



Jurisage AI Privacy Program

Last updated: March 24, 2026

Version 1.2

Main Contact:

Juliano Rabelo

CTO, The Jurisage Group

juliano@jurisagegroup.com

Table of Contents

Introduction.....2

Summary of Generative AI Features in CiteRight..... 2

Data Processing..... 3

 What data is processed?..... 3

 Where is the data processed?..... 3

 Feature Summary and Data Flow..... 3

 Data Location, Retention, and Usage..... 3

 Security Measures..... 3

Privacy and Security..... 4

Version History.....6

Introduction

At Jurisage, security is one of our highest priorities, especially regarding tools that utilize Generative Artificial Intelligence (Gen AI) and Large Language Models (LLMs).

This document outlines our approach to security, data, and Gen AI in our delivery of the CiteRight Services (“**CiteRight**”), a Jurisage product.

Summary of Generative AI Features in CiteRight

As detailed in our Subprocessor List (available at <https://www.citeright.net/legal/terms/subprocessor-list>), we engage with Microsoft Inc. (“**Microsoft**”) as a subprocessor to support certain generative AI features and functionality within our Application Services.

Generative AI-Powered Feature	Availability	Location	Description of Data Transferred to Microsoft
Case law summarization	Opt-out	Case dashboard in the CiteRight web application, available only on app.citeright.net .	Jurisage’s proprietary case law data, selections of summary types, and summary output.
CiteSense	Opt-in	Located within the CiteRight plugin for Microsoft Word.	Document text in order to extract citations.

The table above is a complete list of CiteRight Services product features that utilize Microsoft as a subprocessor, as of the date of this document. Your data is not processed by Microsoft other than via the features listed above.

Data Processing

AI Insights

What data is processed?

The CiteRight AI Insight interface is specifically designed to restrict inputs and prevent customer data from being processed. In the current interface, users can only search for cases that exist in the Jurisage proprietary caselaw database and select how their summary is processed from a dropdown menu of preset options.

Where is the data processed?

Microsoft's servers that run CiteRight AI tools are located in Canada.

CiteSense

Feature Summary and Data Flow

CiteSense is a CiteRight feature that uses the GPT-4o AI model (hosted on Jurisage's Canadian Azure cluster) to extract legal citations from existing Microsoft Word documents and convert them into CiteRight-formatted references. When a user opens a document, the plain-text content is securely sent via API to the Azure cluster, processed by Jurisage's module, and then sent to the secure GPT-4o LLM instance for citation identification and structured formatting.

Data Location, Retention, and Usage

The GPT-4o model is deployed exclusively within Jurisage's Canadian Azure environment. All data processing and storage occur in Microsoft Azure's Canadian data centers. Data has a 30-day temporary retention period. The enterprise OpenAI agreement prohibits the use of client data for training or fine-tuning the model. All processing is isolated per client session, and Jurisage claims no ownership over user data.

Security Measures

Security measures include transport encryption (TLS 1.2 or higher), logical data isolation between customer interactions, prompt injection protection through output validation, and abuse monitoring via Azure with a 30-day retention period that excludes training use. AES-256 encryption is also used.

Privacy and Security

Is customer data retained by Microsoft?

The data submitted to the generative AI features—including document text for CiteSense and search inputs for CiteRight AI Insight—can be securely held by Microsoft for up to 30 days after which time the data is deleted.

What legal protections are in place between CiteRight and Microsoft?

We utilize generative AI services fully hosted on Microsoft Azure, governed by Microsoft's data privacy and security terms. These terms specify the inputs provided to the services and the outputs generated by the services:

- are NOT available to other customers.
- are NOT available to OpenAI.
- are NOT used to improve OpenAI models.
- are NOT used to improve any Microsoft or 3rd party products or services.
- are NOT used to automatically enhance Azure OpenAI models for use in our resource (the models are stateless).

Does CiteRight allow Microsoft to use customer data to train Microsoft services for its own purposes?

No. Microsoft is contractually prohibited from using the document text, items you search, the inputs you select, and the responses you receive to fine-tune or improve their models or services.

Does CiteRight utilize customer data to train or improve its generative AI feature set?

No. The document text, items you search for, the inputs you select, and the responses you receive are not used to fine-tune and/or improve our generative AI feature set.

Does CiteRight use data submitted to its generative AI features to serve other customers?

No. The document text, items you search for, the inputs you select, and the responses you receive are not shared between customers.

Compliance with Subprocessors Policies

Use of the CiteRight Services features powered by Microsoft are subject to compliance with its usage policy, which is available at

<https://learn.microsoft.com/en-us/legal/cognitive-services/openai/code-of-conduct>.

What legal terms apply to my use of generative AI-supported features?

Your services agreement with CiteRight, as well as CiteRight's Generative AI Services Addendum apply to your use of generative AI supported features.

What if my organization wants to disable CiteRight's generative AI features?

Simply contact support@citeright.net or your account manager, and we'll turn the feature off. It can always be reactivated at your request.

Version History

Version #	Date	Changes made
1.2	Mar 24, 2026	Added information about CiteSense
1.1	Jun 14, 2024	Updates to the subprocessor list
1.0	May 29, 2024	Document created