Disclaimer

This is an English translation of the approved Dutch DPIA. In case of discrepancies, the Dutch version prevails.

Executive Summary – Security Measures Joni RAG

Joni uses Retrieval-Augmented Generation (RAG) to temporarily retrieve personal data from JOIN for better answers, but it does not store this data and does not learn from it.

The main security measures are:

- OAuth authentication: Joni uses exactly the same rights as the logged-in JOIN user.
- TLS-encrypted communication between Joni, JOIN, and the AI model.
- Optional logging with sanitizer: all personal data is removed before storage; users are identified only by an internal item key.
- No additional storage of personal data outside JOIN; AI context disappears immediately after answering.
- Access to audit logs is strictly limited to authorized JOIN administrators.

Conclusion: The RAG functionality does not introduce new privacy risks and remains fully within the GDPR and the European AI Act.

Data Protection Impact Assessment (DPIA) – Joni RAG Functionality

1. Purpose and Lawfulness

The AI chatbot Joni supports users in finding and understanding information within the JOIN system. With the new Retrieval-Augmented Generation (RAG) functionality, Joni can also use personal data solely to provide contextual answers to the user who already has access to it.

This processing is necessary to improve the efficiency and usability of JOIN and remains within the existing purpose limitation of the system. No new personal data is collected or permanently processed outside JOIN. Joni acts as an intelligent search and answer layer.

2. Description of the Processing

1. The user asks Joni a question.
2. Joni performs a real-time query via the secured JOIN REST API, using the rights of the logged-in user.
3. Only records accessible to that user are retrieved (including any personal data).
4. The relevant text fragments are temporarily added as context to the AI model's prompt.
5. The AI model generates an answer, after which the context is not stored or reused.

Important conditions:
- No personal data in embeddings ➜ Only anonymous case type names are pre-embedded.
- No training or fine-tuning of the AI model with personal data.
- All personal data remains exclusively in JOIN and is subject to existing access rights.

3. Data Subjects and Categories of Data

Categories of data subjects:
- Citizens, employees, or other natural persons whose data is stored in JOIN.

Categories of personal data:
- Depending on the question: names, addresses, contact details, case data, and possibly special categories of data as described in existing data processing agreements.

4. Necessity and Proportionality

The processing is necessary to support JOIN users more efficiently in their work. No new data is collected; Joni only uses information that the user could already view through the normal JOIN interface. Processing is limited: personal data is not stored but only used temporarily for the specific answer.

5. Risk Analysis

Risk	Explanation	Likelihood	Impact
Unauthorized access via Joni	AI displays data to someone without rights	Low	High
Unintended disclosure in answers	AI may include irrelevant personal data in an answer	Medium	Medium
Storage or reuse of sensitive data	AI might retain data and reuse it later	Low	High
Profiling or secondary use	Data from AI conversations is used elsewhere	Low	High

6. Mitigating Measures

- ***Access control:*** Joni only retrieves data through the JOIN API, which enforces user rights.
- ***No persistent storage:*** personal data is not stored in Joni or embeddings and disappears immediately after answering.
- ***No AI training on personal data:*** no fine-tuning with sensitive data occurs.
- ***Secure communication:*** API traffic is TLS-encrypted, and logs contain no substantive personal data.
- ***Functional limitations:*** Joni shows only short context fragments, not complete files unless explicitly requested.
- ***Awareness and policy:*** users are informed about what Joni can and cannot do, and that AI answers never provide more rights than the JOIN interface.
- ***Penetration test:*** Joni and its components have been subjected to a penetration test by an external party.

7. Technical and Organisational Security Measures

Authentication and authorization
Joni always logs in as the same user who is authenticated in JOIN. OAuth is used, ensuring Joni never has broader access than the user themselves. This prevents unauthorized access to personal data.

Communication security
All communication between Joni, the JOIN REST API, and the AI components is TLS-encrypted. No sensitive data is placed in URLs or headers that could leak through logging or monitoring.

Logging and audit trail
Chat questions and AI answers may optionally be logged for audit purposes. A sanitizer removes all personal data before storage. To identify the user in the log, no name is used but an item key, which only a JOIN administrator can trace back to a specific user.

Storage limitation
Personal data temporarily offered to the AI model through RAG is not stored or reused. After answering, this data immediately disappears from the AI model's memory.

Administration and access
All administration activities related to Joni are restricted to authorized system administrators. Logs with anonymized data are only accessible to a limited number of authorized staff.

8. Residual Risk and Conclusion

With the above setup, privacy risks are low to acceptable because:
- Only authorized users access the same data as in JOIN.
- Processing is temporary and purpose-bound.
- No new storage or training dataset is created.

This DPIA demonstrates that the use of RAG in Joni complies with the GDPR and the European AI Act. Additional technical and organizational measures have been implemented, such as OAuth-based authentication, encrypted communication, anonymized logging, and strict access control.