

AI, WITH REVIEW

Assistance that drafts. People decide.

Every AI feature in ZYVORA is off until your administrator switches it on, writes nothing to a record, and ends in something a named person accepts, edits or dismisses.

The rules that never bend

Draft only	The AI orders nothing, prescribes nothing, signs nothing and contacts no patient. For a note draft, the clinician saves their own revision and the review must refer to it.
Off by default	An administrator enables AI for the organisation, and each assistant separately.
Private by policy	The model address must be on your own network unless an operator approves a host. The model runs where you put it.
Identifiers removed first	Names, record numbers, phone numbers, emails, dates and Indian ID formats become tokens before text leaves the application. This reduces exposure; it is not anonymisation.
Checked against the source	Drafts are checked against the words they came from. The checks are blunt: they catch common failures, they do not prove a draft is right.

What it can draft

Consultation note	A transcript with numbered lines; every statement cites the lines it came from; the clinician accepts, edits or rejects each one and signs. Orders the clinician said they would place are proposed separately and need approval.
Notes and documents	SOAP drafts, summaries, coding suggestions (unconfirmed), fields extracted from scanned documents only where the quoted source text really occurs, medicines read from a scanned prescription.
Hospital protocols	Answers drawn only from your own approved protocols, with the passage cited.
Results	A plain-language explanation draft for one verified, non-critical result; any number not in the result is refused; never sent by the system.
Four assistants	Discharge, claims, handover and patient-reply drafts. They read the record as the person asking, cite what they read, and propose; a person accepts or dismisses. Urgent patient messages get no draft at all.

Honest limits. No model has been clinically validated and no accuracy figure is claimed. The ED triage suggestion is a small trained classifier with safety rules, evaluated on synthetic data only. We fine-tuned and tested a consultation-draft model in three rounds on our own GPU against a pass rule written in advance; none passed, and the results are published in the repository. No speech model has been verified for Indian languages.