacy Officer** and is the individual responsible for overall compliance. Dr. Moloo's dual background in medicine and information security uniquely qualifies him to oversee Scribeberry's privacy program. In practice, this means Dr. Moloo makes final decisions on privacy-related matters, such as approving policies, handling any access requests or breach investigations, and ensuring the PIA is followed. He is the primary point of contact for any privacy inquiries (both internally and from users or patients).

Scribeberry has formally documented the **roles and responsibilities** of the Privacy Officer to ensure clarity in duties.

Organizational Roles: Scribeberry is a small organization, but key roles are defined. **Amaan Rattansi**, co-founder and CTO, shares responsibility for implementing security and privacy controls alongside Dr. Moloo. Essentially, Dr. Moloo and Mr. Rattansi jointly **develop and implement Scribeberry's information security and privacy policies**. Mr. Rattansi acts as an alternate contact for privacy/security matters in Dr. Moloo's absence, ensuring continuity. Other Scribeberry staff (e.g. software engineers) are considered **affiliates under the HIA**, and each has clearly defined access limits and obligations. For instance, technical staff may have access to system infrastructure but **no access to decrypted PHI**. Scribeberry's structure is deliberately designed such that **only the authorized end-user clinicians access patient data**, and **employees do not** (employees work with de-identified or encrypted data whenever possible). This separation of duties mitigates insider risk. An organizational chart (see Attachment E) illustrates the reporting relationships – with the Privacy Officer at the top, reporting directly to Scribeberry's executive leadership (which, in this startup, is also Dr. Moloo).

Privacy Decision-Making: Decisions about health information privacy at Scribeberry are made in a collaborative but well-governed manner. Dr. Moloo leads a **Privacy and Security Committee** (informal, given the company's size, but consisting of key leadership and our external privacy advisor) that meets regularly (at least bi-weekly) to discuss any privacy issues, review risk assessments, and address new threats. This ensures ongoing attention to privacy. Significant decisions (e.g. adopting a new third-party service or changing data handling practices) are vetted for privacy impact and must be approved by Dr. Moloo. Scribeberry consults external privacy experts as needed (e.g. our auditor, Ms. Ingrid Ruys, for guidance on complex HIA questions). Privacy issues raised by users or regulators are escalated to Dr. Moloo immediately. All staff are empowered to raise privacy concerns to management without fear of reprisal, reflecting a culture where privacy is a priority.

In summary, Scribeberry's management structure ensures **accountability**: a specific individual (Privacy Officer) is accountable for compliance, and all affiliates are aware of their responsibilities to protect health information. This fulfills HIA Section 62(2) requirements by clearly assigning responsibility for HIA compliance within the organization. Any gaps identified in privacy roles or decision-making processes are promptly addressed by management as part of our continuous improvement.

2. Policy Management

Scribeberry has developed a suite of **privacy and security policies** to support HIA implementation (fulfilling HIA Section 63). These policies provide the rules and procedures that all Scribeberry affiliates must follow when handling health information.

Policy Development: Policies were initially drafted during Scribeberry's startup phase, leveraging guidelines from the OIPC and best practices from Alberta Health. Dr. Moloo and Mr. Rattansi authored the policies, with review by our privacy auditor to ensure completeness. The

policies have been approved by Scribeberry's leadership (Dr. Moloo in his custodian role) and are considered official company policy. Where possible, we utilized templates or guidance from authoritative sources – for example, we referenced the OIPC's "PIA Requirements" guide and the Alberta Medical Association's sample policies, tailoring them to Scribeberry's context.

Key Policies: The following core policies are in place (a full index is attached in Section E):

- **Privacy Policy & Confidentiality Policy:** Outlines Scribeberry's commitment to patient privacy, rules for confidentiality, and expectations for all staff regarding PHI handling. It codifies the "need-to-know" and "least amount necessary" principles for collection, use, and disclosure.

- **Information Security Policy:** Details technical safeguards (encryption standards, password requirements, access controls, network security, etc.). This includes sub-policies on data encryption, secure development, and device security.

- **Access Control Policy:** Defines role-based access, ensuring each affiliate's access to systems or data is limited to what they require for their job. It includes provisions for unique user IDs, strong authentication, and immediate revocation of access when an affiliate leaves or changes role. The independent Security Assessment confirmed our access controls are robust and above average.

- **Records Management Policy:** (Discussed in Section B.3) Covers data retention and secure disposal.

- **Incident Response Policy:** (Discussed in Section B.4) Lays out how to respond to and report privacy breaches.

- **Training and Awareness Policy:** (Section B.3) Describes the training program and requirements for staff and contractors.

- **Sanctions Policy:** Specifies disciplinary measures for any breach of privacy or security rules by an affiliate. (It is integrated with our HR policies to enforce consequences for non-compliance, as required by HIA Section 60(1)(c) "duty to establish sanctions".)

These policies are living documents. Scribeberry conducts a **regular review (at least annually)** of all privacy-related policies to ensure they remain effective, relevant, and compliant with any changes in law or practice. Our internal review cycle is aligned with our annual Security Risk Assessment updates. For instance, we update the encryption policy to reflect the latest TLS standards. We also update policies sooner if there is a significant operational change (e.g. introducing a new feature or a new third-party service). Revised policies are approved by the Privacy Officer and communicated to all staff.

Communication and Enforcement: All Scribeberry affiliates are made aware of these policies. Upon onboarding, each employee/contractor receives the policy manual and must sign an acknowledgment of understanding and an oath of confidentiality. Policies are easily accessible via a secure internal portal provided through Delve (similar to Vanta/Drata). We maintain a **central repository/dashboard for policy documents** to ensure staff always have the latest versions. When a policy is updated, an announcement is sent and staff are required to confirm they've read the changes. Scribeberry promotes a culture of compliance – management regularly highlights key policy requirements in team meetings, and we include policy knowledge checks in training refreshers. To **ensure affiliates follow the policies**, Scribeberry employs monitoring (audit logs, periodic compliance audits – see Section D) and any deviations are addressed. If an affiliate were to violate a policy (e.g. an unauthorized attempt to access data), the Sanctions Policy would be invoked – consequences range from retraining and warnings up to termination of contract, depending on severity.

In summary, Scribeberry's policy management process ensures we have a **complete set of up-to-date privacy policies** that are effectively communicated and enforced. This provides the necessary administrative framework to comply with HIA Section 63 and to guide all staff in protecting health information.

3. Records Management

Proper **records management** is critical to privacy, encompassing how Scribeberry handles the retention, storage, and disposal of health records and related information. Although Scribeberry itself is not a primary health records repository (it serves as a transient drafting tool), we have defined practices to manage any health information we do hold in compliance with HIA and other regulations.

Retention of Health Information: Scribeberry's system is designed to **minimize long-term retention** of identifiable health information. PHI is retained only as long as needed to serve the user (clinician) in preparing and transferring the clinical note. In practice, the **transcribed notes are stored temporarily** in the user's Scribeberry account (cloud storage) so the user can review and edit them, and possibly access them from multiple devices. Once the clinician has finalized the note and copied it into the official medical record (e.g. EMR) or if they decide the note is no longer needed, they can delete it from Scribeberry. **When the user deletes a note, it is permanently deleted from Scribeberry's encrypted storage**. Scribeberry's policy is that **no personal health information will be retained on our system beyond its useful purpose** of aiding documentation.

To enforce minimal retention, Scribeberry implements automated checks. If any PHI data were to persist in the system beyond a certain period of inactivity, we would purge it. For example, we may configure the system to **auto-delete any notes that remain in the system after a defined period (e.g. 30 or 60 days)** unless the user actively chooses to retain them longer. (This ensures "stale" data does not accumulate on our servers.) Additionally, **no audio recordings are stored at all** – the system performs real-time transcription and does not keep audio files, so

there is no audio data to manage or dispose of. This significantly reduces the volume of sensitive data retained.

Storage and Backup: While PHI is stored, it resides in **encrypted form in Scribeberry's cloud database** (hosted in Azure/Google Cloud Canadian data centers). The encryption is strong (AES-256) and keys are managed such that **Scribeberry personnel cannot decrypt the content** (only the application can present it to the authenticated end-user). This approach effectively keeps the PHI "locked away" even during its short retention. We do maintain backups of our databases for disaster recovery, which means PHI could reside in encrypted backups. Those backups are stored securely in the same Canadian regions and are subject to the same retention limits – backup archives containing PHI are purged on a rolling schedule consistent with our primary retention (and in any case, are encrypted). Scribeberry's **Contingency Plan** (see Attachment E) covers how we ensure continuity of operations (including data backup and restore procedures) without compromising privacy.

Disposition (Secure Destruction): When PHI data is deleted by the user or as part of our retention policy, Scribeberry permanently destroys it. Deletion processes include overwriting or cryptographic wiping of the data in our databases so that it cannot be recovered. For example, when a user deletes a note in the app, the corresponding encrypted record is deleted from the database and the encryption keys are also scrubbed. In the unlikely event that any PHI is stored in temporary logs or caches, we have scripts in place to routinely clear those. If Scribeberry were to ever cease operations or a client clinic leaves the service, we have a procedure to export any remaining data to the client (if needed) and then securely purge all client PHI from our systems.

Physical Records: Scribeberry operates almost entirely electronically. We do not print or create physical copies of any patient information in our custody. Thus, physical records management (like file room storage) is not applicable.

Through these practices, Scribeberry ensures compliance with any applicable **records retention requirements** while also adhering to HIA's principle of holding information only as long as necessary (aligning with Section 60 obligations to safeguard info, which includes proper disposal). Because the authoritative medical record remains with the health provider (outside Scribeberry), Scribeberry's aim is to **avoid becoming a secondary repository of health records**. By promptly destroying data after use, we reduce privacy risk and the burden of long-term retention compliance. We document all deletion/destruction events in an audit log, so we have evidence of when data was removed. Our approach to records management has been clearly communicated to our users (clinicians), so they understand that Scribeberry is not retaining their patients' data indefinitely.

4. Training and Awareness

Scribeberry ensures that all affiliates (employees, contractors, and any others with access to our systems) are **properly trained and aware of their privacy obligations**. We recognize that

training is a key safeguard under HIA to prevent accidental breaches and to ensure consistent practices.

Training Program: Upon hire, every Scribeberry workforce member undergoes **privacy and security orientation training**. This training covers the fundamentals of the Health Information Act (HIA), including custodian duties, definitions of health information, and the importance of need-to-know access. It also reviews our internal privacy policies and procedures in detail, and the specific responsibilities of the individual's role. For example, developers learn about data handling procedures in coding, and support staff learn about appropriate handling of any user queries that might involve PHI. We have developed our own training content tailored to Scribeberry's technology (with reference to external resources like OIPC webinars for general HIA concepts). The initial orientation is completed **before the employee is granted access to any health information systems**.

We supplement formal training with **awareness activities**: all staff must complete an annual privacy refresher course online, and we share periodic security bulletins (e.g. reminders about phishing or locking screens). We also conduct scenario-based discussions in team meetings (e.g. "What would you do if you found PHI in a log file?") to keep awareness high.

Training Frequency and Updates: Training is not a one-time event. Scribeberry **offers training on a continuous basis**:

- New employees: orientation within the first week (and no system access until completed).
- Ongoing: annual refresher training is mandatory for all staff.
- Ad hoc: if a significant policy change or new risk is identified, we issue an interim training update. For instance, when mandatory breach reporting laws were updated, we held a special training session on how to recognize and report a breach.
- Developers receive specialized security training regularly (including how to code securely to protect PHI, and recognizing privacy-by-design principles).

We review and update our training materials annually alongside policy updates. This ensures new regulatory requirements or lessons learned from incidents are incorporated. Training content is approved by the Privacy Officer.

Documentation of Training: Scribeberry **keeps records of all training completed**. Each participant signs or electronically acknowledges completion of each module. We maintain a training log that records the date, participant, and content of each training session. This log is reviewed by management to ensure everyone is up-to-date. If someone misses a training (e.g. annual refresher), management is alerted and that person's system access might be temporarily suspended pending completion – underscoring how seriously we take training compliance.

Affiliate Agreements: All affiliates also sign a **Confidentiality and Non-Disclosure Agreement (NDA)** as a condition of working with Scribeberry. This agreement reiterates that any collection, use, or disclosure of health information by the affiliate must comply with our policies and the HIA. The NDA further specifies that unauthorized access or sharing of PHI is forbidden and grounds for discipline. By having this in place, we create a clear, enforceable expectation of privacy conduct.

Remedies for Non-Compliance: We have established **sanctions** if an affiliate fails to follow privacy procedures. For minor breaches of policy (no actual breach of PHI), the individual may receive remedial training and a warning. For serious violations (for example, intentionally accessing data they shouldn't), disciplinary action up to termination will occur, consistent with our Sanctions Policy and contractual terms. Thus far, we have not encountered any internal violations, but these measures are in place as a deterrent.

In summary, Scribeberry's training and awareness program ensures that **every person handling health information on our behalf is knowledgeable and vigilant** about privacy. Because under HIA Section 62(2) any action by an affiliate is considered an action by the custodian, we take great care to train and supervise our affiliates. This comprehensive training regime significantly reduces the likelihood of human error leading to a privacy incident, and it fosters a privacy-conscious culture within Scribeberry.

5. Incident Response (Privacy Breach Management)

Despite robust safeguards, Scribeberry acknowledges the possibility of privacy breaches or security incidents and has a **Privacy Incident Response Plan** in place to meet legislative requirements (including HIA's mandatory breach reporting provisions). Our goal is to respond swiftly and effectively to any incident, mitigating harm and preventing recurrence.

Breach Response Policy: Scribeberry's Incident Response Policy outlines the procedures to follow when a potential privacy breach is detected. It defines what constitutes a breach (e.g. unauthorized access, loss or theft of devices containing health information, malware attack resulting in exposure of data, etc.) and a step-by-step action plan:

1. **Identification and Containment:** Any staff who discovers or suspects a breach must immediately notify the Privacy Officer (Dr. Moloo) and our Security Lead (Mr. Rattansi). The first priority is to contain the incident – e.g. disconnecting affected systems from the network, revoking compromised credentials, or otherwise stopping any ongoing unauthorized access. Because our system is cloud-based, containment might involve revoking keys or tokens, and working with our cloud providers' security teams if needed

2. **Assessment:** The Privacy Officer will lead an investigation to determine the scope of the breach, what information was involved, which individuals are affected, and the risk of harm to those individuals. We document all findings. Risk assessment includes evaluating sensitivity of the data and whether it could be misused (e.g. identity theft,

embarrassment, etc.).

3. **Notification:** In compliance with HIA's **mandatory breach reporting (Section 60.1)**, if the breach creates a risk of harm or embarrassment to the individual, we will notify:

    - The **Minister of Health** (Alberta Health),
    - The **Information and Privacy Commissioner of Alberta** (as required by HIA and via the OIPC breach reporting form),
    - The **affected individual(s)** (patient(s) whose information was involved).

4. Notifications will be done **without unreasonable delay** (target within 7 days of confirming breach details, or sooner as required by regulation). We'll use the OIPC's Breach Report Form and follow the "Key Steps in Responding to Privacy Breaches" guidance. Affected individuals will be contacted via their preferred method (phone and/or in writing), including information about what happened, what information was affected, steps we've taken, and advice on what they can do (if applicable).

5. **Remediation:** We will take all necessary steps to remediate the situation. That may include recovering lost data, strengthening technical controls to prevent a similar attack, or providing credit monitoring if identity info was breached. We involve our technical team to patch any system vulnerabilities immediately. Our incident response plan assigns specific tasks – e.g. the IT lead ensures any compromised server is rebuilt and secured, while the Privacy Officer handles communication and documentation.

6. **Documentation and Reporting:** Every incident (even small ones) is documented in our incident log. If the incident is reportable under HIA (risk of harm), we file an official report with OIPC as described. We also would update this PIA or our policies if the incident revealed any gaps. In the case of any breach, after immediate actions, Scribeberry's team does a **post-incident review** to learn from it. The findings are used to improve our safeguards and response plan.

Vendor Incidents: Since we rely on third-party processors (Azure, OpenAI, etc.), our contracts stipulate that those parties must **notify us** immediately if they detect any incident involving Scribeberry's data. For example, if Azure suffered a data center breach affecting our servers, or if OpenAI identified any unauthorized access to data we submitted, they are obligated to inform Scribeberry. We would then treat it like an incident on our side – containing (perhaps by suspending use of that service), notifying as required, etc. Vendor responsibilities in incident scenarios are clearly defined in our Information Manager Agreements (e.g. requiring cooperation in investigations and notifiable breach clauses).

Practice and Readiness: Scribeberry's team **drills** the incident response process periodically. We hold table-top exercises simulating a breach scenario so that everyone knows their role and

the process becomes second nature. According to our independent security audit, our incident response plan is *"comprehensive and well-practiced, showcasing preparedness for potential cybersecurity threats, with regular mock sessions carried out"*. This proactive stance means we can respond quickly to real incidents.

To date, Scribeberry has **not experienced any privacy breach**. However, we remain vigilant. We have internal monitoring (detailed in Section D) to catch suspicious activity early. If an event triggers our monitoring alarms (e.g. multiple failed logins, unusual data export), our incident process would initiate to check for a potential breach.

In fulfilling HIA's breach response requirements, Scribeberry is committed to **transparency and prompt notification**. Individuals' right to know about incidents affecting their information is respected fully. By having this structured incident management process, we ensure compliance with HIA Section 60.1 and its associated regulation (Health Information Regulation (HIR) Section 8.1) regarding **duty to notify**. We also align with best practices recommended by the OIPC ("Key Steps in Responding to Privacy Breaches"). Ultimately, our incident response capability is a crucial safety net that minimizes harm if any privacy incident were to occur.

6. Access and Correction Requests

Under the HIA (Sections 7, 10, 13), individuals have the right to **access their own health information** in the custody of a custodian and to request corrections to that information. While Scribeberry operates as a behind-the-scenes service (and typically the official medical record is maintained by the clinician or clinic outside of Scribeberry), we have procedures to handle any **patient access or correction requests** that might be directed to Scribeberry,

Access Requests: If an individual (or their authorized representative) were to contact Scribeberry to request access to health information about them that Scribeberry holds, we would direct them to their clinician, who would follow HIA's access provisions. First, we would verify the identity of the requestor and ensure they have the right to the information (for example, if it's a patient or someone with the patient's written authorization per HIA Section 104). Then, we would determine if we have any responsive records:

- Typically, Scribeberry **does not maintain a comprehensive medical record** for any patient – the information we handle is transient and ultimately resides in the clinician's records. We would likely refer the individual to the appropriate healthcare provider custodian (e.g. their doctor or clinic) who holds the official record.

- We would be unable to retrieve any existing notes or data related to that patient. Note that because our system isn't organized by patient (data is compartmentalized by clinician user account), we may need the requesting individual to provide specifics like date of service or clinician name to help locate their info. If data is encrypted such that we cannot read it ourselves, we will in all cases likely enlist the user (clinician) to obtain the decrypted content for the patient.

Scribeberry would respond to the requester **within 30 calendar days**, as required by HIA Section 13(1). We strive to respond openly, accurately, and completely (the HIA duty to assist, Section 10).

If we have no information on the individual (which might often be the case because data was not retained), we would formally reply that we have no records responsive to their request. We would also assist the individual by directing them to the custodian who likely has their records (their healthcare provider), fulfilling our duty to assist by helping them get access through the correct channel.

Exceptions to Access: Scribeberry is aware of the limited grounds on which access may be refused under HIA Section 11. Examples include if disclosure could result in harm to the individual or another, or if the information contains someone else's personal information that cannot be severed. In Scribeberry's case, our records (if any) would typically be just the patient's own info provided by their clinician, so most exceptions wouldn't apply. It's unlikely we'd have any meaningful information, but if we did, we would follow HIA requirements.

Correction Requests: If an individual believes their health information in Scribeberry's custody is incorrect or incomplete, they may request a correction (HIA Section 13). Upon such a request, Scribeberry would identify the custodian (their clinician/user of Scribeberry) and direct them to:

- Verify identity and authority (same as for access).
- Determine if we still have the information in question. Often, any note we had is now part of a provider's medical record. In that scenario, we would likely not have continuing control to make a correction – we'd advise the individual to request the correction from their healthcare provider custodian.
- The clinician must either correct the information as requested, or if they decline, notify the individual in writing of the reason and their right to have a statement of disagreement attached (as per HIA Section 14). Grounds to refuse a correction could be if the clinician believes the record is accurate.

Logging and Timelines: We treat requests with the same 30-day response timeline. All requests and outcomes are logged. Given the nature of Scribeberry, we anticipate very few direct patient requests; nonetheless, we have these procedures in place to fulfill the law.

Expressed Wishes: Although not directly an "access or correction" request, HIA Section 58(2) requires custodians to consider an individual's **expressed wishes** regarding how their health information is disclosed. In Scribeberry's context, a patient might express to their healthcare provider that they do not want their information used with an AI scribing service. We support honoring such wishes. Scribeberry advises its clinician users that if a patient objects to the use of Scribeberry during their visit, the clinician should respect that and not use the app for that

patient's encounter (and document manually instead). This effectively "masks" the patient's data from being collected by Scribeberry.

Practically, because we don't have a patient index, we would communicate with the clinician involved to make sure they refrain from using the service for that patient. Any expressed wishes that come to us are documented and considered whenever making decisions about information disclosure. Thankfully, Scribeberry does not contribute to provincial electronic health record systems like Netcare, so Section 56.4 (masking in Netcare) is not applicable, but the principle of respecting patient preferences still stands.

Requests from Clinician Users: As a note, clinicians using Scribeberry may sometimes request access to logs or their own data in the system. These are handled as internal administrative requests rather than HIA formal requests, but we have processes to provide clinicians with any of their data (for example, a clinician can export all their notes).

Overall, Scribeberry is committed to **assisting individuals in exercising their privacy rights**. Even though in most cases the custodian with the complete health record will be the healthcare provider, Scribeberry acknowledges its responsibility for any PHI it holds and has established straightforward procedures to meet HIA Part 2 (Access to Individual's Health Information) requirements. We have not yet received any direct patient access or correction requests, but our staff are trained on how to route and handle such requests appropriately.

Section C: Project Privacy Analysis

This section analyzes how Scribeberry meets specific HIA requirements in the context of the AI scribe system. It provides a detailed listing of the health information involved, describes the information flows (with a diagram and legal authority table), and explains our approaches to notice, consent, data matching, and third-party agreements.

1. Health Information Listing

Types of Health Information: Scribeberry processes the following types of individually identifying health information (within the meaning of HIA) in order to perform its scribing function:

- **Registration Information:** Personal identifiers and basic demographic details of patients may be handled. This can include the patient's **name**, contact information (address, phone number), **date of birth**, medical record or ID numbers, and provincial health number (PHN), if mentioned or input. These identifiers might appear in the dictated conversation or be input by the clinician to tag the note. *Purpose:* Registration information is used to correctly identify the patient in the clinical note and ensure the documentation is linked to the right individual. For example, including the patient's name

or chart number in the note helps the clinician later integrate the note into the patient's chart or EMR. It is essential for patient identification and to avoid misfiling of notes.

- **Diagnostic, Treatment, and Care Information:** This is the core content Scribeberry handles – essentially any information about a patient's health status, care, or treatment that is discussed during the encounter. Examples include: **medical history** (e.g. past illnesses, family history), **symptoms and observations** (the patient's complaints, vital signs, exam findings), **diagnoses** (provisional or confirmed diagnoses made by the clinician), **treatment details** (medications prescribed, dosages, investigations ordered, procedures performed), **care plans** and recommendations, **clinical assessments** and impressions, and **notes on progress** or follow-up plans. It may also include specific **clinical documents** or data points such as lab results, imaging findings, allergies, immunizations, etc., if these are discussed and dictated into the note. Essentially, any information that would normally be included in a patient's chart note or consult letter can be processed by Scribeberry's system. *Purpose:* The purpose of using this information is to create accurate and comprehensive **clinical documentation** of the patient's visit. This supports continuity of care – the information is captured in structured text so that the clinician (and the patient's health record) has a clear record of what transpired, what was found, and what the plan is. Scribeberry's use of this information is in direct service of providing health services (documenting the care provided to the individual).

- **(Optional) Scheduling/Billing Information:** In general, Scribeberry does *not* handle scheduling or billing data as part of its core functionality. Appointment times, billing codes, or insurance details are not typically processed by the AI scribe. These would be managed by other systems. We include this category for completeness: the only possible scheduling info could be **appointment date/time** if it's mentioned in the note (e.g. "Follow-up in 2 weeks on [date]"). **Billing information** (insurance numbers, costs) is not collected or used by Scribeberry. *Purpose:* N/A for Scribeberry's primary workflows, aside from minimal mention in notes (like stating a follow-up appointment).

Scribeberry ensures that **unique identifiers** (such as name, PHN, chart number) are only collected when necessary and are properly protected. For instance, if a clinician dictates the patient's name or ID in the note for clarity, that is captured, but Scribeberry itself does not require a PHN for any processing logic – it is simply part of the note content if present. We limit the health information to the **amount essential for the intended purpose**. In practice, that means Scribeberry processes what the clinician provides – we do not augment or pull in any extra data about the patient beyond what is dictated or entered. There is no linking of the notes with external health databases as part of our service (no automatic fetching of patient data from, say, Netcare or other sources).

Use and Disclosure of Each Type:

- Registration info is *used* within Scribeberry to tag the note and is included in the note output. It might be *disclosed* back to the clinician in the note and potentially via the cloud

sync (still under Scribeberry's controlled environment).

- Diagnostic/treatment info is *collected* from the clinician's input (which in turn comes from the patient), *used* by the AI to generate documentation, and *disclosed* through the system's outputs to the clinician and via our third-party AI processor (OpenAI) under strict contractual control (see flow analysis below).

- We do not disclose any of this information to unauthorized parties. Scribeberry itself does not send PHI to any external recipients except our contracted processors. We do not use PHI for any secondary purpose (like research or marketing). Notably, we **do not use any PHI to train AI models** – all AI processing is done on-the-fly with the patient's data kept private.

The above list of information types is consistent with the categories previously identified (registration information, and diagnostic, treatment, and care information). All are considered **health information under HIA** as they identify the patient and relate to health and health services provided. This table (if formatted):

Type of Health Information	Description/Examples	Purpose of Collection/Use
Registration Information	Patient identifiers (name, DOB, PHN, contact info)	Identify patient; label and file documentation correctly
Diagnostic, Treatment & Care Information	Clinical content of encounters (history, exam findings, diagnoses, treatments, notes on care)	Document and communicate details of health care provided; support clinical decision-making and continuity of care
(Scheduling/Billing Information)	*(Not generally used by Scribeberry; minimal if any)* Appointment details, insurance info	*Not applicable (documentation tool only)*

Legal Authority: Scribeberry only collects, uses, and discloses these types of health information for purposes that are authorized under the HIA – primarily the provision of health services to the individual. Detailed authorities for each flow are mapped in Section C.2.b (Legal Authority table). In summary, **HIA Section 20(b)** permits us to collect health information directly from the individual (via the clinician as intermediary) for providing a health service to them, and **Section 27(1)(a)** permits using health information for that purpose (treatment) without consent. Any disclosure to our information managers (e.g. OpenAI) is under **Section 66** and controlled by contract (which is an extension of the custodian's use, rather than a secondary purpose). We ensure no health information is used for anything beyond the stated purpose of clinical documentation, in compliance with HIA's collection/use/disclosure limitations.

2. Information Flow Analysis

This section describes how health information flows through the Scribeberry system, from the point of collection to processing and storage, and then provides a diagram and a table linking each flow to its purpose and legal authority.

Flow Description: The sequence of information flow in Scribeberry's AI scribe service is as follows:

- **(1) Collection from Patient via Clinician:** The process begins with a patient visiting a clinician (e.g., a doctor's appointment). As the patient shares information about their health (symptoms, history) and the clinician conducts an exam and gives assessments/plan, this **health information is collected from the patient**. Traditionally, the clinician would write or type notes; with Scribeberry, the clinician either **dictates** key information or allows the conversation to be recorded by the Scribeberry app (with patient awareness). So, the patient's PHI is captured **through the clinician's use of the Scribeberry application**. If using voice, the app's microphone picks up the conversation (or the clinician can summarize into the app).

- **(2) Secure Transmission to Transcription Server:** The audio stream (which contains the patient's identifiable health information in spoken form) is **encrypted in transit (TLS)** and sent to Scribeberry's **self-hosted transcription server** for real-time processing. This transcription server is a controlled component of Scribeberry – it does **speech-to-text conversion internally**. Importantly, no third-party cloud (like Google Speech API) is used; the audio data remains within Scribeberry's infrastructure. The audio is processed on-the-fly and not retained. Within seconds, the voice data is converted to text.

- **(3) Return of Transcribed Text:** The transcription server returns the **textual transcript** of the conversation to the Scribeberry application interface (again via an encrypted channel). At this point, the clinician can see a draft transcript of what was said (some models of use allow the clinician to see or hear the draft). The text at this stage is raw and may contain more information than needed (every utterance). It is temporarily held in memory/storage in our system to feed into the next step.

- **(4) Processing by AI (Azure/GCP/OpenAI):** The transcribed text (containing the patient's health information and context) is then sent from the Scribeberry backend to the **cloud AI service** to generate a structured clinical note. This is a **disclosure of PHI to our contracted AI processor(s)** for the purpose of transforming the data into a useful format (e.g., a SOAP note: subjective, objective, assessment, plan). The data is transmitted over a secure TLS connection. At the third-party end, their system (under the signed agreements) processes the text input and produces an output (the draft note). These parties do not store or use the PHI beyond this processing – Scribeberry's contract ensures the data is only used to generate the output and not retained or learned from. The processing happens in memory and the result is immediately returned.

Importantly, none of the data is retained for use of AI training.

- **(5) Return of Generated Note to Scribeberry:** The AI-generated **draft clinical note** (text) is returned to Scribeberry's system, which then delivers it to the clinician's device through the app. The note typically has a structured format summarizing the encounter. Throughout this transfer, data remains encrypted. At this stage, the patient's health information has been **successfully transcribed and summarized**.

- **(6) Clinician Review and Storage:** The clinician reviews the draft note on their device. They can make edits or corrections within the Scribeberry app (any edits are also just text and stay within the secure session). The final version of the note is then **available to the clinician to use** – usually they will copy or export it into their official medical record system (EMR) or print it for the patient's chart. Scribeberry assists with export (e.g. copy to clipboard or integration if available). Meanwhile, the note is also **stored in the Scribeberry cloud database** so that the clinician can access it later or from another device. This storage is **encrypted at rest**, and the data is indexed under that clinician's account. *Crucially, no one at Scribeberry can read this note in the database since it's encrypted.* The storage's purpose is to allow synchronization across devices and serve as a short-term archive for the user. The note is visible **only to the clinician (or anyone they explicitly share it with)** when logged in. Even Scribeberry administrators, if they looked at the database, would see ciphertext or otherwise have no capability to retrieve the plain content.

- **(7) Deletion/Disposition:** As described in Section B.3, the clinician may delete the note from Scribeberry once it's safely recorded in their EMR. The clinician has an in-app delete option; upon using it, the note entry is securely erased from Scribeberry's storage. After this, **no PHI remains in Scribeberry's custody** for that encounter. (If the clinician forgets to delete, Scribeberry's retention policy will auto-remove it after a set time.) In either case, the end of the flow is that the PHI is removed from Scribeberry's system once it's served its purpose.

Throughout all these steps, **encryption and access control are applied**. From the moment data enters Scribeberry (audio or text) until it leaves, it is encrypted in transit. When at rest in our system (like the stored note), it is encrypted with strong algorithms. **Scribeberry staff do not access the content at any point** – the system is engineered so that only the authenticated end-user (clinician) can invoke the processes and view the outputs. The **data flow remains within a closed loop** involving just the patient, clinician, Scribeberry system, and the third-parties aforementioned as a processor. Scribeberry does not broadcast or share the information elsewhere.

We note one ancillary flow: **user account data** (the clinician's own info). Clinicians register for Scribeberry with their name, email, etc. That personal information is stored in our system too, but since it's not patient health info, it's out of scope of HIA (though we protect it under

PIPA/other privacy laws). We mention it only to clarify that user authentication happens (with credentials exchange) but that does not involve patient health data.

Now, to **visualize the above flows**, we provide the information flow diagram below. Each numbered arrow corresponds to a flow described above.

Figure: Information Flow of Scribeberry PHI Processing. The diagram illustrates how patient health information moves through the Scribeberry system.

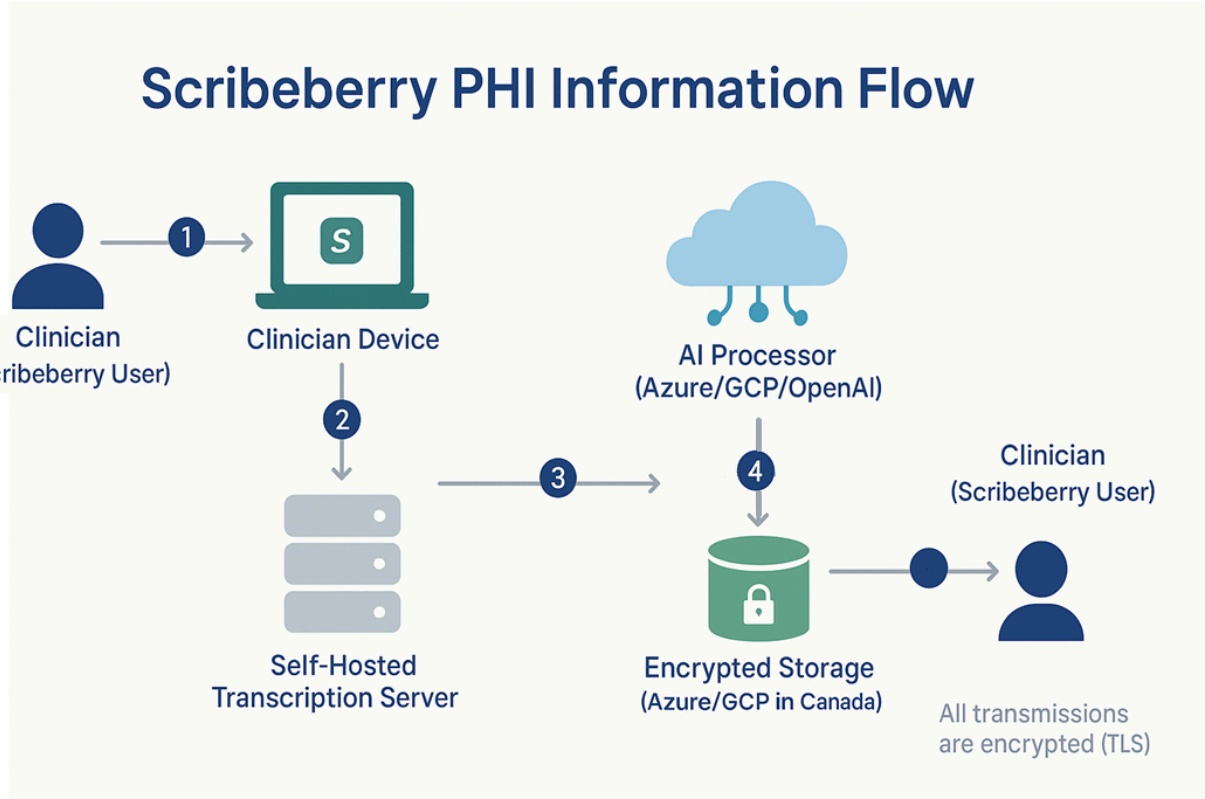

a. Information Flow Diagram

b. Legal Authority and Purposes Table: The following table enumerates each major information flow in the diagram, describes it, identifies the type of information involved, the purpose, and cites the legal authority under which the flow is permitted. Scribeberry, as custodian, ensures each flow is supported by a provision in the HIA (and any other relevant legislation for out-of-province disclosures).

Flow #	Description of Information Flow	Type of Information	Purpose of the Flow	Legal Authority *(HIA and other applicable)*

1	**Collection from patient via clinician into Scribeberry app.** The patient's health information (history, symptoms, etc.) is collected during the encounter and input to Scribeberry (spoken or dictated by clinician).	Registration info (name, etc.) and diagnostic, treatment & care information provided by patient.	To gather essential information needed for providing health services (assessment and documentation of the clinical encounter). This is the initial collection of PHI for the purpose of creating a medical record of the visit.	HIA **§20(b)** – Collection is limited to authorized purpose (health service to individual); **§21(1)** – Permitted collection from the individual with their implied consent; **§22(3)** – Notice of collection provided (see Section C.3). Also aligns with patient consent as part of seeking care.
2	**Use of PHI within Scribeberry for transcription.** The audio of the encounter, containing PHI, is transmitted to and processed by Scribeberry's own transcription server (internal use). No disclosure outside custodian yet.	Diagnostic, treatment & care information in audio form (identifiers may be included in speech).	To convert raw patient information (speech) into text – a necessary intermediate step to document and use the information effectively for care. This is an internal processing (use) of the info by Scribeberry's system to fulfill the documentation purpose.	HIA **§27(1)(a)** – Use of health information to provide health services to the individual (the transcription is part of the service of documenting the patient's care). The info is used only to the extent necessary (in line with §(60) need-to-know).

3	**Internal disclosure of transcribed text to Scribeberry app and clinician.** The transcribed PHI (text) is made available within the system back to the clinician via the app interface (still within Scribeberry's custody).	Diagnostic, treatment & care information (text form, includes any identifiers spoken).	To allow the clinician to review the captured information and ensure accuracy before finalizing the note. This step keeps the clinician in the loop of their patient's info as it's being processed.	HIA **§27(1)(a)** – (Same rationale as flow 2) Still considered a use for providing health service. If considered disclosure to the clinician: **§35(1)(a)** – Disclosure to the individual who is the subject of the info or their authorized custodian (the clinician as caregiver) for continuity of care. (However, since Dr. Moloo is custodian for this PIA, this is essentially internal.)
4	**Disclosure to Third-Party (information manager) for processing.** The transcribed text (PHI) is sent to a secured cloud service (a third-party processor under contract) to generate the structured note.	Primarily diagnostic, treatment & care information (the content of the encounter, which may include identifiers within the text).	To leverage specialized AI technology to analyze and summarize the health information into a useful clinical note format. OpenAI/GCP/Azure acts as a data processor performing a function (summarization) that Scribeberry itself would otherwise do – enabling the efficient creation of the documentation.	HIA **§21(1)** – Allows collection of health info from a custodian (clinician) by an information manager; **§27(1)(a)** – Use of info for providing health service (the processing is part of service delivery); **§35(1)(g)** – Disclosure to an **information manager** in accordance with HIA **§66**. (Additionally, compliance with **PIPEDA** in Canada and **HIPAA** in the US via BAAs ensures lawful processing in those jurisdictions.)

5	**Return of processed note from AI to Scribeberry (use).** The structured clinical note (PHI) generated by AI is returned to Scribeberry's custody and delivered to the clinician's app.	Diagnostic, treatment & care information in summarized text form (contains patient's health details, likely identifiers in header).	To provide the result of the AI processing back to the clinician for use in patient care (the clinician will use this note as part of the patient's health record). Essentially completing the service, enabling the clinician to review and confirm the documentation of the visit.	HIA **§27(1)(a)** – Use of the health info (the returned data) for continuing to provide health services (documentation) to the patient. Also covered by **§66**
6	**Storage of PHI note in Scribeberry's cloud database (use).** The final note is stored (in encrypted form) on Azure/GCP servers in Canada under Scribeberry's control. It is accessible to the clinician user on their account.	Registration info (e.g. patient name) and care information contained in the note.	To maintain the note for a short period so the clinician can access it from multiple devices, edit it, or retrieve it later if needed. It synchronizes the information for the user's convenience and ensures no loss of data between sessions. (This is part of delivering the service to the user.)	HIA **§27(1)(a)** – Use of info for the original purpose (the note remains within custodian's custody for care). Legal authority: **§66(1)** and **§35(1)(g)** allow storage with an info manager (Azure/GCP) as long as contractual safeguards per HIR 7.2 are in place. We have Azure's **Foundational PIA and compliance** for Alberta which covers their cloud offering, and we piggyback on that with our own controls.

7	**Disclosure to clinician or other custodian (external to Scribeberry).** The clinician exports or copies the note from Scribeberry into the patient's official medical record, or shares it with the patient or another healthcare provider. (This occurs outside the Scribeberry platform.)	Could include the full note with all PHI. The receiving entity is usually the original custodian (the clinician's clinic) or the patient themselves.	To ensure the documentation becomes part of the patient's health record for ongoing care. The note may be used for referral letters, follow-ups, or given to the patient. This step is essentially the **integration of the information back into the healthcare system** beyond Scribeberry.	HIA **§35(1)(a)** – Disclosure to the individual (patient) or **§35(1)(b)** – Disclosure to another custodian (e.g. the clinic's EMR, or a referring physician) for continuing care, which is allowed as it's for the purpose of providing health services. *Note:* This flow is typically initiated by the clinician user, who is a custodian in their own right. In the context of this PIA, if we consider Dr. Moloo's clinic (Hello Mental Health) separate, that clinic receiving the note would be a custodian-to-custodian disclosure authorized under HIA (implied consent for treatment). Since the PIA is focused on Scribeberry, this flow is indirect – Scribeberry enables but does not itself perform it. Nonetheless, it falls under permissible disclosures for care continuity (no consent required by HIA §§27, 35).

Explanation: The table above demonstrates that each flow of information has a defined purpose tied to providing health services or managing information as allowed by the HIA.

We also ensure compliance with any other legislation: For instance, if any personal information (non-health, like provider's notes about workflow) were involved, Alberta's PIPA might apply, but primarily this is health info under HIA. For U.S. flows, HIPAA is relevant – Scribeberry being HIPAA-compliant means we treat patient data with the same protections and allowed uses (treatment, payment, operations under HIPAA). Our cross-border data flows are addressed with contracts (BAAs) and are consistent with PIPEDA (for handling data across provinces and into the U.S.). Essentially, no flow violates any jurisdiction's privacy laws.

In conclusion, the information flow analysis shows that **Scribeberry collects, uses, and discloses health information only for the purposes of providing or supporting health services to the individual**, and each such action has a clear authority under the HIA. We **map**

each flow to a specific HIA section to ensure there are no unauthorized uses or disclosures. This mapping exercise also doubles as our internal compliance check – if a proposed new flow cannot find support in legislation, we would know it's not allowed. The diagram and table will be kept updated if Scribeberry's processes change, and they serve as a reference for staff and any interested party (like OIPC) to understand exactly how data moves in our system.

3. Notice

Under HIA Section 22(3), when collecting health information directly from the individual, custodians must provide notice of the purpose, legal authority, and contact person. In Scribeberry's context, **patients are made aware that their health information is being collected and processed using an AI scribe system**, and the required elements of notice are given to them prior to or at the time of collection.

Method of Notice: Since Scribeberry is typically used by healthcare providers during patient visits, we have empowered and required those providers (our users) to deliver the collection notice to patients. Scribeberry has developed a **Patient Information and Consent Form** (see Attachment E.4) that clinicians can use with their patients. This form serves as both a notice and (if desired) a consent for using Scribeberry. The notice can also be given verbally or via signage:

- Many clinics using Scribeberry post a **privacy notice poster** in waiting or exam rooms stating: *"This clinic uses an electronic medical scribe (Scribeberry) to assist in documenting your visit. Your health information may be processed by this tool, which uses secure Canadian cloud services. We ensure confidentiality and compliance with Alberta's Health Information Act."*

- If not posted, the clinician will verbally inform the patient at the start of using Scribeberry. For example, the clinician might say: *"I will be using a secure medical scribe application to help take notes today. It will record what we discuss and generate the note. Only I can see the information, and it's protected and kept confidential."* This gives the patient an opportunity to ask questions or object (in which case the clinician can choose not to use the system for that patient).

Contents of Notice: In accordance with HIA 22(3), our notice includes:

- **Purpose of Collection:** We explain that the information is being collected for the purpose of documenting the patient's health care encounter and assisting the provider in creating the medical record. (E.g. "to create an accurate clinical note of your visit today.")

- **Contact Information:** The notice provides the name, business address, and telephone number of a person who can answer questions about the collection. We list **Dr. Zaahir Moloo, Privacy Officer, Scribeberry Ltd.** along with an address and a phone/email contact. If the clinic prefers, they can list their own clinic privacy officer too, but

Scribeberry has ensured Dr. Moloo is available to field any queries on behalf of our system.

An example from our **Collection Notice** form reads: *"Why is information collected? – To help your doctor document your care using an AI scribe tool (Scribeberry). / Questions? – Contact Dr. Zaahir Moloo (Privacy Officer) at [address], [phone]."*

Timing: Patients are provided this information **before or at the time their information is first collected into Scribeberry**. Ideally, when a clinic implements Scribeberry, they inform their active patients through an intake form update or a posted notice. New patients are given the info as part of registration. In practice, because Scribeberry's use is often evident (the clinician might hold a device or mention it), that's another opportunity to ensure the patient has seen the notice and agrees to proceed.

Documentation of Notice: If a patient signs the consent form, that is kept on file by the clinic (and a copy can be provided to Scribeberry if needed, though we generally trust our users to handle that). If notice is only posted and verbal, no signature is needed, but Scribeberry encourages clinicians to note in their chart "Patient informed of use of Scribeberry; no objections" to have a record.

Patient Questions: Should a patient have questions or concerns after reading the notice, they can reach out using the contact info. In such cases, Dr. Moloo (or delegate) would explain in more detail how Scribeberry works, what safeguards are in place, and address any privacy concerns. Transparency is key – we do not hide the fact that an AI and cloud service is involved. In fact, Scribeberry's privacy policy (available on our website) further explains data practices in plain language, which patients (or anyone) can consult. Reference: Https://trust.scribeberry.com

Compliance with Notice Requirement: By providing this notice, we fulfill Section 22(3) of HIA which mandates informing the individual of the purpose, authority, and contact before collection. Our approach is multi-modal (written and verbal) to ensure patients truly are informed. We believe most patients appreciate knowing that an advanced tool is being used, and given the reassurance about privacy, they usually accept or even welcome it. If a patient were to object (expressed wish), as noted earlier, the clinician would not proceed with Scribeberry, thus respecting their autonomy.

To summarize, **Scribeberry ensures that individuals are given clear notice of the use of their health information in this new system**, addressing why it's needed, under what law it's allowed, and who to contact with concerns. This notice is often integrated into the clinic's overall privacy notice to patients, which might also mention other systems (like EMRs or eHealth systems). Scribeberry has provided template language to clinics so that the notice is accurate and comprehensive. We consider this notice procedure a critical component of fair information practices and patient transparency.

4. Consent and Expressed Wishes

Consent:, Scribeberry recognizes that some aspects might make patients or providers more comfortable with explicit consent. Our provided **Patient Consent Form** (Attachment E.4) is optional for clinics – it doubles as notice and obtains the patient's signed consent to use the AI scribe. If a patient later withdraws consent, that equates to an expressed wish not to use Scribeberry (which we handle as below).

Any health information use beyond the direct care purpose would require consent. **Scribeberry does not engage in such secondary uses**. For instance, if we ever considered using de-identified data for research or to improve our algorithms, we would either ensure it's fully anonymized (thus not subject to HIA) or we would seek proper research ethics approvals and patient consents under HIA Division 3. Currently, we have a firm policy: *no PHI is used to train AI or for any purpose other than delivering the note to the user*, eliminating the need for any additional consent.

Scribeberry and its user clinicians will **honour any such expressed wishes**. Our policy (communicated to client clinics) is that if a patient objects to the use of the Scribeberry system (for any reason – privacy, comfort, etc.), the clinician should not use it for that patient's encounter. The clinician can revert to traditional note-taking or another method that the patient is okay with. This is similar to how a patient might decline having a medical student in the room; here they decline the AI scribe.

What happens in such a case: The clinician indicates in the app or setting that this patient is opted out (or simply doesn't activate the app). Scribeberry does not receive any PHI from that encounter. If any minimal data had been collected before the objection (say the start of an encounter), the clinician can delete it immediately. Thus, the patient's wish to not have their data disclosed beyond the clinician is respected.

If a patient's expressed wish is more nuanced (e.g., they allow the scribe for most but say "please don't include my mental health history in the AI system"), the clinician would exercise judgment to either exclude that portion from Scribeberry (maybe pause the app during sensitive portions) or again choose not to use it to avoid partial data. We want to avoid fracturing the note integrity, so more often it's all or nothing per encounter as per patient comfort.

Masking and Netcare: As noted, Scribeberry is not integrated with Netcare, so the specific scenario of Global Person-Level Masking in Netcare doesn't directly apply. If a patient had masked their record in Netcare, that doesn't restrict Scribeberry's usage because Scribeberry is a separate tool under the custodian's control, not an information exchange. Nonetheless, we keep aware of those regulations – should Scribeberry ever interface with such systems, we would incorporate respect for those expressed wishes accordingly.

Documentation and Consideration of Wishes: If a patient expresses a wish to not use Scribeberry, the clinician should document that preference in the patient's chart (so that future visits also honour it). Scribeberry will have a mechanism for clinicians to flag an account or note with "Patient does not consent to AI scribe" if needed. So even if a different provider in the same clinic sees that patient, they get a heads-up.

In summary, Scribeberry's stance is to fully **respect individuals' expressed wishes regarding the handling of their health information**. Our design gives control to the clinician (and thereby the patient) to opt out or limit use as needed. We treat any such wishes seriously and incorporate them into our processes, ensuring compliance with HIA Section 58(2).

5. Data Matching

"Data matching" in the context of HIA refers to combining data from multiple sources to create new identifiable health information (often for analytics or research). Scribeberry's system **does not engage in any data matching** as defined by the HIA. Our service operates on a per-encounter, per-patient basis, using the information from that single encounter to produce a note for that same encounter. We do not combine one patient's data with another's.

To be explicit:

- We are **not linking or cross-referencing** different data sets to generate new insights about an individual. The AI is using only the data from that individual's session provided by the clinician. We may however, for example, take a patient's current visit information and pull in their hospital records or pharmacy records to combine – that kind of integration is at the sole discretion of the clinician user using Scribeberry. It is not automated.

- We do not create new individually identifying information by combining non-identifying data. The output note is derived from identifiable input and remains about the same person, so no new identity is created.

We have confirmed that none of the use cases of Scribeberry meet the HIA definition of data matching (HIA §1(1)(g)). We are simply transforming data (speech to text, text to formatted text) about one individual in one context. This transformation does not produce any new identifying information beyond what was already provided about that individual.

If in the future Scribeberry considered any analytics (like analyzing patterns across many patient notes to improve service or for research), we would do so only on **de-identified or aggregate data**. Even then, combining multiple patients' de-identified notes to find common trends does not create new identifiable info, so it likely remains outside the "data matching" provisions. We remain aware of sections 68–72 of HIA which set rules for any data matching activities (especially if we ever worked with another custodian's data sets or for research). At present, those sections are not triggered by Scribeberry's operations.

Thus, we **do not engage in data matching** as part of this project. There are no plans to do so.

For this PIA, since no data matching is occurring, we simply note compliance by absence: we adhere to HIA's data matching rules by not performing any unauthorized combining of data sets. Scribeberry focuses solely on the data provided in each instance by custodians in real time for

direct care. This approach avoids the complexities and privacy risks associated with data matching altogether.

6. Contracts and Agreements

To operate Scribeberry's service while protecting privacy, we have entered into formal **contracts** with all third parties and affiliates that may handle health information on our behalf. These agreements delineate roles, responsibilities, and privacy safeguards. Below is an overview of the key contracts:

- **Microsoft Azure:** Azure provides the cloud infrastructure (servers, databases) in Canada where Scribeberry's application and data reside. We have a comprehensive **data processing agreement/Business Associate Agreement** with Microsoft.. It includes clauses ensuring Azure will only process data as instructed by Scribeberry, implement strong security, restrict access to authorized personnel, and report any privacy breaches to us immediately. It also incorporates Canada-specific requirements for information managers storing data out-of-province (for example, Azure acknowledges HIR 8(4) obligations about complying with Alberta law despite data possibly being accessible outside Alberta). Microsoft's adherence to Canadian privacy laws is well documented. The contract also covers data return or destruction if services end.

- **Google Cloud Platform (GCP):** Similarly, we use GCP for some services (potentially redundancy or certain components). We have a signed **data processing addendum** that mirrors the protections of the Azure agreement. Google contractually commits to storing Canadian data in Canadian data centers for Canadian users, and to abide by confidentiality and security standards. Both Azure and GCP agreements ensure compliance with provincial laws where applicable and PIPEDA federally.

- **OpenAI (and Anthropic) API:** For the AI language model service, we have a specialized **Business Associate Agreement (for HIPAA)** and a data processing agreement. This BAA is crucial since PHI is sent to OpenAI; under it, OpenAI agrees not to use the PHI except to provide the service to us, not to disclose it, to safeguard it per industry standards, and to assist us in any compliance needs (like providing audit info or breach notifications). It also ensures that any data input is not used to train their models or any other purpose. We have verified OpenAI's compliance with handling PHI in a HIPAA-compliant manner. (Anthropic is another AI vendor we have available via a BAA, though currently OpenAI is primary; the same contractual conditions apply if we use Anthropic).

- **Transcription Engine (if licensed from a vendor):** Our self-hosted transcription server uses our own proprietary speech-to-text software. If that software is from a vendor, we have an agreement with that vendor as well that no data leaves our server to them (i.e., the processing is local, so perhaps not needed. If cloud-based, we have an agreement

too, but we specifically chose a self-hosted solution

Each of these include maintaining confidentiality, restricting access, using information only for listed purposes, not retaining information beyond our instructions, providing safeguards, and assisting the custodian with any access requests or amendments if needed.

Employee/Contractor Agreements: All Scribeberry employees and any independent contractors (affiliates) with potential access to health info sign **confidentiality agreements** and are bound by our employment contracts or contractor agreements that include privacy provisions. These agreements:

- Require them to follow all our HIA-related policies and procedures.
- Prohibit any unauthorized use or disclosure of health information.
- Require them to report any privacy incidents immediately.
- Impose consequences for breach (tying to our sanctions policy).
- Survive termination (they must continue to keep info confidential even after they leave). Additionally, contractors that assist with our operations (like an IT support person) also sign a **non-disclosure agreement (NDA)** specifically covering any sensitive data they might incidentally see. For example, if we hired a cleaning service for our office where someone might see something on a screen, we would have an NDA with them (though our screens are locked, but this is just a precaution mentioned by HIA guidelines).

Agreements with Clinics (Users): Scribeberry provides its service to healthcare providers (clinics, doctors). In those relationships, typically the clinician is the custodian and Scribeberry would be the information manager. Indeed, we have a standard **Service Agreement** that we sign with each clinic or provider using Scribeberry, in which Scribeberry commits to safeguarding PHI on behalf of that provider. That agreement ensures both parties' obligations are clear. This PIA, however, covers Scribeberry on a broad scale, but we mention that such agreements exist for completeness. Those contracts mirror many of the terms we have with our vendors (ensuring Scribeberry's duties: only use info for providing the service to that provider, not for secondary use, assist with access requests, report breaches, etc.). We include references to these contracts in our attachments (though they are more about our role as info manager).

Other Service Providers (Non-affiliates): We have considered if there are any other services that might **incidentally encounter PHI**. For example, if we used an external email service to send communications that included PHI, or a cloud logging service. We have avoided sending PHI through such channels. Any support tools we use are configured not to log PHI. For instance, our error logging software captures no user data. If we ever needed to share a screen

or data with a third-party support (like Azure tech support) that included PHI, we would execute an NDA or ensure a confidentiality clause covers that. Currently, we do not foresee non-affiliate service providers encountering PHI. Even so, we have basic NDAs in place with our general vendors (like our accounting firm has an NDA, though they see no PHI, just company data – but it's a good practice as mentioned in guidelines).

Summary of Key Contract Terms: Across all these agreements, the common privacy and security terms are:

- **Limitation of Use/Disclosure:** Third parties can only use the health information for the purposes we instruct (e.g. OpenAI only to generate the note, Azure only to store it). They cannot mine it, sell it, or use it for their own purposes.

- **Safeguards:** They must have appropriate administrative, technical, and physical safeguards (often specified to meet standards like ISO 27001 or SOC2). For example, Azure and Google are certified for high security and encryption; OpenAI under BAA isolates our data and doesn't log it.

- **Confidentiality:** Any personnel of theirs who handle our data must be under confidentiality obligations.

- **Breach Notification:** They must notify Scribeberry immediately of any security incident or breach involving our data. Our contracts usually require notification within a specific short timeframe (e.g. within 24 hours of discovery).

- **Sub-processors:** If those vendors use sub-contractors (sub-processors), they must impose the same data protection obligations on them and often must inform us or get consent for sub-processors (especially in EU-style DPAs; we align with that).

- **Audit and Oversight:** Scribeberry has rights to inquire or audit compliance. Practically, for giants like Azure, we rely on their third-party audits and certifications, but contractually we have the right to get evidence of compliance. For smaller vendors like maybe OpenAI, our BAA allows us to request documentation of their safeguards.

- **Termination and Post-Termination:** Upon termination of the service, the vendor must return or destroy the PHI they have. E.g., if we stop using OpenAI, they must delete all our data. Azure/GCP would purge storage when we delete or end contract.

- **Compliance with Law:** They must comply with applicable privacy laws (HIPAA for US, PIPEDA/Provincial laws for Canada) as needed. For instance, Azure's contract references compliance with Canadian provincial health privacy requirements.

- **Jurisdiction Clauses:** Since some data is outside Alberta, our agreements include clauses to address cross-border. We ensure there is no ambiguity that our data remains under our control and subject to HIA standards even if stored in, say, the US. We also

inform patients of this reality in notices, as required by best practices.

By establishing these contracts, Scribeberry has **formalized the accountability chain**: all parties that touch PHI are contractually bound to protect it to the same level Scribeberry is required to. This significantly reduces privacy risk, as recognized by our independent security assessment, which noted *"Dependable third-party management: Scribeberry maintains partnerships with third-party vendors, fortified with stringent BAAs and regular security audits"*. Our relationships with Microsoft, Google, and others were found compliant with regulations and best practices.

Additionally, Scribeberry's **Terms and Conditions and Privacy Policy** for the service (which are agreed to by clinician users) explicitly mention these safeguards and that we have done PIAs and agreements. This gives our users confidence in the product's compliance.

In summary, **all necessary contracts and agreements are in place** to meet Canadian and regional privacy legislative requirements:

- Information Manager Agreements with cloud and AI providers
- Affiliate agreements (employment/contractor) ensure compliance
- Non-disclosure agreements cover any peripheral service providers.
- Data sharing with clients (custodians) is governed by contracts

These agreements are on file, and we are prepared to provide them (or excerpts) to OIPC if required for review. Scribeberry will also ensure to update or add agreements if the scope of the project changes (for example, if we engage a new third-party), we will execute an agreement **before** any health information is shared

---

Section D: Project Privacy Risks and Mitigation Plans

In deploying the Scribeberry system, we have conducted thorough risk assessments to identify potential threats to privacy and have implemented a comprehensive set of **administrative, physical, and technical safeguards** to mitigate those risks. We continuously monitor compliance and plan for regular reviews of this PIA, in line with privacy obligations. Below we discuss the key privacy and security risks we considered and the measures in place to address them. We also describe how access to health information is controlled, how compliance is monitored, and our commitment to keeping the PIA up to date.

Risk Analysis and Mitigation Summary

Scribeberry's independent Security Risk Assessment (SRA) found that the overall risk level to ePHI is **significantly low**, with strong controls in place. We will summarize primary risk areas and mitigations:

- **Unauthorized Access (External threats):** There is a risk that hackers or malicious external actors could attempt to gain access to Scribeberry's systems or intercept data in transit. **Mitigations:** We use state-of-the-art **encryption** for all data in transit and at rest, so even if data were intercepted, it would be unreadable. Specifically, TLS 1.3 with strong cipher suites protects data in transit, and AES-256 encryption protects data at rest on servers (backed by Azure/GCP's robust encryption infrastructure). We enforce strong authentication for system access: clinicians log in with secure passwords (and we are implementing optional two-factor authentication for user accounts). Internally, admin access to servers is restricted to a few key personnel and requires SSH keys or VPN access plus multi-factor authentication. Firewalls are in place on all servers to only allow necessary traffic (e.g., app communications, API calls) – everything else is blocked. We also have intrusion detection/prevention systems (IDS/IPS) monitoring our cloud environment for suspicious activities, and automated threat mitigation from Azure's security center. Frequent security testing (vulnerability scanning, penetration tests) is conducted to identify and patch any weaknesses. The SRA specifically noted our encryption protocols are *advanced, meeting industry standards*, and that these measures ensure maximum security of sensitive data. In essence, **data is protected in transit and storage**, and system entry points are heavily guarded.

- **Unauthorized Access (Internal threats):** The risk that a Scribeberry employee or an unauthorized internal person could access patient data. **Mitigations:** Scribeberry enforces a strict **role-based access control** model. By design, employees have **no direct access to PHI** – for example, developers use dummy data for testing, and production databases with real data are not directly browsable by them. Only the Privacy Officer and CTO have the ability to access production data stores, and even then, the data is encrypted which prevents viewing content. We've implemented technical measures so that even database administrators cannot query sensitive fields in plaintext. Additionally, all staff and contractors are bound by confidentiality agreements and trained on privacy (so they understand the seriousness of unauthorized access). System logs are monitored to detect any unusual access patterns. For instance, if an admin account tried to retrieve large amounts of data, alerts would trigger. We also restrict internal access by environment: the transcription server and AI integration run autonomously without human intervention, so staff don't handle that data. Background checks are performed on employees who will have any privileges (ensuring trustworthy personnel). Finally, as a policy, **we do not allow employees to use real patient data for any purpose** other than what the user (clinician) explicitly needs – no troubleshooting with real PHI unless absolutely necessary and with approval. To date, we've had zero

incidents of internal snooping, and our controls aim to maintain that.

- **Data Breach from Vendor Systems:** Risk that our cloud providers (Azure/GCP) or OpenAI could suffer a breach, exposing Scribeberry's data. **Mitigations:** We selected reputable, security-certified vendors (Azure and GCP are industry leaders with extensive security programs, holding certifications like SOC2, ISO 27001, and in Azure's case, undergoing their own PIA for Alberta). We encrypt data before it goes to these services, adding an extra layer beyond their controls.**(continued)**

Our BAAs/IMAs also commit the vendors to **physical security** measures (their data centers are highly secure facilities with restricted access, CCTV, guards, etc. – Azure and Google meet Tier IV data center standards). This mitigates physical intrusion risks. Scribeberry's own office doesn't store PHI locally, but we secure our workstations and enforce clear screen policies anyway.

- **Data Loss or Corruption:** Risk that PHI could be lost (e.g., server crash) or corrupted (e.g., ransomware). **Mitigations:** We perform **regular backups** of data (with encryption) and have a disaster recovery plan. Data is stored redundantly across multiple servers/availability zones in Azure, reducing chances of total loss. We also protect against ransomware by using advanced anti-malware on our servers and applying strict access controls (minimizing entry points for an attacker). Our disaster recovery (Contingency) plan ensures we can restore service and data quickly if something goes wrong. We test backups periodically. Also, since clinicians typically transfer notes to their own systems, there is an additional copy of information outside Scribeberry; thus even in worst-case, the provider still has their records.

- **Inaccurate or Unauthorized Alteration of Data:** Risk that data might be improperly modified, either accidentally or maliciously (integrity risk). **Mitigations:** Audit logs record any changes to notes or system data – who made an edit and when. Clinicians can trust that what they see is what they said; if the AI made an error, the clinician can correct it, but the system itself won't arbitrarily alter stored data. We implement checksums and integrity checks on data storage to detect corruption. All configuration changes to the system require approval and are tracked (to ensure no one can insert code that might tamper with data). The role-based access also ensures only the clinician (or authorized delegate) can edit their patient's notes.

- **Non-compliance or Human Error by Users:** Risk that clinician users might misuse the system or not follow proper privacy practices (e.g., leaving their device unlocked with the Scribeberry app open). **Mitigations:** We train our client users via materials and support on proper use (like advising them to enable device locks, not to use Scribeberry on public Wi-Fi without VPN, etc.). The application auto-times out after a period of inactivity to prevent a session left open. We have user terms that contractually require them to use the service in compliance with privacy laws as well. We also designed the app to minimize any PHI displayed unnecessarily (for instance, the app might abbreviate patient

name on screen to first name only to reduce exposure if someone glances at it). Ultimately, clinicians are custodians too, so they are motivated to use it responsibly.

- **Cross-border Data Flow Risks:** Only USA users have some data that may flow to the USA (OpenAI, or if a U.S. user's data is stored in U.S. servers). There are theoretical risks related to foreign jurisdiction (e.g., Patriot Act access). **Mitigations:** All data is encrypted such that even if a foreign authority sought it, they would need to obtain keys from us (and we would resist any unlawful disclosure). We have contractual clauses that require vendors to notify us of any government demands. For Canadian patients, we architected to keep data in Canada to avoid this. We are transparent about these flows so that if a patient or provider is uncomfortable, they can choose not to use the service (informed decision). Additionally, Scribeberry being subject to PIPEDA means if U.S. law enforcement requested data, we would only comply if the request goes through proper Canadian/Alberta channels (e.g., via treaty) – we wouldn't volunteer anything just because a US entity asked. This mitigates legal risk around jurisdiction.

Overall, our risk mitigation strategy was validated by an independent audit which highlighted **advanced data protection, strong access controls, and a proactive incident management** as key strengths.

Administrative Safeguards

We have implemented numerous **administrative controls**:

- **Policies & Procedures:** As detailed in Section B.2, all privacy and security policies are current, documented, and accessible to staff. They provide clear instructions on how to handle PHI and what is not allowed.

- **Training & Awareness:** Section B.3 described our comprehensive training program, which ensures every affiliate is aware of their duties and the correct procedures. This reduces human error and insider misuse.

- **Authorized Access Management:** We maintain strict lists of who is authorized to access what. E.g., only specific team members have admin rights. Any new employee's access is approved by the Privacy Officer in line with the principle of least privilege.

- **Confidentiality Agreements:** All staff have signed confidentiality/oath agreements and understand penalties for violations.

- **Vendor Management:** We vet and contractually bind all service providers (see Section C.6). We also review their compliance reports annually.

- **Privacy Impact Assessment Review:** Administratively, we plan to review this PIA on a regular basis (at least every year or when changes occur). We treat the PIA as a living

document – any system change triggers a privacy review step in our project management.

- **Breach Response Plan:** As in Section B.4, we have a documented breach response procedure and have trained for it. It clearly assigns roles and steps to contain and notify, etc. This administrative readiness is crucial if something goes wrong.

- **Monitoring compliance:** The Privacy Officer conducts periodic audits (see Monitoring below) to ensure staff follow procedures and that no unauthorized activity is happening.

- **Sanctions:** There are defined sanctions for any privacy violations by staff, which are known to employees (ranging from retraining to termination). This deterrent helps enforce compliance.

These administrative measures ensure that on an organizational level, privacy is managed actively and effectively, reducing the risk of policy failures or unaddressed vulnerabilities.

Physical Safeguards

Physical protections are also in place, though much of our infrastructure is cloud-based:

- **Office Security:** Scribeberry's office (where development and support happen) is in a secure building with access controls. Only authorized staff can enter the work area. We keep no paper health records. Any scribbles or notes containing PHI (rare) are immediately shredded with a cross-cut shredder. Workstations auto-lock after a few minutes of idle time.

- **Device Security:** All company laptops and devices are encrypted and require password/PIN or biometric access. Staff are instructed not to store any PHI locally, but if any existed, the whole disk encryption and remote wipe capability protect it.

- **Data Center Security:** Our information managers (Azure, GCP) handle physical security of servers: their data centers have multi-factor authentication for entry, security guards, surveillance, and environmental controls. We have reviewed Azure's physical security attestations. They are far more secure than any on-premise clinic server typically would be.

- **Secure Disposal:** We have no physical PHI to dispose of, but if we retired a hardware device that had any cached PHI, we would securely wipe it. We dispose of obsolete media via certified e-waste services with certificates of destruction.

- **Contingency Site:** In the event our office is inaccessible, staff can work remotely securely (VPN to our cloud environment). There is no physical single point of failure in

terms of access to data – it's all in the cloud which has its own redundancies.

Physical risks (like theft of a server or unauthorized person entering) are thus minimal. The most PHI someone could physically steal would be perhaps a logged-in clinician's device; however, as noted, sessions timeout and device-level security would mitigate unauthorized access.

Technical Safeguards

We have heavily invested in technical security controls:

- **Encryption:** As repeated, all PHI is encrypted in transit (HTTPS/TLS) and at rest on servers. Encryption keys are managed in a Key Vault with tight access. We also encrypt sensitive fields at the application layer.

- **Authentication & Access Control:** Each user has a unique account (unique ID for audit). Passwords must meet complexity and are hashed with a strong algorithm (bcrypt). We support multi-factor auth. Access within the system is role-based: e.g., a normal user cannot access admin functions. There are no shared accounts. Admin interfaces are protected behind VPN and additional login. We maintain an access control list and review it quarterly to remove any unnecessary accounts.

- **Audit Logging:** The system generates detailed logs of access to health info – e.g., when a clinician views or edits a note, when data is sent to OpenAI, when an admin logs into a server, etc. These logs are timestamped and immutable. We use a SIEM (Security Info and Event Management) system to aggregate logs and flag anomalies (like an admin logging in at odd hours or multiple failed logins). Regular audits of logs are conducted (see Monitoring below).

- **System Security:** Servers and software are kept updated with the latest patches. We have automated patch management for both OS and our application. Anti-malware and endpoint protection is deployed on all servers (with real-time protection and scanning).

- **Network Security:** Our cloud network is segmented – e.g., the database is not publicly accessible, only the application server can talk to it. Firewalls (security groups) strictly allow required traffic (for example, only allow OpenAI API calls to known IPs). We also employ **intrusion detection/prevention** systems. Azure provides threat monitoring that alerts on suspicious network patterns. We've configured rate limiting to thwart brute force attacks on logins.

- **Backup and Recovery:** We do daily encrypted backups stored off-site (within cloud). We periodically test restoration from backup (at least quarterly) to ensure data integrity. We have documented restoration procedures as part of our contingency plan.

- **Development Security:** Our code is developed following secure coding practices. We also underwent an independent code and architecture review during our security assessment to ensure no glaring vulnerabilities. We utilize environment-based secrets management (no hard-coded credentials). We also run static code analysis and dependency vulnerability scanning in our build pipeline.

All these technical controls mitigate the risk of unauthorized access, data breaches, or loss of data. The independent SRA praised the **"advanced encryption**" and our **"role-based access, continuous monitoring, regular reviews"** that ensure only authorized personnel access ePHI. These align perfectly with HIA Section 60 requirements for technical safeguards and secure storage.

Monitoring, Audit, and Ongoing Compliance

Scribeberry recognizes that implementing safeguards is not enough – we must **continually monitor and verify** that those measures are effective and being followed. Our monitoring and audit activities include:

- **Continuous Privacy Monitoring Dashboard:** We have implemented a third-party continuous monitoring platform (GetDelve) which publicly shows aspects of our security posture. Internally, this tool monitors our compliance with privacy controls in real-time – it checks that our policies, training, and technical measures remain in place and flags deviations. This transparency is an extra layer of accountability.

- **Regular Audits:** Dr. Moloo (Privacy Officer) or a delegate conducts **quarterly privacy audits**. These involve reviewing access logs (did anyone other than the expected user access data?), ensuring that all active employees have up-to-date training, verifying that any new hire signed NDAs, checking that backups have been performed, etc. We use an audit checklist derived from the OIPC's PIA Requirements and our own policies. For example, we audit a random sample of logins to confirm they were legitimate, and we verify no anomalies like an admin querying database tables with PHI. We document these audits and any findings.

- **User Auditing:** On the user side, if we suspect misuse by a clinician (though they are ultimately responsible to their patients), we have the ability to audit their usage if needed – e.g., if a clinic reported a concern that a user accessed data they shouldn't (not common in our one-user-one-patient model, but a theoretical scenario).

- **Automated Alerts:** As mentioned, our SIEM sends alerts to the security team's email/phone if certain triggers occur (multiple failed admin login, a new device logging in from an unusual country, a spike in data download, etc.). These are investigated immediately according to our incident response plan.

- **Annual SRA Review:** Scribeberry will repeat a formal Security Risk Assessment at least annually. This reassessment will catch new threats (cyber threats evolve) and ensure our controls remain sufficient. The SRA documentation, including risk register and mitigation plans, is maintained and updated with each review.

- **Policy Review:** As noted, we review policies annually and when needed. We also plan a **yearly PIA review**. The Privacy Officer will revisit this PIA document and verify if any changes in our practices require an update (HIA Section 64(2) implies PIAs should be updated for changes). If we introduce a new feature or partner, we will perform a mini-PIA analysis and submit an amendment to OIPC if the changes are significant. For instance, if in 6 months we enable a new AI engine or start integrating with a patient portal, we will update the PIA accordingly. We've calendared a PIA review meeting every 12 months where we go through the checklist (like the one in the Alberta template) to ensure nothing is overlooked.

- **Compliance Checks:** We monitor legal and regulatory developments (e.g., Alberta may update HIA or regulations; other provinces might have new requirements). Our privacy consultant/auditor helps keep us informed. We adapt our compliance program proactively (e.g., Law 25 in Quebec – we adjusted some consent features accordingly). As noted on our Trust Center, Scribeberry stays compliant across all provinces and federal laws.

Through these monitoring efforts, we verify that affiliates follow policies and we catch issues early. If an audit finds any gap, we address it promptly (e.g., retraining an employee, patching a system). Senior leadership is kept informed of audit results, and privacy is a standing agenda item in management meetings. We believe this ongoing vigilance is vital because security is not a one-time setup but an ongoing process.

PIA Maintenance and OIPC Notification

We consider this PIA not as a one-off task but part of our operational process. Scribeberry is committed to **reviewing this PIA regularly and updating it as needed**. Specifically:

- We will review the PIA at least annually (as stated).

- We will update the PIA if there is a change in our information systems or practices that could impact privacy (for example, a new data element collected, a new use of info, a new third-party integration, or expansion to a new jurisdiction).

- If updates are made, we will submit the PIA amendment or new PIA to the OIPC **before** implementing the change, as required by HIA Section 64(1). In fact, we already followed this approach when we made certain changes – for example, we prepared a PIA amendment for adding new features and submitted it (this refers to an earlier PIA

amendment submission we've done).

- All versions and file numbers from the OIPC will be tracked in the cover page of the PIA (for reference, Hello Mental Health's initial PIA was one reference; this independent PIA will be a new file number; any subsequent amendments will cite this one).

Additionally, even if no changes occur, if after a couple of years the OIPC has new guidance or we find improvements to be made, we may voluntarily update the PIA to reflect any enhancements in our practices.

We also maintain a dialogue with the OIPC. If ever in doubt whether a planned change requires an amendment, we would reach out to HIA helpdesk or OIPC contacts for advice (to ensure we remain in compliance with Section 64 requirements).

Compliance Attestation: Finally, we align our privacy risk management with broader compliance. We are aiming for SOC 2 Type 2 certification by 2025, which will externally validate our security controls over time. We have also effectively complied with HIPAA and provincial privacy laws, as noted in the independent audit's conclusion that *"Scribeberry meets industry standards for ePHI security and demonstrates adherence to current standards of data protection"*. This multi-framework compliance approach (HIPAA, PIPEDA, HIA, etc.) ensures we don't operate in a silo; rather, our safeguards meet or exceed the most stringent of applicable rules.

In conclusion, Scribeberry has identified possible privacy risks and implemented a **multi-layered defense** (administrative, physical, technical) to mitigate them. We actively monitor these controls and remain agile in improving our privacy posture. We are confident that the residual risks are low and acceptable given the strong measures in place. Our plan for ongoing monitoring, auditing, and regular PIA review ensures that privacy and security will continue to be maintained throughout the lifecycle of the Scribeberry service. We will keep the OIPC informed as required and will promptly address any issues should they arise, thereby upholding our duty under HIA Section 60 to protect health information.

Section E: Policies, Procedures, and Attachments

In support of this PIA, Scribeberry is providing documentation of its privacy and security policies and related materials. This section lists the key policies, procedures, and other attachments that demonstrate our compliance with the HIA. All policies are current and have been reviewed and approved by the custodian (Dr. Moloo). They are available upon request.

1. HIA-Related Policies and Procedures

(The following table enumerates the major policies/procedures and their relevance. All documents are attached or referenced by title and can be produced to OIPC. Page references indicate where in our policy manual the topic is covered, if needed.)

Policy / Document Title	Description and Key Safeguards Covered	Reference (Available on Request)
Privacy Policy and Confidentiality Policy	High-level organizational privacy policy outlining Scribeberry's commitment to HIA compliance, definitions of roles, and the principle of least privilege. Includes procedures for maintaining confidentiality, steps for obtaining patient notice/consent, and handling of expressed wishes. *Contains:* duty of custodian and affiliates, use/disclosure rules, privacy officer contact info, etc.	Included as "Notice of Privacy Practices – Scribeberry" (Attachment E.1)
Information Security Policy	Detailed security policies covering administrative, physical, and technical safeguards. *Contains:* access control standards (unique IDs, password policy, 2FA), device security, encryption requirements (AES-256, TLS 1.2+), network security (firewalls, IDS), vulnerability management (patching timelines), incident response linkage, and change management for IT systems.	Excerpts included in "Scribeberry Privacy & Security Controls" document (Attachment E.7)
Access Control and User Management Procedure	Procedure for granting, reviewing, and revoking access rights of Scribeberry staff and system users. *Contains:* role definitions, approval process for new access, quarterly access reviews, and audit logging practices. Ensures compliance with HIA 60 by limiting access to authorized individuals only.	Internal procedure document (Attachment E.7, Section 19-25)
Training, Awareness, and Sanctions Policy	Outlines the training program for all affiliates and the sanctions for non-compliance. *Contains:* training frequency (orientation, annual), training content summary, documentation of training completion, and disciplinary process (sanctions) for privacy breaches. Ensures continuous education and accountability.	Attachment E.7 (Privacy Controls) covers training (Sec. 22-24) and sanctions (Sec. 11, Sec. 24).

Incident Response Plan & Breach Protocol	Step-by-step procedure for responding to privacy or security incidents. *Contains:* incident definition, reporting chain (who notifies whom), actions for containment, risk assessment, notification templates (for OIPC, Minister, individuals), and post-incident review process. Also includes form of OIPC Breach Report ready for use. This plan ensures compliance with HIA Section 60.1 (mandatory breach reporting) and is aligned with OIPC's Key Steps guidelines.	"Scribeberry Incident Response Plan" (Attachment E.5)
Records Management and Retention Policy	Describes retention schedules and secure disposal methods for health information. *Contains:* policy that PHI is only retained temporarily (sync purpose) and deletion procedures, backup retention durations, and process for data destruction upon request or end of purpose. Aligns with HIA data minimization and ensures no unauthorized retention.	Incorporated in Security Policy; see Attachment E.7 (Sec. 16).
Information Manager Agreements (with third-parties)	Copies or key excerpts of agreements with Azure, GCP, OpenAI, etc.,. *Contains:* confidentiality clauses, breach notification, use limitation, data location specifics. Shows how vendors are contractually bound to protect PHI.	"Azure Privacy and Security Terms (excerpts)" – Attachment E.8; "OpenAI BAA Agreement (excerpts)" – Attachment E.9.
Client (Clinic) Service Agreement (IMA)	Template of the contract between Scribeberry and client clinics (custodians) for using the service. *Contains:* roles (clinic is custodian, Scribeberry is info manager), obligations of each (Scribeberry safeguarding data, assisting with access requests, etc.), and affirmations of HIA compliance. This is the agreement under which clients engage Scribeberry's services.	"Scribeberry Information Manager Service Agreement (Template)" – Attachment E.10.
Audit and Monitoring Procedure	Internal procedure describing how and when audits are conducted. *Contains:* log review schedule, items to check (user access, admin actions), use of continuous monitoring tools, and reporting of audit results to management.	Attachment E.7 (Privacy Controls doc, Sec. 7-12 covers risk monitoring and review).

	Ensures ongoing compliance verification as per HIA.	
Contingency and Business Continuity Plan	Plan for ensuring availability and restoration of systems in case of outages or disasters. *Contains:* backup procedures, recovery time objectives, emergency contacts, and alternate processing arrangements. Important for protecting integrity and availability of health info (HIA s.60 requires reasonable safeguards against loss).	"Scribeberry Contingency Plan" (Attachment E.6).

(Attachments E.1 through E.10 are provided in a separate document bundle on request. They contain the full text or relevant excerpts of the policies and agreements listed above. Please note some contain proprietary security details and may be marked Confidential.)

2. Other Supporting Documents

- **Privacy Impact Assessment (this document):** This PIA is to be submitted to OIPC. It will also be shared internally and with any interested stakeholders (clients, etc.) as evidence of our privacy due diligence.

- **Compliance Certifications and Audit Reports:** We include our independent Security Risk Assessment Report Executive Summary (Jan 2024) as evidence of third-party evaluation (Attachment E.11). We will also share our future SOC 2 audit report once completed, to demonstrate control effectiveness over time.

- **Patient Information Sheet / Consent Form:** The template given to clinics to inform patients (Attachment E.4 as mentioned). This is not a required HIA form but provided for completeness.

- **OIPC Accepted PIA Reference:** If applicable, reference to Dr. Moloo's Hello Mental Health PIA which initially covered Scribeberry as an amendment. This independent PIA supersedes that context for Scribeberry-specific deployments.

By maintaining and attaching these documents, Scribeberry demonstrates that it has the necessary policies and agreements **in place and in practice**. These documents will be kept up to date (Section D outlines our review frequency).

Before submitting this PIA, we have cross-checked it against the **OIPC Submission Checklist** to ensure all required elements are included:

- Cover letter and cover page ✅
- Complete sections A through E ✅
- All policies and procedures listed and attached ✅
- Sign-off by custodian ✅ (Dr. Moloo has reviewed and signed this PIA).

With all components in order, Scribeberry Ltd. is confident that this PIA provides a full account of the project and its privacy protections. We are prepared to answer any further questions the OIPC may have during the review.

Conclusion: Scribeberry's AI medical scribe system offers innovative benefits to healthcare providers while maintaining rigorous privacy and security standards. Through careful planning, robust safeguards, and adherence to the HIA and related laws, we ensure that patient health information remains protected. We respectfully submit this PIA for your review and look forward to the OIPC's feedback and acceptance.

Attachments: *(List of attached documents as described in Section E, available upon request)*

- Attachment E.1: Scribeberry Notice of Privacy Practices (Privacy Policy)
- Attachment E.2: [Reserved for any additional policy]
- Attachment E.5: Scribeberry Incident Response Plan
- Attachment E.6: Scribeberry Contingency Plan
- Attachment E.7: Scribeberry Privacy & Security Controls Documentation (comprehensive summary of safeguards)
- Attachment E.8: Excerpts from Azure/GCP Information Manager Agreements
- Attachment E.9: Excerpts from OpenAI BAA/Data Processing Addendum
- Attachment E.10: Template Client Service/Info Manager Agreement
- Attachment E.11: Independent Security Risk Assessment Report (Exec Summary)