

IT AND SECURITY QUESTIONS

The questions your IT team **will ask.**

QUESTION	ANSWER
Where does it run, and where does data live?	On your infrastructure: Docker Compose, Kubernetes (Helm chart) or k3s. One Rust service, one PostgreSQL database (14 to 18), a React web app. The data stays in your database.
How do users sign in?	OpenID Connect authorization code with PKCE through your identity provider; ZYVORA never sees a password. Sessions are server-side, stored hashed, 30 minutes idle and 8 hours at most. An optional edge can require a client certificate.
Who can see and do what?	Row-level security is forced on every tenant-owned table, and the application connects as a role that cannot bypass it. A cross-tenant id returns "not found". Thirteen roles, deny by default; a denied attempt is itself audited.
Is there an audit trail?	Yes: every read and change of a record writes an append-only, hash-chained entry in the same transaction (metadata only), with a verify endpoint. It is tamper-evident, not tamper-proof against a database owner.
How do we connect other systems?	HL7 v2 and ASTM intake (optionally over mutual TLS), read-only FHIR R4, webhooks, and read-only API keys. The integration sheet lists what has been run against real equipment.
How do we operate it?	A Prometheus metrics endpoint (off unless a token is set) with starter alert rules; backup, restore and upgrade commands in the installer (upgrade rolls back on failure). Scheduled backups and recovery drills are yours to arrange.
Can we use our own AI model?	Yes, any OpenAI-compatible server on a private address; off by default; identifiers removed first. Zyvora's Nuvora is the intended backend and its wire contract has been checked.
Is it certified?	No. It is not certified or attested against HIPAA, GDPR, ISO 27001, SOC 2, IEC 62304, NABH or JCI. A gap list is published. There has been no independent penetration test, no load test and no screen-reader testing.

Pilot it with synthetic data, in a separate environment, with your own identity provider, and bring your security team in early. Vulnerabilities are reported privately to the security contact on the security sheet.