Security & Compliance Overview

JUNIOR
PATENT AI AGENT

Junior's Privacy & Security Highlights

- Junior models do not train on public internet or customer data. All models are custom-trained by the user *de novo* in closed-loop environments with no feedback to training pipelines.
- Each Word document initiates an isolated session. Siloed architecture ensures zero data sharing or co-mingling between sessions or tenants.
- We do not have access to your encrypted materials. Sensitive data is encrypted in transit (TLS 1.3) and at rest (AES-256).
- All users and admins can purge each Junior instance.
- Native Microsoft Word integration eliminates external platform dependencies.

Requirement	Junior Patent AI Agent Solution
Data Protection	ISO 27001 Certified, End-to-End Encryption
Compliance Assurance	SOC 2 Type I & II and GDPR Compliance
Access Control	Multi-factor Authentication (MFA) + Role-based Access Control
Trusted Infrastructure	AWS, Anthropic, OpenAI, Cloudflare, Stripe, Google Cloud
Business Continuity & Reliability	99.95% uptime over the past 12 months and redundant infrastructure

Security White Papers & Documentation

Comprehensive security and compliance documentation is available at junior.law/:

- Data Privacy Agreement
- ISO 27001 Compliance, SOC2 Certification, GDPR Compliance, Sub-Processors

Contacts

Yuri Eliezer, CEO. Email: yuri@junior.law
Mark Burazin, CTO. Email: mark@junior.law

Address:
Equipat IP LLC. 2870 Peachtree Rd NW #484, Atlanta, GA 30305

Data-Secure, Session-Isolated AI for Confidential Work

Uploaded confidential disclosure is **NOT** stored or used to train or engineer anything outside of the session, is encrypted by a local key tied to the local machine, and is removed once the user clears the session or within 30 days of session inactivity.

(a) Each session of Junior is 'siloed' and 'secured' in two ways: 1) it's tied to a local machine, and 2) it's tied to a specific document on the local machine. In this way, you can have the benefits of fine-tuning the output based on uploaded confidential disclosure, but not have any of that leave the specific session and cross-pollinate into other documents (e.g., patents). You can also transfer the encryption key from one machine to another machine (e.g., home computer -> work computer).

(b) As long as any session is active, all user-uploaded data associated therewith is encrypted and stored on AWS / US-Based servers. We do not see nor do we store any information on the actual word document or prompts - ie., we can't see what patent you're writing, we don't know what the inputs/outputs are. The encryption key is tied to the local machine that is running the session, but data is stored in AWS. This data can be deleted by the user once they clear the session (easily from within MS word), or within 30 days of session inactivity. No one at Junior can decrypt this information.

(c) Company-wide prompt libraries are also stored in AWS for each separate Company Account / Practice Group. This data is not used to train any general aspect of Junior or any public LLMs. Rather, it is called upon the private session of the specific document to fine-tune the outputs generated by the LLM for that private session. The Admin user of the Company Account / Practice Group can delete this data anytime. This gives the Practice Group chair control of what data can/cannot be used. Your company, your prompt libraries.

(d) The session does call upon enterprise grade GPT and ANTHROPIC models as selected by user designation. This is limited to transient instances of communication and does not get stored anywhere beyond the transport layer (communicating with the LLMs), where it lives and dies upon completion of the communication.

(e) The user has full control here as to which model is used, and can also provide the API keys to their own hosted models - Junior can plug and play the session right into any private models, while still giving your team the leverage of their company wide fine-tuning and prompt libraries.