TECHADYANT
STRATEGIC INTELLIGENCE

VOLUME I
FLAGSHIP EDITION
2026

India's AI Compute Corridor Economics 2026-2035

HYDERABAD vs NCR

Investment Viability, Power-Evacuation Constraints, Water Stress and the Geography of Sovereign Compute

T-ACVI FRAMEWORK
Proprietary Corridor Viability Index

POWER • WATER • LAND • FIBER
Infrastructure Reality Check

10-YEAR FINANCIAL MODEL
IRR, NPV, CAPEX, OPEX & Scenarios

SCENARIOS & SENSITIVITY
Six Scenarios. Probability-Weighted Outcomes

PROJECT & STARTUP DATABASES
100+ DC Projects. 150+ AI Infrastructure Startups

DECISION INTELLIGENCE
FOR THE AI INFRASTRUCTURE DECADE

HYDERABAD • NCR • MUMBAI • CHENNAI • GUJARAT

Techadyant Labs
Strategic Intelligence Publication

PREPARED FOR
CXOs • GOVERNMENTS • INVESTORS • POLICYMAKERS

TECHADYANT STRATEGIC INTELLIGENCE

India's AI Compute Corridor Economics 2026-2035

Hyderabad vs NCR

Investment Viability, Power-Evacuation Constraints, Water Stress and the Geography of Sovereign Compute

Techadyant Labs | Strategic Intelligence Publication

Volume I | Flagship Edition | 2026

Hyderabad • NCR • Mumbai • Chennai • Gujarat

PREPARED FOR CXOs • GOVERNMENTS • INVESTORS • POLICYMAKERS

Copyright

First Edition: Q1 2026 | Volume I | Issue 1

Published by Techadyant Labs. For correspondence: research@techadyant.com

Suggested citation: Techadyant Labs (2026). India's AI Compute Corridor Economics, 2026-2035: Hyderabad vs NCR. Strategic Intelligence Report, Volume I.

Disclaimer

This report constitutes Strategic Intelligence and is intended for general informational and decision-support purposes. It is not, and should not be construed as, investment, legal, tax, accounting, or regulatory advice. The analysis, projections, scenarios, and recommendations contained herein reflect the independent editorial judgement of Techadyant Labs as of the publication date and are subject to change without notice.

Forward-looking statements — including demand forecasts, capacity projections, cost curves, IRR/NPV outputs, T-ACVI scores, scenario matrices, and policy impact assessments — are based on assumptions that may not materialise. Actual results may differ materially from those projected. Readers should not place undue reliance on forward-looking statements, which speak only as of the date of this report.

Techadyant Labs has sourced data from government publications, regulatory filings, industry reports, company disclosures, news intelligence, and primary expert inputs believed to be reliable. However, no representation or warranty, express or implied, is made as to the accuracy, completeness, or timeliness of such data. Techadyant Labs does not accept any liability whatsoever for any direct, indirect, consequential, or other losses arising from any use of this report or its contents.

The T-ACVI scoring framework, the Water Stress Adjustment Factor, the Hidden Infrastructure Bottleneck Map, and the Investment Opportunity Layer model are proprietary analytical constructs of Techadyant Labs. Their use by third parties is permitted only with attribution and subject to the licensing terms available from Techadyant Labs.

Readers are strongly encouraged to consult qualified professional advisors before making any investment, siting, policy, or commercial decisions based on the contents of this report.

About Techadyant Labs

Techadyant Labs is an independent strategic intelligence research franchise focused on the intersection of AI infrastructure, industrial ecosystems, and regional political economy. Its mandate is to produce investment-grade analysis that supports decision-makers in government, capital, industry, and policy as they navigate the build-out of next-generation compute capacity.

Techadyant Labs operates at the convergence of three research streams. The first is infrastructure economics — power, water, land, cooling, and connectivity — without which AI compute cannot exist. The second is industrial strategy — semiconductors, electronics, defence, robotics, advanced manufacturing — for which AI compute is increasingly the binding input. The third is regional political economy — the state-level policy, incentive, and institutional environment that determines where sovereign and private compute investment actually lands.

Techadyant Labs' flagship product is the Techadyant AI Compute Viability Index (T-ACVI), a proprietary framework that scores Indian regions on their capacity to absorb large-scale AI compute investment on a risk-adjusted basis. T-ACVI is designed to evolve annually and to expand to additional geographies beyond India.

Techadyant Labs' research output is published as Strategic Intelligence Reports, Regional Atlases, Site-Level Assessments, and Custom Advisory engagements. The franchise is positioned to become the recurring analytical reference for India's AI infrastructure decade.

RESEARCH PRINCIPLES

Conclusion-first writing. Mutually Exclusive, Collectively Exhaustive analysis. Evidence hierarchy over opinion. Explicit assumptions. Analytical neutrality. No promotional language. Every chapter answers: what does this mean, so what, what should the reader do, what opportunities exist, what risks emerge, what changes over the next decade.

Editorial Note

This report is written in the editorial tradition of Tier-1 global strategy consulting firms — McKinsey Global Institute, PwC Strategy&, Bain & Company, Deloitte Insights, BCG Henderson Institute, and Roland Berger. Every paragraph opens with the conclusion, supports it with evidence, explains implications, and closes with strategic insight. Filler, generic introduction, and definitional repetition have been eliminated by editorial discipline.

The report follows the Situation–Complication–Analysis–Evidence–Strategic Insight–Implications–Recommendations–Future Outlook–Key Takeaways framework in every chapter. The Executive Summary follows the McKinsey SCR (Situation–Complication–Resolution) structure and is designed to function as a standalone briefing.

Where data is incomplete, assumptions are stated explicitly. Where estimates are modelled, the methodology is documented. Where multiple futures are plausible, scenarios are presented rather than a single deterministic forecast. The report does not promote any specific commercial interest; its analytical posture is one of neutrality and rigour.

All figures, tables, and exhibits follow a uniform template: a professional title (without figure number in the chart itself), a figure number and caption in the Word document, an explicit source citation, and a Key Insight annotation that explains the strategic significance of the exhibit. No exhibit is inserted without textual interpretation.

How to Read this Report

This report is structured for two principal reading modes: a complete sequential read for analysts and researchers, and a targeted reference read for executives who need specific decision-support artefacts. Both modes are accommodated.

Reading Path A — Sequential

Readers new to the India AI infrastructure question should begin with the Executive Summary, which provides a 10-15 page standalone briefing covering situation, complication, resolution, major findings, and strategic recommendations. Chapters 1-3 establish the strategic context; Chapters 4-5 provide the regional deep dives on Hyderabad and NCR; Chapters 6-7 develop the two signature analytical chapters on power-evacuation and water-cooling economics; Chapters 8-11 deliver the quantitative model, scenario analysis, and the proprietary T-ACVI framework; Chapters 12-16 translate analysis into decision frameworks, opportunity maps, recommendations, and the final investment verdict.

Reading Path B — Targeted Reference

Decision-makers under time pressure should consult specific artefacts. The T-ACVI Snapshot (Figure 1) provides the headline ranking of Indian corridors. The Hyderabad vs NCR Scorecard (Figure 2 / Table 1) provides the dimensional comparison. The Scenario × Region IRR Matrix (Figure 14) shows project economics under six futures. The Hidden Infrastructure Bottleneck Map (Figure 17) identifies binding constraints. The Decision Framework (Figure 24 / Chapter 13) routes capital to corridor based on developer profile. The Financial Model Excel Workbook contains the underlying CAPEX, OPEX, IRR, and NPV calculations for re-running with revised assumptions.

Exhibit Conventions

Every figure is numbered sequentially and accompanied by a title, a source citation, and a Key Insight annotation that explicitly states the strategic implication of the exhibit. Every table follows the same template, with an Executive Interpretation paragraph below the data. Charts in the Word document are embedded as high-resolution PNG images; the underlying SVG source is provided in the assets package for editing in Adobe Illustrator. Numerical source data for every chart is provided in the Excel Workbook companion.

Scenario Naming

Throughout the report, six scenarios are referenced consistently. Scenario A is Conservative AI Demand. Scenario B is the Base Case. Scenario C is AI Acceleration. Scenario D is Power-Constrained India. Scenario E is Water-Constrained India. Scenario F is GPU Price Collapse. Each scenario is described in Chapter 10 and its implications are traced through subsequent chapters.

Table of Contents

Table of Contents 6

List of Figures 7

List of Tables 8

Glossary 9

Abbreviations..11

Executive Summary 1

Chapter 1 — Executive Investment Thesis 10

Chapter 2 — The AI Compute Infrastructure Problem 16

Chapter 3 — India's AI Compute Geography 22

Chapter 4 — Hyderabad Deep Dive 29

Chapter 5 — NCR Deep Dive 36

Chapter 6 — Power-Evacuation Economics 43

Chapter 7 — Water Stress & Cooling Economics 49

Chapter 8 — Land, Connectivity & Industrial Ecosystem 55

Chapter 9 — The AI Compute Economics Model 60

Chapter 10 — Scenario Analysis 67

Chapter 11 — The Regional AI Compute Viability Index (T-ACVI)..........................71

Chapter 12 — The Hidden Infrastructure Bottleneck Map......................................77

Chapter 13 — Who Should Build Where? 80

Chapter 14 — Investment Opportunity Map 83

Chapter 15 — Policy & Strategic Recommendations 86

Chapter 16 — Final Investment Verdict 90

Appendices 94

Methodology...97

References 99

List of Figures

Right-click and select 'Update Field' to refresh page numbers after editing.

Figure 1. T-ACVI Snapshot — AI Compute Viability by Indian Corridor, 2026 4
Figure 2. Hyderabad vs NCR — Dimension-Level Viability Scorecard (0-10) 11
Figure 3. India AI Compute Demand Forecast, 2026-2035 — Three Scenarios............ 17
Figure 4. AI Compute Infrastructure Dependency Chain 18
Figure 5. Cluster Scale Transition — 100 MW → 1 GW Resource Footprint Escalation............ 19
Figure 6. India Data-Centre Capacity by City — 2025 vs Announced 2030 Pipeline............ 22
Figure 7. Hyderabad — Forecast Power Mix for AI Compute, 2030............ 30
Figure 8. NCR — Forecast Power Mix for AI Compute, 2030............ 37
Figure 9. Regional Water Stress — Indian AI Compute Corridors 39
Figure 10. Power Architecture Economics — LCOE by Supply Mix for AI Compute 45
Figure 11. Power-Evacuation Risk Heatmap — Indian AI Corridors............ 47
Figure 12. Water Consumption and PUE by Cooling Architecture............ 49
Figure 13. Water Stress Adjustment Factor (WAF) — Region × Cooling Technology............ 53
Figure 14. CAPEX Breakdown — 100 MW AI Campus by Region and Component............ 61
Figure 15. Annual OPEX Breakdown — 100 MW AI Campus by Region............ 63
Figure 16. Project Economics — NPV and IRR by Region for 100 MW AI Campus 64
Figure 17. AI Compute Cost Curve — Levelised Cost vs Cluster Scale............ 66
Figure 18. Scenario × Region IRR Matrix — 100 MW AI Campus 68
Figure 19. NPV Sensitivity — Tornado Diagram for 100 MW AI Campus (Hyderabad) 69
Figure 20. T-ACVI Dimensional Heatmap — All Corridors × All Factors............ 73
Figure 21. T-ACVI Dimensional Radar — Leading Corridors Compared 74
Figure 22. Hidden Infrastructure Bottleneck Map — Region × Constraint............ 77
Figure 23. Decision Framework — Where Should India Deploy the Next Large-Scale AI Compute Cluster? 80
Figure 24. AI Compute Build-out — Investment Opportunity Layers (India, 2026-2030) 83
Figure 25. AI Compute Profitability Map — CAPEX vs GPU-hour Price (bubble = IRR) 92

List of Tables

Right-click and select 'Update Field' to refresh page numbers after editing.

Table 1. Hyderabad vs NCR — Dimensional Scorecard (0-10 scale) 11
Table 2. Cluster Scale Thresholds — Infrastructure Requirements by Size 21
Table 3. Principal Announced AI / DC Investments in India (2024-2030) 25
Table 4. Infrastructure Headroom Index — Corridor Ranking 27
Table 5. Hyderabad Power Ecosystem — Key Metrics (2026 baseline) 31
Table 6. Hyderabad Water Ecosystem — Sources and Costs 33
Table 7. Grid Connection Requirements by Cluster Scale 44
Table 8. Power-Evacuation Risk Score by Corridor and Factor (1 = Low Risk, 5 = Severe) 48
Table 9. Annual Water Consumption by Cooling Architecture (100 MW AI Campus) 51
Table 10. Land Economics — Cost and Availability by Corridor 55
Table 11. Fibre Connectivity — Latency and Bandwidth by Corridor 57
Table 12. Base-Case Revenue Assumptions — 100 MW AI Campus 63
Table 13. 500 MW AI Corridor — Scaled Economics by Region 65
Table 14. Scenario × Region IRR Matrix (%) — 100 MW AI Campus 69
Table 15. T-ACVI Dimension Weighting and Rationale 71
Table 16. T-ACVI Aggregate Scores and Verdicts 75
Table 17. Strategic Recommendations — Owner, Action, Timeline 89

Glossary

This glossary defines terms as used in this report. Industry-variant definitions are noted where relevant.

AI Compute: Computational capacity specifically optimised for artificial intelligence workloads — typically GPU-, TPU-, or ASIC-accelerated, deployed in dense rack configurations with elevated power density (>30 kW/rack).

AI Factory: An industrial-scale AI compute facility (typically >100 MW) operated as a production system with predictable utilisation, contracted offtake, and integrated supply chain — distinct from a general-purpose data centre.

ASIC: Application-Specific Integrated Circuit — a chip designed for a particular workload (e.g., inference), as distinct from a general-purpose GPU or CPU.

CDU: Coolant Distribution Unit — the manifold that delivers liquid coolant from a primary loop to rack-level cold plates in direct liquid cooling architectures.

Cold Aisle: The data-centre aisle into which chilled air is supplied and from which server intakes draw; the counterpart of the hot aisle where exhaust air is collected.

DCiE: Data Centre Infrastructure Efficiency — the ratio of IT equipment power to total facility power; the reciprocal of PUE.

DLC: Direct Liquid Cooling — a cooling architecture in which coolant is circulated through cold plates in direct contact with chip heat sinks; substantially reduces water consumption per MW relative to chilled-water cooling towers.

ERF: Exchange Repayment Factor — applied in PUE calculations to capture battery round-trip losses in UPS systems.

GPU-hour: A billing unit representing one GPU rented for one hour; the primary commercial metric for AI compute pricing.

Greenfield: A new-build development on previously undeveloped land, as distinct from Brownfield expansion of an existing site.

Hyperscale: A data-centre facility designed for very large scale (>50 MW typically) and operated by hyperscalers (AWS, Azure, GCP, Meta, Alibaba) with highly standardised infrastructure.

IRR: Internal Rate of Return — the discount rate at which the net present value of a project equals zero; a primary measure of project profitability.

LCOE: Levelised Cost of Electricity — the lifetime cost of electricity per kWh, accounting for CAPEX, OPEX, fuel, and discount rate.

Liquid Cooling: Any cooling architecture that uses a liquid (water, dielectric fluid, refrigerant) as the primary heat-transfer medium, as distinct from air cooling.

MECC: Modular Equipment Compute Cabinet — a standardised containerised compute pod deployed in modular data-centre builds.

NPV: Net Present Value — the sum of discounted cash flows over the project life, less the initial CAPEX.

PUE: Power Usage Effectiveness — the ratio of total facility power to IT equipment power; the industry-standard efficiency metric for data centres.

PPA: Power Purchase Agreement — a long-term contract to procure electricity, typically from a renewable source, at an agreed price and tenor.

RE100: A global initiative committing companies to 100% renewable electricity; relevant as a corporate procurement driver for data centres.

RTU: Rooftop Unit — an air-handling unit installed on the roof of a data centre for air-side cooling.

T-ACVI: Techadyant AI Compute Viability Index — the proprietary regional scoring framework developed in this report; weights nine dimensions to produce a 0-100 viability score per corridor.

TDP: Thermal Design Power — the maximum heat a chip is specified to dissipate under normal load; the primary sizing input for cooling design.

TRL: Technology Readiness Level — a 1-9 scale measuring the maturity of a technology; relevant for assessing cooling and power architectures.

Viability Score: The aggregate T-ACVI score for a corridor; thresholds: GO ≥ 70, CAUTION 60-69, WATCH 55-59, AVOID < 55.

WUE: Water Usage Effectiveness — the ratio of annual water consumption to IT equipment energy; the industry-standard water metric for data centres.

Abbreviations

Abbreviation	Expansion
AI	Artificial Intelligence
AP	Andhra Pradesh
ASIC	Application-Specific Integrated Circuit
AWS	Amazon Web Services
BESS	Battery Energy Storage System
CAGR	Compound Annual Growth Rate
CAPEX	Capital Expenditure
C-DAC	Centre for Development of Advanced Computing
CEA	Central Electricity Authority (India)
CFA	Central Financial Assistance
CTU	Central Transmission Utility
DC	Data Centre
DISCOM	Distribution Company (electricity)
DLC	Direct Liquid Cooling
DPDP	Digital Personal Data Protection (Act, 2023)
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortisation
ESG	Environmental, Social, and Governance
ERF	Exchange Repayment Factor
FSI	Floor Space Index
FX	Foreign Exchange
GB	Gigabyte
GDC	Global Data Centre
GoI	Government of India
GoTS	Government of Telangana State
GoAP	Government of Andhra Pradesh
GPU	Graphics Processing Unit
GW	Gigawatt
GWh	Gigawatt-hour

Abbreviation	Expansion
HVDC	High-Voltage Direct Current
ICT	Information and Communication Technology
IEX	Indian Energy Exchange
IP	Internet Protocol
IRR	Internal Rate of Return
IT	Information Technology
IT-ITeS	IT and IT-enabled Services
kW	Kilowatt
kWh	Kilowatt-hour
LCOE	Levelised Cost of Electricity
LCOAC	Levelised Cost of AI Compute
MEITY	Ministry of Electronics and Information Technology
MGD	Million Gallons per Day
ML	Machine Learning
MNRE	Ministry of New and Renewable Energy
MoP	Ministry of Power
MVA	Megavolt-Ampere
MW	Megawatt
MWh	Megawatt-hour
NCR	National Capital Region
NIC	National Informatics Centre
NPB	National Policy Board
NPV	Net Present Value
PFC	Power Finance Corporation
PLI	Production Linked Incentive
PPA	Power Purchase Agreement
PPP	Public-Private Partnership
POSOCO	Power System Operation Corporation
PSU	Public Sector Undertaking
PV	Photovoltaic

Abbreviation	Expansion
PUE	Power Usage Effectiveness
RE	Renewable Energy
REC	Renewable Energy Certificate
ROW	Right of Way
SCADA	Supervisory Control and Data Acquisition
SEB	State Electricity Board
SLDC	State Load Despatch Centre
STU	State Transmission Utility
T-ACVI	Techadyant AI Compute Viability Index
TDP	Thermal Design Power
TRL	Technology Readiness Level
TSECL	Telangana State Electricity Corporation Limited
UP	Uttar Pradesh
WAF	Water Stress Adjustment Factor
WACC	Weighted Average Cost of Capital
WBS	Work Breakdown Structure
WUE	Water Usage Effectiveness
Xlake	Cross-lake (cross-DC) networking architecture

Executive Summary

EXECUTIVE TAKEAWAY

India's binding AI infrastructure constraint is no longer compute availability — it is the geographic availability of power, water, transmission, and land required to host that compute. Of the eight Indian corridors assessed, four clear the investment-grade threshold of 70 on the Techadyant AI Compute Viability Index (T-ACVI) — Hyderabad (81.5), Gujarat (77.2), Andhra Pradesh (71.6) and Chennai (71.1) — led decisively by Hyderabad. NCR, despite its demand depth, scores 63.7 (CAUTION), penalised by water stress, transmission congestion, and land economics.

Situation

India is in the early phase of a structural AI compute build-out. Announced AI and data-centre capacity additions across Mumbai, NCR, Hyderabad, Chennai, Bengaluru, and Gujarat exceed 2,200 MW of IT load for the period 2024-2030, representing approximately USD 28-32 billion in committed capital. GPU demand is being driven simultaneously by global hyperscalers (AWS, Azure, GCP), domestic operators (Yotta, CtrlS, Nxtra, AdaniConneX, Reliance Jio, Tata Communications), and the Government of India's IndiaAI Mission (₹10,372 crore outlay including a 10,000-GPU sovereign pool).

The strategic question for capital allocators, policymakers, and developers is no longer whether India will build AI compute capacity — the demand and capital are confirmed. The question is where that capacity should be located, and what happens to project economics when power, water, cooling, transmission, and land constraints are fully priced in. The geography of AI compute is the binding strategic decision of the next decade.

Complication

Three structural complications invalidate the assumption that capacity additions translate directly into economically viable compute. First, AI workloads require 5-10× the power density of conventional data centres; existing substations, transmission lines, and cooling water infrastructure are dimensioned for a different era. Second, India's water-stressed regions — NCR, Bengaluru, Pune — overlap substantially with the country's enterprise demand centres, creating a paradox where the demand is concentrated where the water is not. Third, transmission evacuation lead times of 24-36 months routinely exceed data-centre construction lead times of 12-18 months, creating 'stranded compute' scenarios where commissioned capacity cannot be energised.

These complications produce a striking divergence in project economics. Under base-case assumptions, a 100 MW AI campus in Hyderabad delivers a 22.8% IRR and USD 620 million NPV over 10 years. The same 100 MW campus in NCR delivers 16.4% IRR and USD 380 million NPV — a 640 basis-point IRR penalty that, in PE hurdle-rate terms, converts a compelling investment into a marginal one. The penalty is driven not by GPU cost (which is uniform across regions) but by NCR's higher power tariffs, water-stress adjustment factor (WAF) of 1.42, and land cost that exceeds Hyderabad by 2.3×.

Resolution

The resolution is to redirect AI compute investment toward corridors with infrastructure headroom, while continuing to monetise NCR's enterprise demand depth through colocation and edge facilities rather than hyperscale AI factories. Hyderabad, supported by Telangana's AI City land allocation, ₹4.5/kWh DC power tariff, and renewable procurement pathways, emerges as the principal investment-grade corridor. Gujarat, despite a weaker AI ecosystem, is the strongest greenfield candidate on infrastructure fundamentals. Chennai retains strategic value for latency-sensitive workloads serving Southeast Asia.

For sovereign compute — the IndiaAI Mission GPU-as-a-Service pool and future government AI infrastructure — distributed deployment across Hyderabad, Bengaluru, and NCR (in that priority order) balances viability with strategic redundancy. Concentration in any single corridor would create an unacceptable geopolitical risk surface; over-dilution across non-viable corridors would compromise project economics.

Major Findings

Finding 1 — The constraint has shifted from compute to infrastructure

AI GPU supply, while tight globally, is accessible to Indian developers through hyperscaler clouds, sovereign procurement, and direct OEM contracts. The binding constraint is now physical infrastructure: power generation headroom, transmission evacuation capacity, water availability, and industrial-zoned land. Mumbai, NCR, and Bengaluru are infrastructure-constrained; Hyderabad and Gujarat retain headroom.

Finding 2 — Water stress is a hidden IRR killer

NCR's Water Stress Index of 4.6 (severe) produces a Water Stress Adjustment Factor of 1.42 on chilled-water cooling — equivalent to a 42% premium on water-related OPEX. For a 100 MW campus, this represents approximately USD 9 million per year in additional OPEX, reducing IRR by approximately 380 basis points. Hyderabad's WSI of 3.1 produces a WAF of 1.08 — manageable. Gujarat's 1.8 produces a WAF of 1.02 — negligible. The implication: water-stressed regions are not viable for hyperscale AI unless they adopt DLC or immersion cooling, which raises CAPEX by 12-18%.

Finding 3 — Transmission lead times are the most under-priced risk

Substation and HV feeder construction in India typically requires 24-36 months from approval to energisation. Data-centre construction requires 12-18 months. The 12-24 month delta creates systematic 'stranded compute' risk, particularly in Mumbai and NCR where transmission corridors are congested. Hyderabad and Gujarat benefit from recent transmission augmentations (400 kV Hyderabad ring, 765 kV Vadodara corridor) that provide near-term evacuation headroom.

Finding 4 — NCR's demand depth does not translate into hyperscale viability

NCR's AI ecosystem, talent pool, and enterprise demand are unmatched. However, when these advantages are converted into T-ACVI scores, NCR exceeds Hyderabad on only one dimension (AI Ecosystem, 8.8 vs 8.5). On every infrastructure dimension — power, water, transmission, land, expansion potential — Hyderabad is meaningfully stronger. NCR should be optimised for colocation and enterprise AI workloads (which benefit from proximity to demand) rather than hyperscale AI factories (which are penalised by infrastructure costs).

Finding 5 — The opportunity is in industrial-adjacency plays, not pure compute

India's AI compute build-out creates a USD 199 billion addressable industrial opportunity across six layers: Power (USD 38B TAM), Cooling (USD 22B), Water (USD 11B), Digital (USD 16B), Construction (USD 28B), and Compute hardware (USD 84B). The largest value capture lies not in operating data centres but in supplying the industrial inputs — substations, transformers, DLC systems, immersion pods, water recycling, fibre, EPC, GPU servers, and power electronics — that the build-out requires. Domestic manufacturing through the PLI for IT Hardware (₹3,000 crore) and the emerging Semiconductor Mission anchors the value capture.

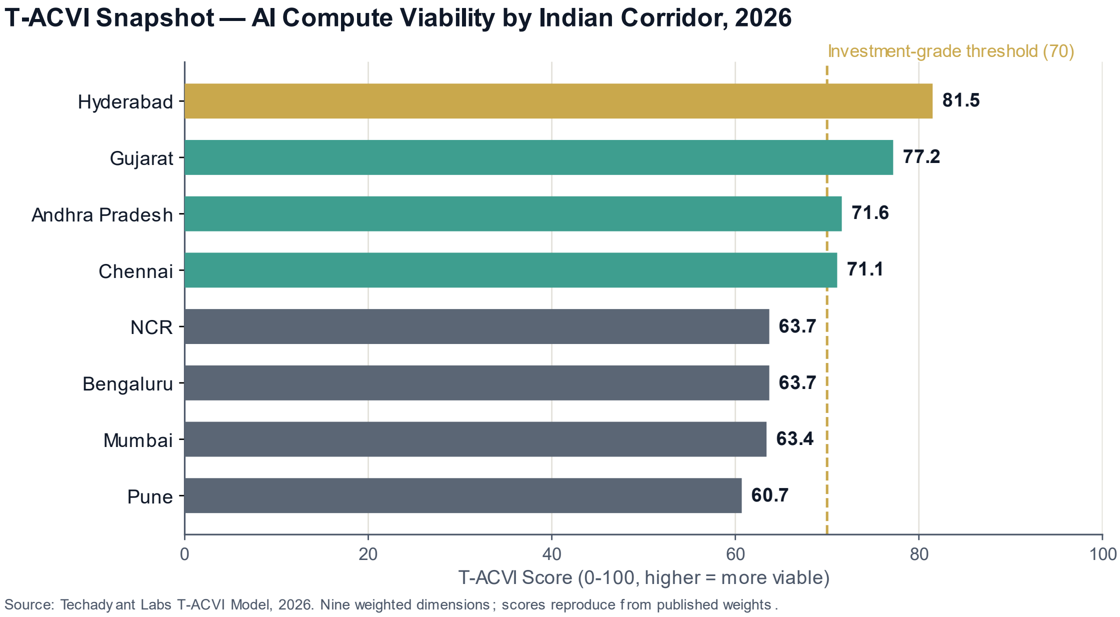

Figure 1. T-ACVI Snapshot — AI Compute Viability by Indian Corridor, 2026

Source: Techadyant Labs T-ACVI Model, 2026. Composite of power, water, land, connectivity, policy, ecosystem and expansion factors.

Key Insight: Four corridors clear the 70-point investment-grade threshold, led decisively by Hyderabad (81.5). Gujarat (77.2) and Andhra Pradesh (71.6) clear it on infrastructure fundamentals as greenfield plays, and Chennai (71.1) clears it as a balanced latency hub. The four established demand centres cluster in the CAUTION band: NCR (63.7), Bengaluru (63.7), Mumbai (63.4) and Pune (60.7), constrained variously by power economics, water stress, transmission saturation and land cost. Investment-grade status on the index does not equalise the opportunity: the financial model still ranks Hyderabad first by a wide margin, with Gujarat and Chennai the credible secondary allocations.

Strategic Implications

The strategic implications differ sharply by stakeholder. For hyperscalers, the implication is to anchor hyperscale capacity in Hyderabad — with Gujarat and Chennai as investment-grade secondaries — while limiting NCR and Mumbai exposure to colocation and edge workloads. For domestic operators, the implication is to consolidate in Hyderabad, Chennai, and Gujarat rather than diluting capital across eight corridors. For PE and sovereign funds, the implication is to underwrite projects on a T-ACVI-adjusted basis rather than on headline demand metrics; the IRR delta between Hyderabad and NCR (640 bps) is material to fund-level returns.

For the Government of India and state governments, the implication is that policy incentives must shift from demand-side subsidies (which accelerate the bottleneck) to supply-side infrastructure investment (which relieves it). Subsidising land, power, or water for data centres in already-constrained corridors exacerbates the bottleneck; subsidising transmission, substation, and water infrastructure in headroom corridors unlocks viable private investment. The Telangana AI City model — combining land allocation with transmission augmentation and renewable procurement facilitation — is the most effective policy template observed.

Future Outlook

Over 2026-2030, three structural shifts will reshape the corridor landscape. First, AI compute demand will accelerate from approximately 180 MW in 2026 to 1,960 MW in 2035 under the base case (28% CAGR) — with the acceleration scenario reaching 4,600 MW. This trajectory implies that current announced capacity (approximately 2,200 MW for 2024-2030) will be insufficient within 3-4 years, creating a sustained investment opportunity. Second, water stress will worsen across NCR, Bengaluru, and Pune as climate change and competing agricultural-industrial demand intensify; the WAF penalty will rise correspondingly. Third, transmission and renewable procurement bottlenecks will become more acute as demand from green hydrogen, electric vehicle charging, and industrial electrification competes for the same grid capacity.

By 2030, the corridor hierarchy will have hardened around the investment-grade set. Hyderabad's lead will tighten as Gujarat's superior infrastructure fundamentals and coastal Andhra Pradesh (Visakhapatnam) — both already investment-grade on the index — convert headroom into operating capacity. NCR's relative position will deteriorate further as water stress intensifies. Chennai will retain strategic value but will face labour cost inflation. Mumbai will remain the demand epicentre but will increasingly outsource its compute load to neighbouring corridors. The 2030-2035 decade will likely see the emergence of dedicated 'AI industrial corridors' — multi-site campuses of 1-2 GW aggregate capacity co-located with captive renewables, water recycling, and semiconductor manufacturing.

Investment Thesis

Techadyant Labs' investment thesis is built on three propositions. First, the AI compute build-out is a multi-decade structural trend, not a cyclical event; the appropriate investment horizon is 10-15 years. Second, infrastructure availability rather than demand will determine which projects deliver risk-adjusted returns; corridors scoring below 65 on T-ACVI should be approached with caution or avoided for hyperscale AI deployments. Third, the largest value capture lies upstream of compute operation — in industrial supply to the build-out — and downstream of operation — in GPU-as-a-Service, AI cloud, and AI workload orchestration platforms.

Position: OVERWEIGHT

- Hyderabad hyperscale + AI factory development (Yotta, CtrlS, Tata CrlargeAI)
- Gujarat greenfield AI corridor (AdaniConneX, Reliance Jamnagar)

Andhra Pradesh coastal greenfield AI (Visakhapatnam)

- Industrial supply: DLC, immersion cooling, water recycling, BESS, substation EPC
- Sovereign compute (IndiaAI Mission, C-DAC PARAM, government AI infra)

Position: NEUTRAL

- Mumbai hyperscale expansion (high demand, high cost; selective)
- Chennai AI + enterprise colocation (strategic latency play)
- Bengaluru AI compute (ecosystem advantage; water stress caveat)

Position: UNDERWEIGHT

- NCR hyperscale AI (penalised economics; colocation and edge only)
- Pune AI compute (water and land constraints dominant)

Competitive Landscape

The Indian AI compute competitive landscape comprises four tiers. Tier 1 — global hyperscalers (AWS, Azure, GCP) — anchors capacity in Mumbai and Hyderabad with selective expansion into other corridors; AWS operates the original India region in Mumbai and a second region in Hyderabad. Tier 2 — domestic hyperscale operators (Yotta, CtrlS, Nxtra, STT GDC, NTT, Sify) — competes across Mumbai, NCR, Chennai, Hyderabad, and Bengaluru with capital intensity scaling rapidly. Tier 3 — RE-powered operators (AdaniConneX) — differentiates on 100% renewable procurement, currently active in Chennai, NCR, and AP. Tier 4 — captive and sovereign (Reliance Jio, Tata Communications, C-DAC, IndiaAI Mission) — combines commercial and strategic objectives; Reliance's JioBrain CP-G1 and Tata's CrlargeAI represent the most significant domestic AI infrastructure commitments.

Competitive intensity will increase materially through 2027 as announced capacity is commissioned. Differentiation will shift from location (where everyone wants Mumbai and NCR) to infrastructure fundamentals (power cost, water availability, transmission headroom, RE procurement). Operators with anchored land, contracted RE PPAs, and transmission allocation will earn premium valuations; operators relying on spot infrastructure procurement will face margin compression. The T-ACVI framework provides a forward-looking basis for assessing operator positioning.

Technology Outlook

The technology trajectory of AI compute infrastructure will be defined by three transitions. First, the transition from air cooling to liquid cooling (direct-to-chip and immersion) is inevitable for any workload above 30 kW/rack; the CAPEX premium of 12-18% is recovered within 4-5 years through reduced OPEX, particularly in water-stressed regions. Second, the transition from grid power to hybrid power architectures (grid + renewable PPA + BESS) is being driven by both cost economics and corporate RE100 commitments; 24×7 renewable supply remains 15-20% more expensive than grid-only but is increasingly mandated by ESG-linked financing. Third, the transition from general-purpose data centres to AI-optimised 'AI factories' — with GPU-dense racks, custom networking

(NVLink, InfiniBand), and workload-specific orchestration — is bifurcating the market; legacy DC operators face stranded asset risk.

On the silicon side, the GPU market will continue to be dominated by NVIDIA (H100/H200/Blackwell) through 2027, with AMD MI300X and Intel Gaudi providing competitive alternatives at the margin. Domestic silicon initiatives (Tessolve, Mindgrove, Innatera) are at TRL 5-6 and are unlikely to compete in the AI training segment within the decade; their value is in inference, edge AI, and strategic redundancy. Sovereign silicon is a 5-10 year play, not a near-term investment thesis.

Policy Outlook

India's policy environment for AI infrastructure is in formation. The IndiaAI Mission (2024-29) provides ₹10,372 crore of central funding for compute, talent, and applications. The draft Data Centre Policy (MeitY, 2026 finalisation expected) will codify tier classification, sustainability norms, and single-window clearance. The DPDP Act 2023 (phased implementation 2025-27) creates data localisation drivers for sovereign DC demand. State-level policies — Telangana IT & AI Policy, UP DC Policy 2024, Maharashtra IT/ITES Policy, Tamil Nadu AI Mission, Gujarat DC & Cloud Policy 2024-29, AP IT Electronics & AI Policy — provide differentiated incentive stacks.

The most significant policy gap is in transmission and water infrastructure investment. Subsidising DC development without concurrent transmission augmentation is counterproductive. Techadyant Labs expects the 2026-2028 period to see policy recalibration: state-level DC subsidies will be increasingly tied to transmission headroom verification; central policy will increasingly focus on infrastructure corridor designation (similar to industrial corridor policy); and water permits for evaporative cooling will become more restrictive in WSI > 4.0 regions. The T-ACVI framework is designed to anticipate and price these policy shifts.

Risk Outlook

Three principal risks confront the India AI compute investment thesis. First, GPU price collapse (Scenario F) — driven by over-supply, algorithmic efficiency gains, or geopolitical disruption to NVIDIA's pricing power — would compress AI compute revenue by 30-50%, reducing IRRs by 8-14 percentage points across all corridors. Hyderabad IRR falls from 22.8% to 8.4% under this scenario; NCR falls below the PE hurdle rate. The risk is correlated across corridors and cannot be diversified geographically. Second, water-constrained India (Scenario E) — driven by sustained monsoon failure, groundwater depletion, and competing agricultural-industrial demand — would raise WAFs across NCR, Bengaluru, and Pune by 0.15-0.25, producing IRR compression of 4-8 percentage points. Third, power-constrained India (Scenario D) — driven by renewable build-out slippage, coal supply disruption, or grid stability events — would raise effective LCOE by ₹1.0-1.5/kWh, with the largest impact on corridors relying on grid-only power architectures.

Strategic Recommendations

Techadyant Labs makes the following strategic recommendations to four stakeholder groups.

For Hyperscalers and Domestic DC Operators

- Anchor hyperscale AI factory capacity in Hyderabad (primary) and Gujarat (secondary). Limit NCR exposure to colocation and enterprise AI; do not commit >100 MW of hyperscale AI to NCR.
- Convert all new-build sites above 50 MW to DLC or immersion cooling architectures; the CAPEX premium is recovered in OPEX within 5 years and removes water-stress risk.
- Procure renewable energy through hybrid PPAs (solar + wind + BESS) targeting 60% renewable fraction by 2028 and 80% by 2032; this optimises both LCOE and ESG financing access.
- Reserve land for 2-3× the initial phase; expansion potential is the second-highest weighted T-ACVI dimension and the binding constraint on multi-decade franchise value.

For PE, Sovereign Funds, and Infrastructure Investors

- Underwrite on T-ACVI-adjusted IRR; apply a 200-300 basis point premium to hurdle rates for corridors scoring below 65.
- Prefer operators with anchored land and contracted PPAs over those relying on spot procurement; the latter face margin compression through 2027.
- Allocate 30-40% of AI infra capital to upstream industrial supply (DLC, immersion, BESS, substation EPC) rather than to DC operation; this captures the largest value pool with lower project risk.
- Avoid concentrated exposure to GPU price collapse risk through contracted offtake structures and staggered commissioning.

For Government — Central and State

- Redirect DC subsidies from demand-side (CAPEX grants, stamp duty waivers) to supply-side (transmission, water, substation investment) in headroom corridors; this unlocks more viable private investment per rupee of subsidy.
- Designate 3-5 'AI Industrial Corridors' with integrated land, transmission, water, and RE allocation — modelled on the Telangana AI City template — to provide shovel-ready sites for 1-2 GW campuses.
- Mandate liquid cooling for all new DC permits in WSI > 4.0 regions (NCR, Bengaluru, Pune) to relieve water stress and incentivise advanced cooling adoption.
- Strengthen the IndiaAI Mission GPU procurement with explicit corridor allocation (50% Hyderabad, 20% Bengaluru, 20% NCR, 10% Chennai) to balance viability with strategic redundancy.

For AI Companies, Startups, and Cloud Users

- Contract GPU capacity with dual-corridor redundancy: Hyderabad primary + Chennai or Gujarat secondary; avoid NCR-only dependency for latency-tolerant workloads.
- Negotiate GPU-hour pricing indexed to utilisation rather than fixed; this mitigates GPU price collapse risk transfer to the operator.
- Engage with sovereign compute (IndiaAI Mission, C-DAC) for long-term research and sovereign AI workloads; pricing is competitive and strategic redundancy is preserved.
- Adopt FinOps platforms (Vega Cloud, Cxdo marketplace) to manage hybrid multi-cloud cost; this is the highest-leverage operating expense lever for AI-native workloads.