ARCHITECTURE COMPARISON

Cloud AI SaaS vs. Librarian On-Premise — *where your firm's data goes under each deployment pattern*

CLOUD AI SAAS

Typical architecture of a hosted legal AI product

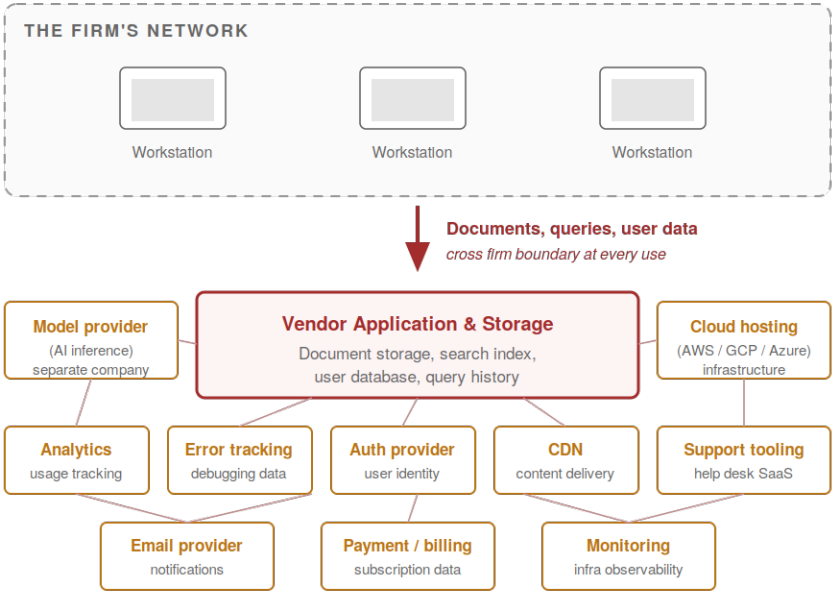

Each box represents a separate organization processing the firm's data.
Most legal AI vendors disclose 8–15 sub-processors in their DPAs.

DATA CROSSES ~10–15 ORGANIZATIONAL BOUNDARIES

- Each boundary is a disclosure point requiring its own contractual relationship.
- Each sub-processor change requires customer notification under GDPR Article 28.
- Breach of any single sub-processor is a breach affecting firm data.

LIBRARIAN

On-premise deployment

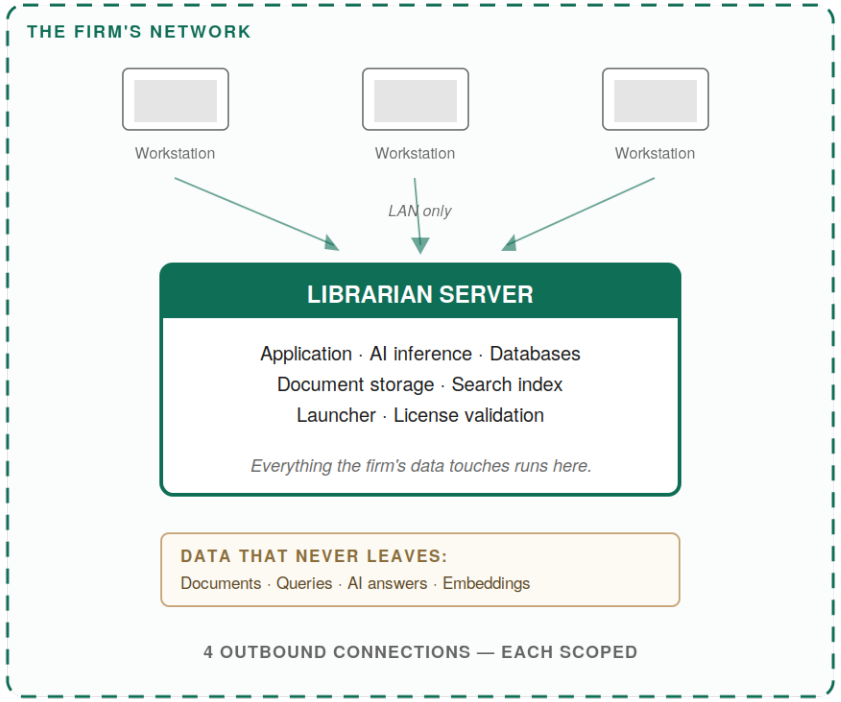

DATA CROSSES 0 ORGANIZATIONAL BOUNDARIES

- Zero sub-processors. The firm processes its own data on its own hardware.
- Four outbound connections, each with a specific and disclosed purpose.
- No possibility of sub-processor breach affecting firm data.

Both architectures are in active use across legal AI in 2026. The choice is not about which is "better" in the abstract — it is about which risk surface matches your firm's confidentiality requirements.