




PRAVERSE TECH PVT LTD

ENTERPRISE AI · HEALTHCARE · PHARMA · PHOTONICS

Engineering human-centred intelligence

Enterprise AI for healthcare, pharma and regulated industries — and the high-speed interconnect layer for India's sovereign AI data centres.

praversetech.com · inquiry@praversetech.com · +91 93132 47264 · Vadodara, Gujarat, India



One company, four focused business lines

2023

Idea began

Vadodara

Headquarters · Gujarat, India

2026

Praverse Tech Pvt Ltd incorporated

1

Healthcare product in development (MENNIE)



Enterprise AI & Agents

Agents, copilots, RAG search and production ML that live inside real workflows.

Agents · RAG · MLOps



Pharma & Regulatory AI

FDA 483 responses, CAPA, data integrity and inspection readiness — every output expert-reviewed.

GMP · 21 CFR 211 · ICH



Healthcare AI

MENNIE, our voice-first patient assistant, plus imaging and screening support.

MENNIE · Imaging · Edge AI



NEW

Photonics & AI Data Centres

High-speed copper and optical interconnect for India's sovereign AI data centres, with our USA technology partner.

800G · 1.6T · India partner

Built in India. Delivered to India, the US and Europe



India-first engineering, built to global standards for domestic and export markets.



Headquarters

Vadodara, Gujarat, India



Markets

India · United States · Europe



Technology partner

Our USA technology partner —
interconnect



Pharma advisory

SWATI Spentose · V-Group, est.
1952

Most enterprise AI stalls between the demo and the workflow

Where AI usually stalls

-  The model works in a notebook — and nowhere else
-  No evidence trail, so regulated teams cannot sign off
-  Built without domain experts, so it answers the wrong question
-  Black-box outputs that teams must trust, not verify
-  No plan for monitoring, drift or changing data

How we build instead

-  We learn the workflow before we design the system
-  Every output links to the source it came from
-  Human review is designed in, not bolted on
-  Audit trails by default: who asked, what surfaced, who decided
-  Production deployment with monitoring from day one

AI should not replace experts. It should amplify them

“

The system drafts. Your team decides.

That hierarchy is not a policy layered on top of the architecture — it is the architecture.



Evidence-based review

Every suggestion links to the document, record or observation it came from. No invented answers.



Audit trails by default

Who asked, what the system surfaced and what the expert decided — captured for every interaction.



Expert judgment first

The system analyses, drafts and flags. Qualified people approve. High-stakes decisions stay in human hands.

Enterprise AI systems, not demos

What clients hire us to design, build and run in production.



Machine learning & applied AI

Full ML lifecycle — model design to monitored production, with MLOps built in.

NLP · Predictive ML · MLOps



AI agents & copilots

Domain agents that analyse documents, write reports, search knowledge and automate steps.

Agents · Copilots · Automation



Document intelligence

RAG and knowledge systems that turn unstructured files into sourced answers.

RAG · Knowledge graphs · Search



Vision & medical imaging

Computer vision for fundus and ultrasound imaging, inspection and multimodal systems.

Vision · Imaging · Multimodal



Decision intelligence

Dashboards and predictive decision support that give teams clear, explainable answers.

Dashboards · Forecasting




Federated & edge AI

Privacy-preserving training and on-site inference where data cannot leave its home.

Federated · Edge · DPDP-aware

Agents that do the work — with an expert in the loop

Every consequential step is reviewed, traceable and reversible.



01

Document review agents

Parse, classify and cross-reference thousands of pages, with evidence attached.

02

Regulatory & research agents

483 response drafting, guidance tracking and compliance support.

03

Data analysis agents

Ask enterprise data questions in plain English; get explainable answers and charts.

04

Knowledge search agents

One question across wikis, drives, tickets and archives — one grounded answer.

Quality systems, with a second brain

In a GMP environment, a wrong answer from an AI system is worse than no answer at all.

So every observation maps to evidence, every CAPA draft traces back to its source document, and every output is a starting point for qualified expert judgment — never a substitute for it.

DESIGNED AROUND

USFDA 21 CFR 211

ICH Q9

ICH Q10

ALCOA+



Evidence-linked

Each suggestion cites its source observation or record.



Inspection-ready trail

Full interaction history captured automatically.



Approved by QA

Quality professionals sign off every output before use.

Eight modules for the quality system

Each output is traceable, reviewable and grounded in evidence.



FDA 483 response assist

Faster, evidence-backed drafts using enforcement-trend patterns.



OOS / OOT analysis

Speeds up root-cause analysis for out-of-spec and out-of-trend results.



Data integrity review

Checks records against ALCOA+ and flags possible anomalies.



Document intelligence

SOP drafting, search and gap detection across the quality system.



CAPA reasoning

Structured root cause and corrective-action drafts with evidence trails.



Quality risk management

Risk signals across deviations, OOS, complaints and recalls.



Human error investigation

Separates systemic causes from individual ones.




Inspection readiness

Continuous readiness scoring, so audits stop being a crisis.

Highlighted: the usual entry points.


Five stages, fully traceable

Every AI interaction in a regulated environment follows the same pipeline.



All AI outputs are drafts that need final review by qualified personnel. No diagnosis, compliance guarantee or automated decision-making is implied.

MENNIE: healthcare that listens



Medical Empowerment and Neural Navigation Intelligence Engine — a voice-first assistant that meets patients at the door, the kiosk or on their phone, and hands clinicians a clear, structured picture.



Voice-first intake

Conversational symptom intake — kiosk, mobile or in person, in multiple languages.



Doctor matching support

Routes patients to the right specialty and the right appointment.



Smart reminders

Medication and follow-up nudges that keep care plans on track.



Clinician-ready summaries

Consultations start with context, not questions.

MENNIE supports care workflows under clinical oversight. It does not diagnose, treat or replace medical professionals.

A staged path from pilot to institution

PHASE 1 · Private beta

Select healthcare partners; real-world workflow testing under close clinical supervision.



Kiosk workflows

Self-service intake and navigation stations for hospital lobbies — voice-first and multilingual.

PHASE 2 · Pilot workflows

Intake, routing and clinician-handoff pipelines in live settings.



Privacy by architecture

Federated learning and edge inference keep health data on-site, encrypted and auditable (DPDP-aware).

PHASE 3 · Institutional rollout


Kiosk deployment, edge inference and multilingual support at operational scale.



Next: assistive robot (R&D)

An assistive form factor, tested in digital-twin simulation before any physical use.

Systems we have already built



Research behind it: Peer-reviewed work in medical imaging, edge AI and simulation. Full publication list available on request.

The physical layer for India's sovereign AI

AI models run on GPUs — but GPUs only work as a cluster when thousands of high-speed links connect them. We are building the India capability to supply, design in and support that connectivity layer.



India partner

Our USA technology partner builds the interconnect; we bring it to Indian projects.




Design-in & integration

Indian architecture, selection, integration and support for AI data-centre projects.




Roadmap to localisation

Indian testing and qualification first; local assembly as volumes justify it.



India's data-centre capacity is set to multiply by 2030

India data-centre capacity, GW (actual vs. 2030 forecasts)



24×

Growth in AI-dedicated capacity, 2025–2030: 275 MW → 6.5 GW (Wood Mackenzie)

38,000+

GPUs onboarded for common compute under the IndiaAI Mission (MeitY, Mar 2026)

₹10,372 cr

IndiaAI Mission outlay over five years (Union Cabinet, 2024)

Sources: MeitY (2020 and 2025 actuals, Mar 2026); Savills India (Jul 2025); Colliers (May 2025); Wood Mackenzie via IANS (Jul 2026). Forecasts differ widely — treat 2030 as a range, not a point.

AI build-outs will land first in a few hubs



~52%

of listed capacity is in Mumbai & Navi Mumbai
— 790 of ~1,520 MW

~65%

of installed IT load sits in Maharashtra and
Tamil Nadu (Wood Mackenzie)

~400 km


Vadodara HQ to Mumbai — close to India's
largest data-centre cluster

What it means for us: sovereign AI clusters will be built in these hubs first, so that is where interconnect demand — and our first deployments — will land.

Source: MeitY reply in Lok Sabha (Mar 2026), operational capacity by city; Pune and Kolkata together ~15 MW. Map boundary: Survey of India outline (DataMeet).

Copper inside the rack, optics across the hall

Every GPU server connects to the fabric. Link reach decides whether copper or optics is used.



Reach values are typical maximums at 400G–800G and vary by product and speed; the AEC figure is our partner's 9 m 800G specification. Indicative, for architecture discussions.

Interconnect from 100G to 1.6T

**Direct attach copper (DAC)**
Passive copper for ultra-low-latency, in-rack links.

**Active copper (ACC)**
Linear redriver signal conditioning for extra reach.

**Active electrical (AEC)**
Retimer/DSP copper — 800G up to 9 m.

**Active optical (AOC)**
Integrated-transceiver cables for dense, longer runs.

**Optical transceivers**
Pluggable modules for multi-mode and single-mode links.

**Fibre cable assemblies**
Patch cords, trunks and high-density breakouts.

SPEED ROADMAP



Form factors: SFP · QSFP · QSFP112 · QSFP-DD · OSFP · OSFP-XD · Ethernet and InfiniBand · Custom assemblies on request

Global technology, Indian capability, sovereign data centres




ROADMAP



Why sovereign: data stays under Indian law (DPDP Act 2023) · resilient domestic supply · engineers and support in India

From problem to production in four stages



Tell us what you are trying to solve. We will tell you honestly whether AI is the right answer — and send a written plan within two business days.

Industry depth, engineering velocity



Three decades of pharma depth

Advisory from a 30-year specialty-pharma career: GMP floors, regulatory strategy and global launches.



Peer-reviewed foundations

Applied work anchored in published research in imaging, edge AI and simulation.



Human-in-the-loop by design

Built for industries where a wrong answer has consequences: traceable, reviewable, reversible.



India-first, globally deployable

Indian engineering talent, built to global standards for domestic and export markets.



Pharma advisory backed by SWATI Spentose
V-Group, est. 1952 · USFDA, EU-GMP, PMDA Japan,
WHO-GMP approvals · 3 units in Vapi, Gujarat

AI engineering, backed by deep pharma expertise



Pratham Shrivastav

Founder & CEO
Praverse Tech Pvt Ltd

- Began Praverse in 2023 as a research question; incorporated it in 2026
- Peer-reviewed work in imaging, edge AI and simulation
- Leads the full stack: pharma AI, MENNIE and photonics



Vishal Jajodia

Strategic Advisor
MD, SWATI Spentose

- 30+ years in specialty pharma, from formulation to global regulatory strategy
- Guides our sovereign AI data centre and photonics direction
- IDMA-APTAR Innovation of the Year 2022



Ashwani Shrivastav

Managing Director
Nexgen Pharma Solutions Pvt Ltd

- Pharma quality expert: QMS, QA, QC and DMF filings in eCTD format
- Inspection support: USFDA, EU-GMP, WHO-GMP, EDQM, ANVISA, MFDS (Korea)
- Brings hands-on GMP and inspection practice to our pharma AI

Four ways to start working with us



Enterprise AI

Custom agents, copilots, RAG systems, vision pipelines and production ML.



Pharma & regulatory AI

483 responses, CAPA, data integrity, SOP intelligence, inspection readiness.



Healthcare AI

MENNIE pilots, patient-flow intelligence, imaging and screening support.



AI data-centre interconnect

Architecture reviews, design-in and supply for 400G–1.6T AI clusters.

1. Share the problem, domain and timeline

2. Receive a written plan within 2 business days

3. Start with a scoped discovery or pilot



PRAVERSE TECH PVT LTD

Let's build what your industry needs

From AI agents to healthcare platforms, regulatory intelligence and AI data-centre connectivity — we turn hard problems into deployed systems.



inquiry@praversetech.com



+91 93132 47264



praversetech.com



Vadodara, Gujarat, India

