PROACTIVE™
Enhancing Digital Experience

AN eBOOK FOR INDIAN ENTERPRISES

The AI-Ready Data Center Blueprint for India

Your staged path from a classic data center to an AI factory, without a forklift rebuild.

Proactive Data Systems Pvt Ltd

Cisco Preferred Cloud and AI Partner · Dell Platinum Partner · NetApp Preferred Partner
35 years in enterprise IT · 1,500+ organisations served · 24/7 service desk in India

Table of Contents

Introduction: Why a Blueprint, Not a Rebuild 03

1. The AI-Readiness Maturity Model 04

2. The Six Dimensions of AI-Readiness 05
   Self-Assessment Scorecard

3. The Staged Roadmap 08

4. The Decisions You Will Face 09

5. Budgeting and Sequencing 10

6. Common Pitfalls to Avoid 11

7. Your First 90 Days 12

Putting the Blueprint to Work 12

Frequently Asked Questions 13

Glossary 14

Introduction: Why a Blueprint, Not a Rebuild

The board has approved the AI. The data center underneath it was built for a different job. That is the position most Indian enterprises are in, and the instinct it provokes, rip it out and build an AI factory, is the most expensive way to proceed and rarely the right one. An existing data center can almost always be made AI-ready in stages, closing the gaps that matter in an order that delivers value early and defers cost until it is justified.

This blueprint is the staged path. It gives you a way to measure where your data center stands today, a structured assessment of the six things that have to change, a phased roadmap to close the gaps without a forklift rebuild, the decisions you will face at each stage, and a 90-day plan to begin. It is written for the CIO or CTO who has to turn an AI ambition into infrastructure that works, on a budget that survives scrutiny.

1. The AI-Readiness Maturity Model

Before you plan a journey, locate your starting point. The AI-Readiness Maturity Model places a data center on one of five levels. Most enterprises find themselves at Level 1 or 2, running pilots on infrastructure that will not carry production. The goal is Level 3, a data center that runs production AI dependably; Level 4 is for organisations whose business now depends on AI at scale.

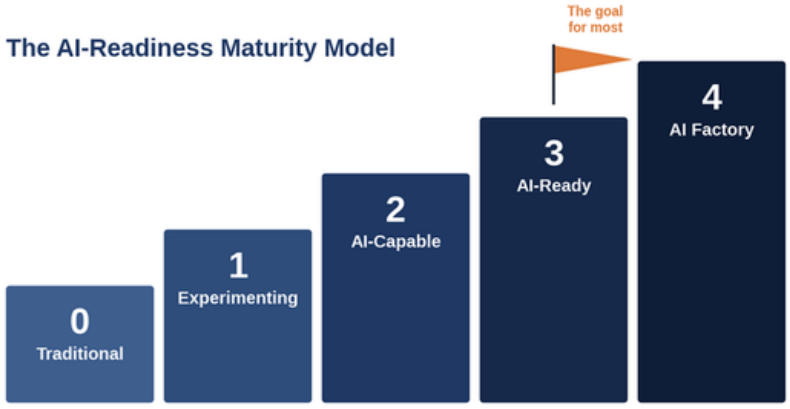


Figure 1: The five levels of AI-readiness. Most enterprises start at Level 1–2.

Level	Description	Typical Symptom
0. Traditional	CPU workloads, air-cooled, no GPU capability	AI is “a cloud thing we haven’t started”
1. Experimenting	Ad hoc GPUs or rented cloud capacity, no design	Pilots run on borrowed or rented hardware
2. AI-Capable	Some GPU servers, but storage, network and cooling unprepared	Pilots work; production stalls or strains
3. AI-Ready	Balanced GPU, storage, network and cooling; governance in place	Production AI runs dependably
4. AI Factory	Purpose-built, scales on demand, operated and optimised	AI scales with the business

Key takeaway: A pilot that ran once on a single GPU is Level 1, not Level 3, however impressive the demo. Locate yourself honestly first.

2. The Six Dimensions of AI-Readiness

AI-readiness is not one upgrade; it is six, and the weakest dimension caps the rest. The assessment below is the heart of this blueprint. For each dimension, judge where you stand and what "ready" looks like.

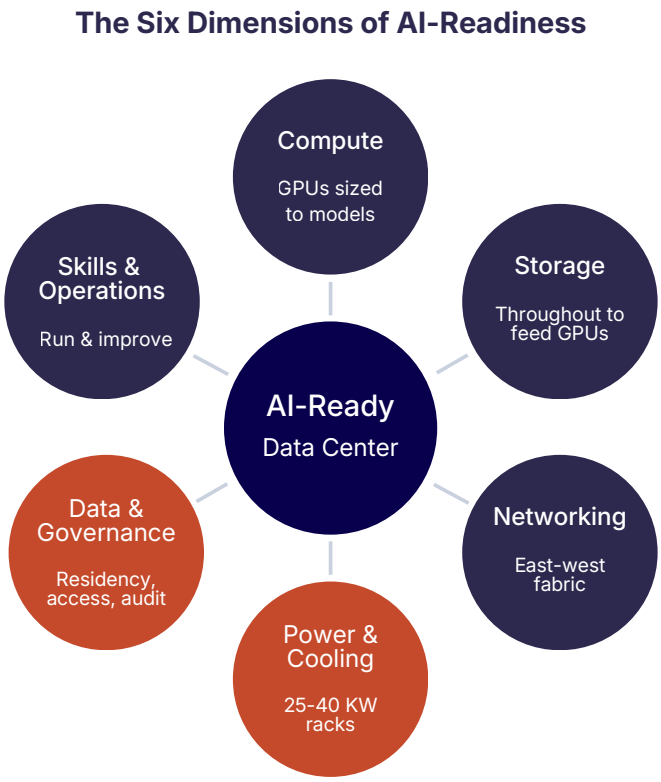


Figure 2: The six dimensions. Power & Cooling and Data & Governance are assessed first.

Dimension	What “Ready” Looks Like	Common Gap
Compute	GPU servers sized to your models and traffic	No GPUs, or GPUs sized by guesswork
Storage	High-throughput storage that keeps GPUs fed	Storage that starves the GPUs
Networking	Low-latency east-west fabric so GPUs scale	A north-south network that throttles clusters
Power & Cooling	Capacity for dense racks, often liquid cooling	A facility that cannot power or cool 25–40 kW racks
Data & Governance	Data pipelines, residency, access control and audit	No governance; unclear where data may live
Skills & Operations	The ability to run and improve the system	Nobody to operate it after launch

Two dimensions sink more AI programmes than the rest. Power and cooling, because AI racks now draw 25 to 40 kW, beyond what most enterprise server rooms can handle, so the facility, not the GPU, is often the real constraint. And data and governance, because a model answering on regulated data inherits that data’s obligations under India’s DPDP framework, whose substantive provisions take effect from 13 May 2027. Assess these two first.



Self-Assessment Scorecard

Score each dimension 0–4 against the maturity model. Total, then divide by six for your overall level.

Dimension	Your Score (0–4)	Evidence / Gap
Compute - GPUs sized to models		
Storage - throughput to feed GPUs		
Networking — east-west fabric		
Power & Cooling — 25–40 kW racks		
Data & Governance — residency, access, audit		
Skills & Operations — run and improve		
Average = your maturity level	___ / 4	

Watch out: Score below 2 on Power & Cooling or Data & Governance? Address those before anything else - they dictate your path.

3. The Staged Roadmap

With a baseline and an assessment, you close the gaps in sequence. The roadmap below delivers pilot capability early and defers heavy capital until a production workload justifies it. Each phase is testable and leaves the estate better than it found it.



Figure 3: Five phases from assessment to scale, with value delivered early.

Phase	Focus	Key Outcome
0. Assess and Baseline	Locate on the maturity model; assess the six dimensions	A gap map and a target level
1. Enable Pilots	Stand up pilot capability and a data foundation	Teams can build and test without waiting
2. Build the Production Base	Size and deploy the first production cluster; upgrade storage, network, power and cooling	A balanced environment that carries production
3. Govern and Secure	Residency, access control, audit, PII handling	A defensible, compliant environment
4. Scale and Operate	Expand capacity, instrument operations, optimise	AI capacity that grows with demand

Worked example (illustrative): A manufacturer running a Level 2 data center in Pune bought a few GPU servers for a vision-inspection pilot, but its 6 kW-per-rack server room was never designed for them. The pilot works; production stalls. Following the blueprint, the team scores the six dimensions, finds power and cooling and governance are the binding constraints, and colocates the production cluster in a Mumbai facility rather than rebuilding the plant's server room. Storage and networking are upgraded to match, residency is designed in, and capacity scales in steps. Eighteen months on, it runs production AI at Level 3, with no forklift rebuild.

4. The Decisions You Will Face

At each stage, a few decisions shape everything downstream. Resolve them deliberately rather than by default.

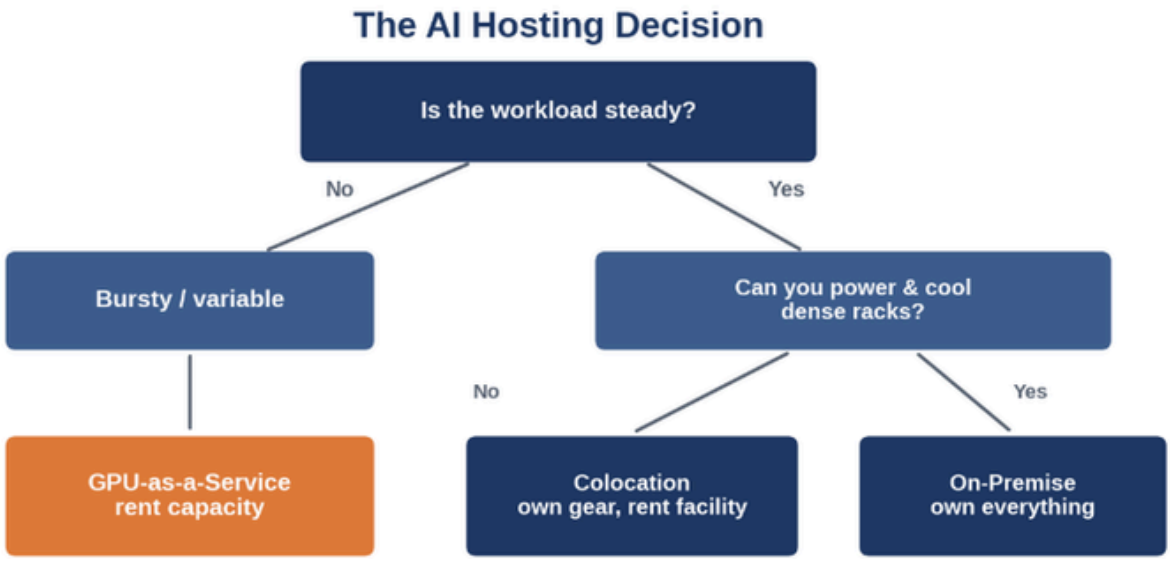


Figure 4: The hosting decision — rent, colocate or build.

- Where to host. Rent GPU capacity for bursty work, colocate owned hardware when you want ownership without a suitable facility, and build on-premise for sustained, sensitive workloads where you have the facility. Most enterprises blend these.
- On-premise or sovereign. For regulated data, decide early whether residency requires an on-premise or sovereign deployment, because it shapes the architecture rather than bolting on later.
- How to size. Size from the workload outward, the models and concurrency set the GPU count, which sets storage and network, which sets power and cooling, and validate with a benchmark.
- Whether the facility can carry it. Confirm power and cooling before committing to GPUs. If the facility cannot, colocation is the bridge, and discovering this early saves a stranded purchase.

5. Budgeting and Sequencing

The blueprint is designed to phase capital, not front-load it. Phase 1 is deliberately light: pilot capability is often cheapest on rented or modest hardware. The capital concentrates in Phase 2, the production base, and should be released only once a pilot has proved a use case worth carrying. Avoid the temptation to buy a large cluster up front “to be ready”; an over-bought cluster running at low utilisation is the most common waste in AI infrastructure. Size to the first real production workload, then scale in steps as demand grows.

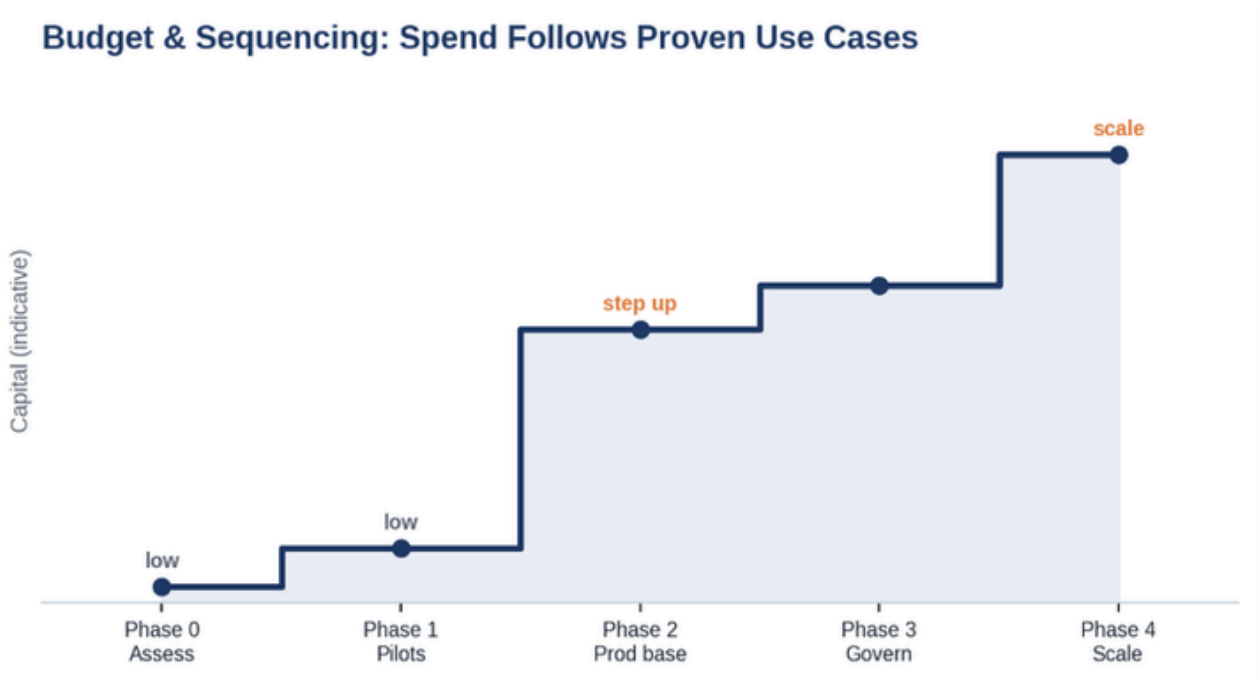


Figure 5: Capital follows proven use cases, concentrating at Phase 2 and scaling in steps.

6. Common Pitfalls to Avoid

Five recur often enough to plan against.

- Buying GPUs before checking the facility can power and cool them, which strands the purchase.
- Sizing the GPUs while ignoring storage and networking, which leaves expensive accelerators idle.
- Treating governance and residency as a later phase, then retrofitting them across a live system at multiplied cost.
- Over-buying capacity to “be ready”, which destroys the utilisation that justifies ownership.
- Building the infrastructure without the operations to run it, so a system that launched well decays unattended.

Key takeaway: Each pitfall is avoidable by following the staged sequence and assessing the weakest dimensions first.

7. Your First 90 Days

You do not need budget approval to begin. In the first 90 days you can complete the assessment and produce a costed plan.

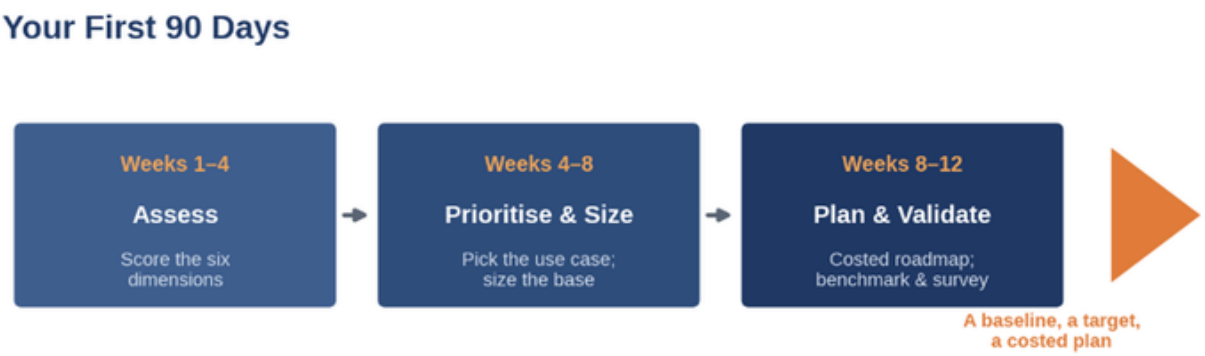


Figure 6: A 90-day path to a baseline, a target and a costed plan.

- Weeks 1–4: locate the data center on the maturity model and assess the six dimensions, paying closest attention to power and cooling and to data and governance.
- Weeks 4–8: identify the priority use case and size the production base it would require, then decide the hosting model and the residency posture.
- Weeks 8–12: produce a phased roadmap with a budget and quick wins, and validate the riskiest assumptions, facility capacity and GPU sizing, with a survey and a benchmark.

Putting the Blueprint to Work

A blueprint is a plan; the value is in the build. Turning this into a working AI-ready data center, with the facility confirmed, the six dimensions balanced, governance designed in and operations in place, is where experience saves the cost a first attempt usually wastes.

Proactive Data Systems designs, builds and operates AI-ready data centers for Indian enterprises, on-premise, hybrid and sovereign. We are a Cisco Preferred Cloud and AI Partner, Dell Platinum Partner and NetApp Preferred Partner, with 35 years in enterprise IT, more than 1,500 organisations served, and a 24/7 service desk in India. We run the assessment, design the staged roadmap, build each phase without a forklift rebuild, and operate the result. To begin, you can ask Proactive for an AI-readiness assessment.

Frequently Asked Questions

- **What Does It Mean for a Data Center to Be AI-Ready?**

  An AI-ready data center can run AI training and inference dependably in production, not just in a pilot. It has GPU compute sized to its models, storage and networking to keep the GPUs fed, power and cooling for dense racks, and the data governance and operations to run it. Readiness is measured across six dimensions, and the weakest caps the rest.

- **Can I Make My Existing Data Center AI-Ready Without Rebuilding It?**

  Yes, in almost all cases. AI-readiness is achieved in stages: enable pilots cheaply, build a production base sized to a real workload, add governance, then scale. The facility's power and cooling is the most common hard constraint; where it cannot carry dense GPU racks, colocation bridges the gap without a full rebuild.

- **What Is the Biggest Obstacle to an AI-Ready Data Center?**

  Usually power and cooling. AI racks can draw 25 to 40 kW, well beyond what most enterprise server rooms can power or cool, so the facility is often the real constraint rather than the GPUs. Data governance and residency under DPDP is the second most common, and both should be assessed first.

- **How Long Does It Take to Become AI-Ready?**

  The assessment and a costed plan take about 90 days, with no budget approval needed to start. Implementation is phased: pilot enablement can be quick, while the production base, governance and scale follow as use cases justify them.

- **Should AI Infrastructure Be On-Premise or Sovereign?**

  For regulated or sensitive data, decide early. On-premise gives maximum control, and sovereign deployment keeps data and models in-country to meet DPDP and sector requirements. The choice shapes the architecture, so it belongs in the assessment phase rather than being retrofitted later.

Glossary

Term	Definition
AI factory	Infrastructure built to produce AI (train and run models) as its primary workload, rather than to store and serve data.
AI-ready	A data center able to run AI training and inference dependably in production across all six readiness dimensions.
Training vs inference	Training builds or adapts a model and is memory-heavy; inference serves the finished model and needs far less per copy.
GPU	A graphics processing unit; the accelerator that does the heavy computation in AI training and inference.
East-west fabric	The low-latency network between servers and GPUs inside a data center, which AI clusters depend on.
Liquid cooling	Cooling that uses liquid rather than air, increasingly needed for dense GPU racks drawing 25–40 kW.
RAG	Retrieval-augmented generation; grounding a model's answers in your own documents via a retrieval pipeline.
Sovereign AI	Training and running AI on infrastructure you control, with data kept in-country and inside your governance boundary.
GPU-as-a-Service	Renting GPU capacity by the hour from a provider, owning no hardware.
Colocation	Owning the hardware but housing it in a third-party data center that provides space, power and cooling.
DPDP	India's Digital Personal Data Protection framework; substantive obligations take effect 13 May 2027.

Disclaimer: *This blueprint provides general guidance on AI infrastructure and references to India's data-protection framework. It is not legal, financial or compliance advice, and is not a quote. Costs, specifications and regulations vary and change. Validate facility capacity and GPU sizing with a survey and benchmark, model costs on your own requirements, and confirm compliance obligations with qualified counsel before committing budget.*

Proactive Data Systems is an enterprise IT infrastructure and services company. For more than three decades, the company has designed, deployed, and operated the technology foundations behind some of India's most demanding enterprise environments, from regulated industries to large distributed operations.

Founded in 1991, Proactive serves more than 1,500 customers, from fast-growing businesses to Fortune 500 enterprises, across banking and financial services, manufacturing, healthcare, public sector, IT and ITES, retail, education, and global capability centres. Our work spans enterprise networking, cybersecurity, cloud and AI-ready data centres, collaboration platforms, structured cabling, and managed services.

A Cisco Preferred Partner under the Cisco 360 Partner Program, Proactive is recognised across the Networking, Security, Collaboration, Cloud & AI, and Services portfolios, and is a Cisco Advanced Customer Experience Specialised Partner. The company is ISO 9001:2015 and ISO/IEC 27001 certified, and engineers solutions in partnership with Dell, Nutanix, NetApp, Palo Alto Networks, Fortinet, VMware, Veeam, and CommScope.

Headquartered in Delhi NCR, with offices in Mumbai, Bengaluru, Pune, Hyderabad, Indore, and Singapore.

To learn more about our solutions, write to us
at pro@proactive.co.in

www.proactive.co.in

PROACTIVE™
Enhancing Digital Experience