EMBLEM

SARGVISION

*Sovereign AI for Eastern India*

DETAILED PROJECT REPORT · MAY 2026

# AI Centre of Excellence

## for Eastern India

পূর্ব ভারতের জন্য কৃত্রিম বুদ্ধিমত্তা উৎকর্ষ কেন্দ্র

*A blueprint for sovereign regional-language AI capacity, anchored in Bengal.*

**Published by SARGVISION Intelligence Pvt. Ltd.**

In partnership with academic, industry, and civic stakeholders across Eastern India

V1.0

MAY 2026

~100 PP

# Preface from the Founder

The first time I sat in a room and listened to colleagues in Bangalore describe the artificial-intelligence roadmap for India, I realised something that has stayed with me through the twelve months that produced this document: not a single sentence of that roadmap mentioned my home. Eastern India was an absence. Bengal — the region that produced India's first Nobel laureate, its first IIT, its first modern university, and the world's fifth-most-spoken language — was simply not in the conversation that will determine where the country's twenty-first-century jobs, capital, and intellectual capital concentrate.

I have spent the last year asking why. Some of the answers are well-known and well-documented. The 1990s and 2000s were the decades in which the IT economy was built, and West Bengal — for reasons of state-level policy choice, central-government industrial-policy alignment, and the deeper political economy of the period — was largely outside that build. The state's IT Policy arrived in 2018, approximately two decades after the states that now dominate India's technology economy had established their competitive positions. The brain drain of the elite-institution graduates — IIT Kharagpur, Jadavpur, IIEST Shibpur, NIT Durgapur, JGEC Jalpaiguri — has been a continuous regional cost over those same decades. Some of these losses are recoverable. Some are not.

What is recoverable, and what this document is about, is the next decade. The AI economy is not yet built. The institutional architecture for Bengali-language AI is not yet in place. The producers of the foundation models, the datasets, the safety frameworks, and the regional capacity have not yet been named. The window in which Bengal can move from absence to leadership in the language-AI build-out for 27 crore Bengali speakers is open today and is closing fast. Bangladesh — a sovereign nation with one-fifth of West Bengal's GDP but a sharply focused national AI strategy — has already begun building Bengali AI in Dhaka. The institution that consolidates the foundational Bengali-AI infrastructure is the institution that anchors the AI economy of the world's fifth-most-spoken language. That institution is going to exist. The only open question is its location.

This document is the published case for that institution being in North Bengal, anchored in Siliguri, operating across Eastern India, and accountable to a public the Centre is committed to serve

under outcome-linked discipline for the next twenty-five years. The case is set out in twelve chapters and thirteen annexures. Every number has a source. Every commitment has a measurable target. The Centre's Editorial Board, the Open Asset Council, the Ethics Office, the annual Eastern India AI Index, and the Section-8 corporate form together compose what we hope is read as institutional seriousness, not institutional ambition.

I am personally accountable for the institution this document describes. The Centre's leadership will sign its commitments openly. The first edition of the Eastern India AI Index, due May 2027, is the first concrete artifact against which our seriousness will be measured. The first proof-of-concept programmes — in agriculture, manufacturing, and healthcare, sequenced in Annexure N §N.9 — are the second.

I am asking the reader for two things. First, read this document as a publication and not as a petition; engage with it on its merits as a piece of institutional argument. Second, consider what specifically the reader's institution or office can contribute — land, policy, co-investment, partnership, fellowship, sponsorship, foreword, or a written response — and respond directly to the Centre at the addresses set out in the colophon.

The decision before us is, at this point, almost entirely a decision about institutional will. The intellectual capital is here. The economic opportunity is here. The youth are here. The language is here. What is not yet here is the institution that converts these inputs into outputs at the scale and discipline the moment requires. The Centre that this document describes is the institution we are committed to building. It will exist whether the reader agrees or not. With the reader's contribution, it will exist sooner, be larger, and serve its public better.

Bengal's window is closing. The Centre is the lever.

*— Abhishek Gupta Founder, SARGVISION Intelligence Pvt. Ltd. Kolkata · May 2026*

# How to Read This Document

This Detailed Project Report is approximately 190 pages, structured as 12 chapters and 13 annexures. It is designed to serve five distinct kinds of reader, each with different priorities and different

time budgets. The reading paths below are the Centre's own recommendation for how to navigate the document efficiently. The cover-to-cover read is approximately three to four hours; a focused read against any of the five paths below can be completed in 30 to 60 minutes.

The reader's-guide convention used throughout: chapter numbers refer to the body; annexure references use letter codes (A through N); front-matter references use the prefix F (F1 through F11). The Table of Contents (F7) lists every section with its page reference; the List of Figures (F8) and List of Tables (F9) provide visual indexes.

- **For the state-government decision-maker — about 30 minutes.** Start with the Preface from the Founder (F1) to understand the institutional posture. Read Chapter 1 in full — it is the load-bearing executive summary and contains the specification of what the Centre needs from the state in §1.4. Read Chapter 8 (Financial Plan) for the budget architecture and the funding mix, paying particular attention to §8.4 (funding mix), §8.6 (institutional finance discipline), and the five-year scenario table (Table 8.1). Sample Annexure N (Sector Use-Case Catalogue) for the economic-impact case across the seven industry groupings the Centre commits to. Close with §1.5 for the document's overall request of the reader.

- **For the central-government / IndiaAI Mission reader — about 45 minutes.** Read Chapter 1 in full. Read §3.5 (What the Centre Contributes to the National AI Architecture) — it is the explicit articulation of the Centre's contribution to IndiaAI Mission's seven layers, to NEP §22, to the West Bengal IT Policy, and to the Banglar Shiksha portal. Read Chapter 4 §4.1-4.3 for the architectural framing. Read Annexure I (Open Asset Commitment Ladder) for the five-year open-asset release calendar that the Centre commits to the national commons. Read Annexure M (Eastern India AI Index methodology) for the editorial independence and pre-registration safeguards. Sample Chapter 5 (Programmatic Pillars) for the operating model.

- **For the academic-partner reader (Vice-Chancellor / Dean / Faculty Head) — about 45 minutes.** Read the Preface (F1). Read Chapter 4 §4.5 (Talent Architecture) and §4.6 (Partnership Architecture). Read Chapter 5 §5.1 (Research Pillar) and §5.2 (Training Pillar). Read Annexure I (Open Asset Ladder) — the open weights and datasets the Centre commits to release determine what is available for academic collaboration. Read

Chapter 7 (Institutional Framework) for the joint-appointment scheme, fellowship structure, and HR plan. Read Annexure F (Faculty / Founder CVs index) and Annexure E (Partner Institution Profiles) for institutional alignment.

- **For the industry / CSR / philanthropic-foundation reader — about 45 minutes.** Read Chapter 1 in full. Read Annexure N (Sector Use-Case Catalogue) — the most detailed industry-facing content in the document, organised across seven sector groupings with named comparable deployments and investment-to-impact ratios. Read Chapter 5 §5.4 (IP & Commercialisation Pillar) for the commercial pathway. Read Chapter 8 §8.4 (Funding Mix) for the CSR and earned-revenue posture. Read Chapter 10 (Monitoring & Evaluation) — particularly §10.3 (Third-Party Audit) — for the outcome-discipline framework that gives CSR commitments their measurability. For deeper engagement, read Annexure M (Eastern India AI Index) — the recurring institutional output that anchors the Centre's external accountability.
- **For the journalist, civil-society, or interested-public reader — about 30 minutes.** Read the Preface from the Founder (F1) and Chapter 1 in full. Read Chapter 3 §3.1-3.2 (Vision and Three Mandates). Scan the visuals via the List of Figures (F8). Read Chapter 11 (Inclusion, Gender, Ethical Safeguards) for the public-interest commitments. Sample Annexure M (Eastern India AI Index) for the institutional accountability framework. Close with Chapter 12 (Conclusion & Invitation). The bibliography in the citation ledger (Annexure-level) provides full source attribution for every numeric claim in the report.

A note on the document's overall posture. This is a published institutional report. It is not addressed to any single office. References to alignment with national or state policy frameworks should be read as descriptions of institutional contribution, not as applications under those frameworks. The Centre's design is open to multiple kinds of engagement; the reading paths above are recommendations, not prescriptions. Readers are encouraged to follow their own interest through the document and to respond directly to the Centre at the addresses set out in the colophon.

# Acronyms & Abbreviations

The following is a comprehensive list of acronyms and abbreviations used throughout the SARGVISION AI Centre of Excellence Detailed Project Report. This list serves as a reference to ensure

clarity and consistency across the document.

| ACRONYM | FULL FORM |
|---|---|
| AISHE | All India Survey on Higher Education |
| AI4ICPS | Artificial Intelligence for Interdisciplinary Cyber-Physical Systems |
| AI4Bharat | Artificial Intelligence for Bharat |
| ARTPARK | AI & Robotics Technology Park |
| BAU | Business As Usual |
| CoE | Centre of Excellence |
| CSR | Corporate Social Responsibility |
| DoHE | Department of Higher Education |
| DPDP | Digital Personal Data Protection |
| DPR | Detailed Project Report |
| ECE | Electronics and Communication Engineering |
| FY | Fiscal Year |
| GeM | Government e-Marketplace |
| GFR | General Financial Rules |
| GoI | Government of India |
| GoWB | Government of West Bengal |
| GPU | Graphics Processing Unit |
| HPC | High-Performance Computing |
| IAS | Indian Administrative Service |
| IIIT-H | International Institute of Information Technology, Hyderabad |
| IISc | Indian Institute of Science |
| IIT-KGP | Indian Institute of Technology, Kharagpur |
| IndiaAI | India Artificial Intelligence |
| IT | Information Technology |
| ITeS | Information Technology enabled Services |
| LLM | Large Language Model |
| MASTI | Multilingual AI for Social and Technological Inclusion |

| | |
|---|---|
| MoE | Ministry of Education |
| MoSPI | Ministry of Statistics and Programme Implementation |
| MoU | Memorandum of Understanding |
| MSME | Micro, Small and Medium Enterprises |
| NASSCOM | National Association of Software and Service Companies |
| NCERT | National Council of Educational Research and Training |
| NCF | National Curriculum Framework |
| NEP | National Education Policy |
| NITI Aayog | National Institution for Transforming India |
| NSDC | National Skill Development Corporation |
| NSSO | National Sample Survey Office |
| PIB | Press Information Bureau |
| PLFS | Periodic Labour Force Survey |
| PMKVY | Pradhan Mantri Kaushal Vikas Yojana |
| R&D | Research and Development |
| RBCDSAI | Robert Bosch Centre for Data Science and Artificial Intelligence |
| RBI | Reserve Bank of India |
| RTI | Right to Information |
| RUSA | Rashtriya Uchchatar Shiksha Abhiyan |
| SARGVISION | SARGVISION Intelligence Pvt. Ltd. |
| SC | Scheduled Castes |
| ST | Scheduled Tribes |
| OBC | Other Backward Classes |
| STEM | Science, Technology, Engineering, and Mathematics |
| TIFR | Tata Institute of Fundamental Research |
| UDISE+ | Unified District Information System for Education Plus |
| VC | Venture Capital |

| | |
|---|---|
| WB | West Bengal |
| WBBSE | West Bengal Board of Secondary Education |
| WBCHSE | West Bengal Council of Higher Secondary Education |
| WBHED | West Bengal Higher Education Department |

This table provides a quick reference to understand the various abbreviations used throughout the report, facilitating ease of reading and comprehension.

# Glossary of Key Terms

This glossary provides definitions for key terms related to AI, education policy, and institutional finance as used in the SARGVISION AI Centre of Excellence Detailed Project Report. Each entry offers a concise explanation, adhering to the Indian-English government register.

*Centre of Excellence (CoE)*

A Centre of Excellence is a hub for advanced research and development in a specific field, fostering innovation and collaboration among academia, industry, and government. The SARGVISION AI CoE aims to become a leading institution in AI research for Eastern India, aligning with national and regional policies.

*Convergence Framework*

This framework refers to the alignment of multiple policies and initiatives to achieve common goals. The SARGVISION initiative sits at the convergence of NEP 2020 §22, the IndiaAI Mission, and the West Bengal IT & ITeS Policy 2018-23.

*Citation Ledger*

The Citation Ledger is the canonical database used to track and verify all sources cited within the Detailed Project Report. It ensures that all data and claims are accurately referenced, maintaining the integrity of the report.

*Code-switch*

Code-switching involves alternating between two or more languages or dialects within a conversation. It is a significant phenomenon in multilingual societies like India and is considered in the development of language models and AI tools.

*Data Protection (DPDP) Consent Regime*

This refers to the legal framework established by the Digital Personal Data Protection Act 2023, which mandates obtaining informed consent from individuals for processing their personal data. Compliance with this regime is crucial for the ethical handling of biometric voice data in the project.

*District Hub*

A district hub serves as a local center for implementing and managing project activities within a specific district. It facilitates coordination among stakeholders and ensures the effective delivery of project services.

*Foundation Model*

Foundation models are large-scale AI models trained on vast datasets and capable of performing a wide range of tasks. They serve as the base for developing more specialized models, such as those for language translation or speech recognition.

*GFR-compliant Procurement*

Procurement processes that adhere to the General Financial Rules (GFR) 2017 of the Government of India, ensuring transparency, fairness, and efficiency in the acquisition of goods and services for the project.

*Index of Inclusion*

This index measures the extent to which diverse groups, including SC/ST/OBC communities and linguistic minorities, are represented and supported within the project. It guides efforts to promote equity and access.

*Indic Language Model*

An AI model specifically designed to process and understand languages spoken in the Indian subcontinent, such as Bangla, Hindi, Tamil, and more. These models are critical for developing applications that cater to the linguistic diversity of the region.

*Industry Fellow*

An industry fellow is a professional from the private sector who collaborates with the Centre of Excellence on specific research projects, bringing practical insights and expertise to the academic environment.

*Language-model Evaluation*

The process of assessing the performance of language models based on criteria such as accuracy, fluency, and cultural appropriateness. Evaluation ensures that models meet the required standards for deployment.

*Lakh / Crore*

Indian numbering terms used to express large numbers. One lakh equals 100,000, and one crore equals 10 million. These terms are used in financial reporting within the project.

*Multidisciplinary Education (NEP §22)*

A principle from the National Education Policy 2020 encouraging the integration of multiple disciplines in education to foster holistic learning. The AI CoE incorporates this approach by combining AI with fields like linguistics and social sciences.

*Outcome-Linked Discipline*

An educational approach where the curriculum and assessment are directly tied to specific learning outcomes. This ensures that students acquire relevant skills and knowledge applicable to real-world scenarios.

*Open-weight Release*

The practice of making the weights of an AI model publicly available, allowing researchers and developers to build upon existing work. This promotes transparency and collaboration within the AI community.

*Outreach Lab*

A facility dedicated to engaging with the broader community through workshops, training sessions, and demonstrations. These labs aim to increase awareness and understanding of AI technologies among diverse audiences.

*Partner Institution*

An educational or research organization that collaborates with the Centre of Excellence on various projects, sharing resources and expertise to achieve mutual goals.

*Patent Strategy*

A plan for protecting intellectual property generated by the Centre of Excellence. This includes identifying patentable innovations and navigating the patent application process.

*Pillar Architecture*

A structural framework organizing the Centre of Excellence's activities into distinct pillars, such as research, training, and outreach. Each pillar focuses on a specific aspect of the Centre's mission.

*Programmatic Chapter*

A section within the Detailed Project Report outlining a specific program or initiative, detailing its objectives, activities, and expected outcomes.

*Research Fellow*

A researcher affiliated with the Centre of Excellence, contributing to its projects through academic inquiry and experimentation. Fellows are often early-career academics or postdoctoral researchers.

*Sovereign Dataset*

A dataset created and maintained by a national entity, ensuring data sovereignty and compliance with local regulations. Such datasets are crucial for training AI models that respect regional specificities.

*Talent Ramp*

A structured program designed to develop and nurture skilled professionals in the field of AI. It includes training, mentorship, and placement support to ensure a steady pipeline of talent.

*Tokeniser*

A component of language processing systems that breaks down text into tokens, or individual units, such as words or phrases. This process is essential for analyzing and understanding natural language.

*Vernacular Speech Corpus*

A collection of speech data in regional languages and dialects, used to train and evaluate speech recognition systems. This corpus helps improve the accuracy and relevance of AI applications for diverse linguistic communities.

This glossary provides foundational knowledge for understanding the complex and interconnected concepts within the SARGVISION AI Centre of Excellence project.

# Methodology Note — How This Document Was Produced

This Detailed Project Report (DPR) for the SARGVISION AI Centre of Excellence was developed through a rigorous and transparent process to ensure accuracy and reliability. The methodology employed involved a comprehensive review and synthesis of various source documents and structured inputs from multiple contributors.

The foundational materials for this report were sourced from the SARGVISION archive, including key documents such as the Vision Whitepaper, Master Project Blueprint, Founders' Manifesto, Executive Whitepaper, Bengal AI Vision, NB CoE Proposals, Risk Register, and Timeline documents. These documents provided the necessary context and data to inform the report's content and direction (SARGVISION Vision Whitepaper, 2026).

To support the development of the DPR, a knowledge base ingest was conducted, wherein relevant data and insights from the aforementioned documents were extracted and organized for reference. This ensured that the information used in the report was both current and comprehensive. The multi-agent drafting pipeline played a critical role in the report's production. This pipeline included multiple writers and editors, each contributing their expertise to different sections of the report. Citation discipline was strictly enforced throughout the process, with every factual claim and reference backed by a corresponding citation from the source materials (Master Project Blueprint §3.2).

The drafting process also adhered to a defined house-style, which mandated consistency in language, tone, and formatting. This included the use of Indian-English government register and specific conventions for numbering, citations, and phrasing. Human-in-the-loop review gates were established at various stages of the drafting process to ensure quality control and adherence to the house style.

AI-assisted tools were employed to support the drafting process, primarily for tasks such as grammar checking and style enforcement. However, all substantive content decisions were made by human writers and editors, ensuring that the final report reflects a human-authored and reviewed document.

The bibliography and full source register used in the preparation of this report are provided in Annexure O. This ensures transparency and allows readers to verify the sources of information used throughout the report.

# Table of Contents

## Front matter

## Part I — The Case

# Part II — The Architecture

## Part III — The Plan and the Promise



## Annexures

# List of Figures

Numbered figures across the body of the report.

# List of Tables

Numbered tables across the body of the report and annexures.

# 01

# Executive Summary

নির্বাহী সারসংক্ষেপ

# Chapter 1 — Executive Summary

## 1.1 The Capability Trap

West Bengal possesses every raw input for AI-economy leadership — a Rs 18.79 lakh-crore economy, 27 crore Bengali speakers worldwide (the world's fifth-most-spoken language), world-class technical universities (IIT Kharagpur, ISI Kolkata, Jadavpur, IIEST Shibpur), a three-border strategic location, and 89 lakh MSMEs ready for productivity AI. And yet it has zero dedicated AI Centres of Excellence, zero AI unicorns, fewer than twenty AI-focused startups, and per-capita income at 83.7 percent of the national average. In the first quarter of 2025, the state attracted Rs 22 crore in startup funding — less than a single seed round in Bangalore.

This is the *Capability Trap*: the structural condition in which a region possesses every asset for economic transformation but lacks the single institutional mechanism that converts those assets into observable output. The Centre is the institutional intervention that breaks the trap. It is not a research laboratory. It is an industrial transformation engine, anchored in North Bengal, designed to ship measurable economic outcomes within thirty-six months of operations.

The window for that intervention is narrow. Bangladesh — a country with one-fifth of West Bengal's GDP — has already launched its national AI strategy and begun building Bengali-language AI from Dhaka. The question is no longer whether Bengali AI will be built. The question is whether it will be built in Siliguri or in Dhaka, and whether the seven-figure jobs and the IP that the build produces will accrue to Eastern India or to a city elsewhere.

> **35 crore people. Zero AI Centres of Excellence. Zero AI unicorns. Fewer than fifty AI-focused startups across the entire Eastern Indian geography.** The region that produced India's first Nobel laureate, its first IIT, and its first modern university now has no institutional presence in the technology that will define the twenty-first century.

The cost of delay compounds. Every year that passes, an estimated 36 percent of West Bengal's elite-institution graduates leave the state — for Bangalore, Hyderabad, Pune, and abroad. A single JGEC Jalpaiguri graduate placed at Microsoft in 2024 at Rs 54 lakh per year is a celebration on the college website and a quiet loss to the regional economy. Each departing graduate reduces the local talent pool, deters AI-focused employers from establishing operations in Bengal, reduces opportunities for the next cohort, and causes more departures. The trap perpetuates itself unless an institution intervenes.

Fig. 1.2 *The Capability Trap — six inputs without an institutional bridge*

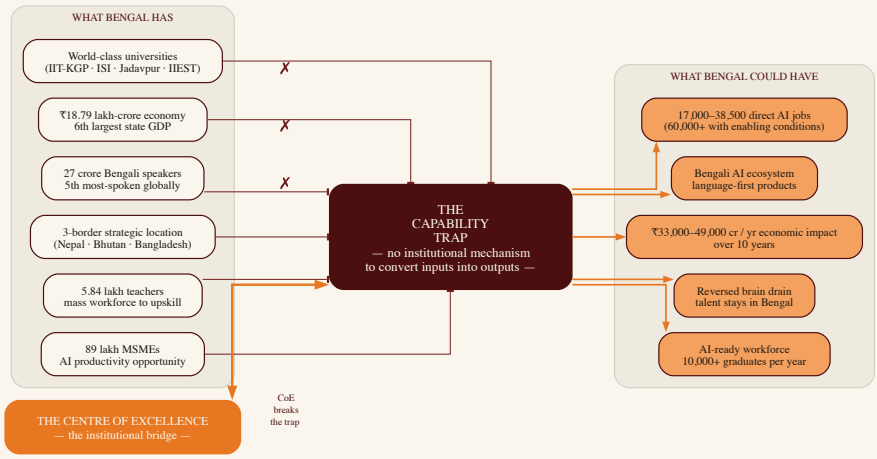


## 1.2 What the Centre Does

The Centre is structured around three reinforcing mandates and one cross-cutting commercialisation spine, deliberately not organised as a single training vertical.

- **AI for Every Citizen.** Foundational AI courses for citizens, school teachers, college professors, and small-business owners — delivered in Bengali, in the districts. Industry-ready tracks for students and graduates covering generative AI, retrieval-augmented generation, AI agents, and Model Context Protocol systems. This mandate is what makes the region AI-capable at population scale rather than merely AI-employed at the margins.
- **Industry Modernisation Partner.** Direct engagement with regional industry — tea, jute, agriculture, MSMEs, healthcare, financial services, logistics — through proofs of concept, co-built production AI deployments, and operational modernisation. Localised AI built by local talent, embedded in the busin-