



AN AJEEVI BOOK

Make & Made in India Data Centers

A practical guide to building, powering and running
data centers with Indian technology

First edition | October 2026 | Draft for review

AJEEVI Technologies

HumanTech for a smarter world | First edition, 2026

www.ajeevi.com

ABOUT THIS BOOK

Read this first

A practical guide to building, powering and running data centers with Indian technology.

HOW TO USE IT

The book has twelve short chapters in four parts. You can read them in order or jump to the part you need.

- **Part A, Context (chapters 1 to 3):** the market, what "Make in India" and "Made in India" mean for a data center, and the rules you must design for.
- **Part B, Build (chapters 4 to 7):** sites and power, design, cooling and sustainability, and what it costs.
- **Part C, Intelligence (chapters 8 to 10):** AI and IoT operations, AI-ready facilities, and edge data centers for cities.
- **Part D, Deliver (chapters 11 and 12):** security and resilience, and a practical roadmap with checklists.
- The contents page and the PDF bookmarks link to every chapter.

DISCLAIMER

This book is for general information. It summarizes public reports, regulations and common engineering practice as of October 2026. It is not legal, financial, tax or engineering advice, and it is not a quotation or a guarantee of results. Regulations, tariffs, incentives and market figures change. Cost and design figures are indicative planning ranges. Always confirm current requirements with qualified professionals and the primary documents before making decisions.

Market figures are quoted from the sources listed at the end, which use different definitions and dates. Where they differ, both figures are given.

DCIM, Video Analytics and Helpdesk are products of AJEEVI Technologies. Other names and marks belong to their owners and are used for identification only.

ABOUT THIS EDITION

Make & Made in India Data Centers. First edition, draft for review, October 2026. Published by AJEEVI Technologies, India. www.ajeevi.com. ISBN: to be assigned.

© 2026 AJEEVI Technologies. All rights reserved. No part of this book may be reproduced or transmitted in any form without written permission from the publisher, except for short quotations with credit.

REVIEWER NOTES

First draft prepared for AJEEVI review. Before publication, please check company statements, add case studies and photographs that AJEEVI can stand behind, confirm figures against the primary documents, and have counsel review the regulatory chapter.

CONTENTS

What is inside

Twelve short chapters in four parts. Read them in order or jump to the part you need.

Read this first	2
Foreword	4
PART A CONTEXT	
1 Why India, why now	6
2 What "Make" and "Made" in India really mean	8
3 The rules you design for	10
PART B BUILD	
4 Choosing a site and securing power	13
5 Design fundamentals	15
6 Cooling and sustainability	17
7 What it costs and how long it takes	19
PART C INTELLIGENCE	
8 Intelligent operations	22
9 Building for AI	24
10 Edge data centers for cities	26
PART D DELIVER	
11 Security and resilience	29
12 A practical roadmap	31
BACK MATTER	
Glossary	33
Sources	34
About AJEEVI	35

FOREWORD

Foreword

A data center made in India, and made for India.

Every digital service a city depends on, from a water bill to an emergency call, ends up in a building full of racks. India is now building those buildings faster than at any time in its history. Operational capacity has grown from roughly 400 MW in 2019 to between 1.6 and 1.8 GW in mid-2026, and about 3.9 GW more is under construction or planned.

Much of what fills those buildings is still designed somewhere else. AJEEVI builds AI, IoT and software for urban infrastructure, so we sit close to the people who consume data center capacity: city command centers, utilities and transport systems. We see the other end of the cable. This book writes down what a data center should look like if it is both **made in India**, built with Indian engineering, manufacturing and intellectual property, and **made for India**, shaped by Indian heat, power, regulation and users.

THE BOOK IN ONE PAGE

1. **Demand is real.** About 1.6 to 1.8 GW operates today, about 3.9 GW is in the pipeline, and one major adviser expects 6 GW by 2029.
2. **Power decides the project.** Getting grid power can take 12 to 24 months. Electrical systems are the largest share of capital cost, and electricity is the largest running cost.
3. **"Made in India" is a spectrum.** Measure it with a local content and IP scorecard instead of a slogan.
4. **Compliance is a design input.** The DPDP Act, CERT-In directions and state policies shape site, logging, security and contracts.
5. **Efficiency is a competitive edge.** Track PUE and water use from day one, and plan liquid cooling for AI racks.
6. **Intelligence pays back.** Sensors, AI and digital twins cut risk and cost, and keeping data and models in India keeps you sovereign.

THE IDEA IN ONE LINE

A data center made in India is built with Indian engineering and ownership, and made for India's heat, power, rules and users.

A note on numbers: market figures come from public reports by JLL, Cushman & Wakefield, CEEW and others, dated as shown. Reports define capacity differently, so figures differ. Cost and design figures are indicative ranges for planning, not quotations.

A

PART A | CHAPTERS 1 TO 3

Context

The market, what Make in India and Made in India mean for a data center, and the rules you must design for.



PART A | CHAPTER 1

Why India, why now

Capacity has grown several times over in seven years, and the next four years are planned to add as much again.

A data center is the physical home of cloud, AI, payments, e-governance and city operations. In 2019 India had about 400 MW of operational capacity. By the first half of 2026, Cushman & Wakefield counted nearly 1.8 GW and ranked India the third-largest data center market in Asia Pacific. JLL put mid-2026 capacity at 1.6 GW. Both agree on the direction: capacity has more than tripled, and the pipeline is larger than the installed base.




Figure 1. India data center capacity and pipeline (IT load). Sources: Cushman & Wakefield, CEEW, JLL.

THE MARKET IN TEN NUMBERS

Indicator	Figure	Source
Operational capacity, mid-2026	About 1.6 GW (JLL); nearly 1.8 GW (C&W)	JLL, C&W
Capacity in 2019	About 400 MW	C&W
Pipeline, under construction and planned	Nearly 3.9 GW	C&W
Target and investment need	6 GW by 2029; about US\$110 billion	JLL
Demand absorbed in H1 2026	101 MW, over 20% above the three-year average	JLL
Share pre-committed by hyperscalers	82% of H1 2026 absorption	JLL
New supply in H1 2026	85 MW; supply growing about 60% a year since 2021	JLL
Typical project size	Grown from about 20 MW to over 50 MW	JLL
Share of national electricity	About 0.5%	CEEW
Water used each year	About 150 billion litres	CEEW

Reports measure different things: IT load or facility power, operational or commissioned capacity, colocation only or all types. Use one definition in your own business case and contracts.

WHAT IS DRIVING DEMAND

- **Cloud and enterprise digitization.** More workloads run in the cloud, and enterprises move out of server rooms into colocation.
- **AI.** JLL reports that recent absorption is driven by high-density AI computing. Racks for AI draw several times the power of a classic rack (chapter 9).
- **Regulation.** Data protection and sector rules push personal and financial data to be stored and handled in India (chapter 3).
- **Government and city systems.** E-governance, digital public infrastructure and integrated city command centers all need reliable, local capacity.
- **Hyperscaler self-build.** JLL notes more than US\$50 billion in commitments for about 1.4 GW of self-built capacity by 2029, close to a third of new capacity.

WHAT MAKES INDIA DIFFERENT

- **Heat and humidity.** Warm, humid air limits how much "free" outside-air cooling works through the year, so cooling is a larger share of energy than in cooler countries.
- **Power.** Grid quality and tariffs vary by state, and policy is pushing hard toward renewables, with a national target of 500 GW of renewable capacity before 2030.
- **Monsoon and coast.** Flooding, corrosion and storms are design inputs in many of the leading cities.
- **People.** India has deep engineering and software talent, which is the foundation for a Made in India approach to design, operations and software.

WHAT THIS MEANS FOR A BUILDER

Plan for a market where good land with power is the scarce resource, where large buyers sign early, and where policy and sustainability expectations are rising faster than the rulebook.

PART A | CHAPTER 2

What "Make" and "Made" in India really mean

Two phrases are used loosely. This chapter pins them down so you can write them into a tender.

Procurement documents often say "Make in India" when they mean four different things. Separating them makes your specification clearer and your evaluation fairer.

	Make in India	Made in India
Meaning	Built, assembled or integrated in India	Designed, engineered and owned in India
Who benefits	Local factories, installers and jobs	Local intellectual property, source code and long-term support
Typical examples	A transformer or switchgear built in an Indian plant; servers assembled in India	Monitoring software, firmware, AI models and IoT devices designed by Indian engineers
Risk if used alone	Foreign design and key parts; updates depend on an overseas vendor	Small volumes and a short track record

The strongest projects aim for both: proven manufacturing at scale, together with Indian ownership of the design and software that run the facility.

SIX LAYERS OF A DATA CENTER

6 Software and intelligence	DCIM, BMS, monitoring, analytics, AI, digital twin
5 IT equipment	Servers, storage, network, racks
4 Cooling (mechanical)	Chillers, CRAH units, liquid cooling, pumps, towers
3 Power (electrical)	Transformers, switchgear, UPS, batteries, gensets
2 Shell and civil	Land, building, structure, fire, access, security
1 Services and people	Design, build, operate, maintain, certify

Figure 2. The six layers of a data center, from services to software and intelligence.

Indian supply is deepest in the lower layers: construction, electrical equipment and services, where a mature manufacturing and contracting base exists. Cooling equipment and IT hardware are growing, helped by incentives. Advanced AI accelerators and the most advanced chips are still largely imported as of this edition. Software and intelligence is where Indian engineering talent is strongest, and where ownership of source code and data matters most.

THE GOVERNMENT RULES THAT APPLY

The Public Procurement (Preference to Make in India) Order, first issued in 2017 and revised in 2020, gives purchase preference to local suppliers in government procurement. It sorts suppliers by how much of the price is local value added:

- **Class-I local supplier:** local content of 50% or more.
- **Class-II local supplier:** local content of 20% or more and under 50%.
- **Non-local supplier:** local content under 20%.

FORMULA

$$\text{Local content (\%)} = (\text{Price of the item} - \text{Value of imported content}) \div \text{Price of the item} \times 100$$

Illustration: a ₹100 power package contains ₹38 of imported parts. Local content is $(100 - 38) \div 100 = 62\%$, so the supplier qualifies as Class-I. Procurement rules also restrict global tenders for smaller government purchases. Always check the latest notification for your product category and the minimum local content set by the concerned ministry.

On hardware, the revised Production Linked Incentive scheme for IT hardware carries an outlay of ₹17,000 crore. Press reports at the time of approval named 27 companies, including global brands such as Dell, Lenovo and HP, which means more servers and laptops are now built in India.

A MADE IN INDIA SCORECARD

Score each bidder from 0 (none) to 5 (full) on each criterion, multiply by the weight, and add up. The weights below are a starting point; change them to suit your project.

Criterion	Question to ask	Weight
Local content	What share of the price is local value added, and can the supplier prove it?	25%
IP and source code	Who owns the design and software, and is the source code available or escrowed?	20%
Support in India	Are spares, engineers and training based in India with defined response times?	20%
Supply-chain resilience	Are there alternate sources for critical components?	15%
Security assurance	Is there independent testing, a software bill of materials and patch discipline?	10%
Data residency	Do data, logs and models stay in India under your control?	10%

PART A | CHAPTER 3

The rules you design for

Compliance is not paperwork added at the end. It decides where you build, what you log and what you promise customers.

DATA PROTECTION: THE DPDP ACT AND RULES

The Digital Personal Data Protection Act, 2023 sets duties for anyone who handles personal data. The Ministry of Electronics and Information Technology notified the DPDP Rules, 2025 on 14 November 2025. JLL reports phased compliance dates of 13 November 2025, 13 November 2026 and a full compliance deadline of 13 May 2027, with penalties of up to ₹250 crore. The Rules set conditions for transferring personal data outside India, and the Act allows the government to restrict transfers to notified countries.

Data center operators are usually processors, not the data fiduciary that owns the customer relationship. In practice, their customers will pass obligations down through contracts: security safeguards, breach reporting, deletion on request and, often, a promise to keep data in India. Design for that from day one.

CYBER SECURITY DIRECTIONS FROM CERT-IN

The Indian Computer Emergency Response Team issued directions on 28 April 2022 that apply to data centers, cloud and hosting providers. In summary:

- Report specified cyber incidents to CERT-In within **6 hours** of noticing them.
- Keep ICT system logs for a rolling **180 days**, stored within Indian jurisdiction.
- Synchronize system clocks with the NTP servers of NIC or NPL.
- Keep accurate subscriber information for 5 years.

GOVERNMENT AND SECTOR RULES

- **Government workloads.** Departments are expected to use cloud services from providers empanelled by MeitY, whose facilities are audited. If you want government customers, design to the empanelment requirements.
- **Payments.** The Reserve Bank of India has required payment system data to be stored only in India. Other regulators for insurance, securities and telecom have their own rules.
- **State policies.** CEEW counts 15 state data center policies. Only 5 include sustainability provisions, so incentives and requirements differ sharply by state.
- **Tax.** JLL notes that the Union Budget 2026-27 provides a tax holiday until 2047 for foreign cloud providers that serve global customers from Indian data centers.

APPROVALS YOU WILL NEED

Approval	Typical authority	Start
Environmental clearance, where thresholds apply	SEIAA or MoEFCC	Early
Building plan approval	Local development authority or municipal body	Early
Fire NOC, provisional then final (NBC 2016, Part 4)	State fire services	Design and handover
Consent to establish and operate	State Pollution Control Board	Before build and before operation
Electrical inspectorate approval	State electrical inspectorate	Before energization
Power connection agreement or open access	State utility or SLDC	Day one
State data center policy registration	State IT department	Day one
Telecom licence, if you sell connectivity	DoT	If applicable

STANDARDS TO CITE IN YOUR DESIGN BRIEF

Common references include the Uptime Institute tier standards and TIA-942 for site infrastructure, the ISO/IEC 22237 series for facilities, ISO/IEC 27001 for information security, ISO 50001 for energy management, and the National Building Code of India for fire and life safety. Cite the exact edition in your tender so bidders price the same thing.

NOT LEGAL ADVICE

This chapter is a general summary. Rules change; confirm current requirements with a qualified adviser before you commit.

B

PART B | CHAPTERS 4 TO 7

Build

Sites and power, design fundamentals, cooling and sustainability, and what it costs.



PART B | CHAPTER 4

Choosing a site and securing power

The best site is the one where the grid, fibre, water and approvals can all arrive on time.

CHOOSING A CITY

Published reports describe the leading markets as follows. Verify current status before you shortlist.

City or region	What the reports say	Watch for
Mumbai	Largest market (C&W); an established core hub (JLL)	Land and power cost, monsoon flooding
Chennai	Established hub with strong recent supply additions (JLL)	Humidity and coastal corrosion
Pune	Established hub with AI compute demand (JLL)	Grid capacity in growth corridors
Hyderabad	Major secondary market (C&W); emerging hub (JLL)	Rapid land price rises
Delhi NCR	Major secondary market (C&W); emerging hub (JLL)	Summer peak loads, air quality
Visakhapatnam	Emerging destination for AI-focused investment (C&W)	Coastal exposure, new ecosystem

POWER: THE LONG POLE

Power drives schedule and cost. IMARC estimates that securing a greenfield power connection can take 12 to 24 months, while the whole build from planning to commissioning for a hyperscale facility takes about 30 to 48 months. Start the power application on day one.

Tariffs matter as much as capacity. IMARC puts landed cost from open access or group captive renewable power at about ₹4 to ₹7 per unit, against about ₹7 to ₹10 per unit from state utilities. CEEW reports that developers see up to about 80% renewable supply as feasible using solar and wind portfolios with battery storage, and expects battery storage to become central to data center power within 3 to 5 years.

WORKED EXAMPLE: THE ANNUAL POWER BILL OF A 10 MW FACILITY

IT energy = 10 MW × 70% utilization × 8,760 h = 61,320 MWh a year. Facility energy at PUE 1.5 = 61,320 × 1.5 = 91,980 MWh.

Tariff	Annual energy cost
₹5.5 per kWh (renewable open access)	91,980,000 kWh × 5.5 = ₹50.6 crore
₹8.5 per kWh (state utility)	91,980,000 kWh × 8.5 = ₹78.2 crore

WHY POWER STRATEGY MATTERS

The ₹27.6 crore a year gap between the two tariffs is why the power strategy is a financial decision, not only an engineering one.

LAND, WATER AND HAZARDS

- **Land.** A 25 to 50 MW facility typically needs 10 to 20 acres; a 100 MW campus needs 20 to 40 acres (IMARC).
- **Water.** A typical 100 MW hyperscale facility can use about 2 million litres a day for on-site cooling (IEA, quoted by CEEW). Check supply, quality and competing local demand.
- **Hazards.** Check flood maps, the seismic zone under the BIS code (IS 1893), heat, lightning and coastal corrosion. Place the disaster recovery site in a different seismic and flood zone.
- **Fibre.** Look for at least two diverse carriers entering the building from different directions.

A SITE SCORING SHEET

Score each shortlisted site from 0 to 5, multiply by the weight and add up. The weights are illustrative; set your own.

Criterion	Weight	What good looks like
Power availability and tariff	25%	Committed capacity, dual feeders, clear tariff path
Connectivity	20%	Two or more fibre routes, close to an internet exchange
Natural hazards	15%	Outside flood plain, lower seismic zone
Water	10%	Reliable supply or a design that needs little
Land cost and expansion room	10%	Room to double without moving
Policy and approvals	10%	State policy incentives, single-window clearance
Talent and ecosystem	10%	Skilled operators and vendors within reach

PART B | CHAPTER 5

Design fundamentals

Design is a set of trade-offs between availability, efficiency, cost and speed. Make them on purpose, and write them down.

START WITH THE WORKLOAD

Before drawing anything, fix four numbers: the IT load in megawatts, the density in kilowatts per rack, the availability the business needs, and the growth phases. CEEW, quoting Uptime Institute data, reports that average rack power has risen from about 1 to 3 kW in the early 2000s to about 8 kW today. A hall of 1,000 racks at 8 kW is therefore an 8 MW IT load. AI racks are far denser (chapter 9).

TIER LEVELS

The Uptime Institute tier system describes how much redundancy a site has. The availability figures below are the commonly cited ones; downtime is calculated as $(1 - \text{availability}) \times 8,760$ hours.

Tier	Concept	Availability	Downtime a year
I	Single path, no redundant components	99.671%	About 28.8 hours
II	Single path, redundant components	99.741%	About 22.7 hours
III	Multiple paths, concurrently maintainable	99.982%	About 1.6 hours
IV	Fault tolerant, multiple active paths	99.995%	About 26 minutes

Uptime Institute certifies design and operations. A tier is not a guarantee of uptime, and people and process matter as much as equipment.

WHICH TIER SHOULD YOU CHOOSE?

Most colocation and enterprise sites target Tier III, where any single piece of equipment can be maintained without shutting the room. Tier IV suits mission-critical systems, at a higher cost (chapter 7). Some AI training clusters accept lower redundancy because jobs save checkpoints and can restart, but inference services that face customers usually do not.

REDUNDANCY IN PLAIN WORDS

Notation	Meaning	Example for a load needing 4 UPS modules
N	Exactly what the load needs	4 modules; any failure means an outage
N+1	One spare unit	5 modules; survives one failure
2N	A full second, independent system	8 modules in two separate paths
2N+1	A full second system plus a spare	9 modules or more

THE ELECTRICAL CHAIN




Figure 3. The path of power from the utility to a rack, with backup generation.

Power enters at medium voltage (typically 11 kV or 33 kV in Indian utility networks), steps down through transformers, passes through the UPS and batteries that bridge short outages, and reaches the racks through power distribution units. Diesel generators start within seconds of an outage and carry the load for longer ones.

- **Fuel.** Specify fuel autonomy in hours and sign a refueling contract with a response time.
- **Batteries.** Lithium-ion takes less space and lasts longer than lead-acid, at a higher upfront price. Compare on total cost over the life of the facility.
- **Metering.** Meter at every level, from the utility feed down to each rack, so you can find waste and bill customers fairly.
- **Power quality.** Specify harmonics, earthing and surge protection in the tender, because the facility is full of electronic loads.

BUILD IN PHASES AND IN MODULES

Build the shell and the main power and cooling backbone for the full plan, but fit out the halls in phases of one to a few megawatts. Prefabricated power rooms, skid-mounted cooling and standard hall designs shorten construction and reduce risk. Contain hot and cold aisles so that hot exhaust air does not mix with cold supply air, a simple measure that raises efficiency.

A DESIGN BRIEF CHECKLIST

- IT load per phase, density per rack, and the growth plan for five years.
- Tier level and redundancy for power, cooling and network paths.
- Cooling approach for today and a path to liquid cooling for tomorrow.
- Target PUE and WUE, and how they will be measured.
- Fire detection and suppression strategy, with the approvals it needs.
- Security zones, access control and video coverage.
- Monitoring and control points, and who owns the data and the software.
- Commissioning plan, with integrated systems testing before handover.

PART B | CHAPTER 6

Cooling and sustainability

Every watt of IT load becomes heat. In a hot, humid country, how you remove it decides your efficiency and your water bill.

MEASURE WHAT MATTERS

PUE = Total facility energy ÷ IT equipment energy. **WUE** = Site water used (litres) ÷ IT equipment energy (kWh).

CUE = CO₂ emitted ÷ IT equipment energy (kWh).

A PUE of 1.5 means that for every 1 kWh used by servers, 0.5 kWh goes to cooling, power conversion and lighting. Targets of 1.3 to 1.5 are common for new builds in warm climates; check what your site's design weather allows. CEEW notes that most state policies still lack enforceable PUE, WUE or carbon benchmarks, and that the European Union already requires energy and water reporting for facilities above 500 kW. Expect similar rules in India and design for them now.

WORKED EXAMPLE: WHAT A BETTER PUE IS WORTH

A 10 MW IT load running at 70% uses 61,320 MWh a year. Improving PUE from 1.6 to 1.4 cuts facility energy from 98,112 MWh to 85,848 MWh.

Result	Value
Energy saved each year	12,264 MWh
Money saved at ₹6 per kWh	About ₹7.4 crore a year
CO ₂ avoided at 0.82 kg per kWh (grid factor used in AJEEVI calculators)	About 10,056 tonnes a year

COOLING OPTIONS

Densities below are indicative and depend on the design.

Option	Typical rack density	Water	Notes
Air cooling with aisle containment	Up to about 15 to 20 kW	Low to medium	Simple and proven; limited for AI racks
Rear-door heat exchangers	About 20 to 40 kW	Medium	Adds capacity to air-cooled halls
Direct-to-chip liquid cooling	About 40 to 100 kW or more	Low to medium	Needs coolant loops and leak detection; air still removes some heat
Immersion cooling	100 kW or more	Low	High efficiency; new maintenance practices and supply chain

CEEW describes the central trade-off: air-cooled systems use little or no water but more energy, while water-based systems are more energy-efficient but depend on a reliable supply. Direct-to-chip and immersion cooling can cut both water and power, though cost and supply constraints have slowed adoption.

WATER

CEEW estimates Indian data centers use about 150 billion litres of water a year, and expects both water and electricity use to more than double by 2030. A typical 100 MW hyperscale facility can use about 2 million litres a day for on-site cooling. Globally, about 43% of data centers operate in areas of high water stress this decade (S&P Global, quoted by CEEW). Practical responses include using treated wastewater for cooling towers, closed-loop designs, hybrid dry and wet systems, and publishing WUE.

CLEANER POWER

- Buy renewable power through open access or group captive arrangements, with a plan for the hours when solar and wind are low.
- Add battery storage. CEEW expects it to become central to data center power architecture within 3 to 5 years, and calls for clear rules on how storage earns revenue.
- Join demand response where utilities allow it, shifting flexible loads away from grid peaks.
- Report honestly. Publish PUE, WUE and the renewable share of your supply each year.

SUSTAINABILITY CHECKLIST

Meter every system. Set PUE and WUE targets in the design brief. Choose a cooling path that handles today's racks and tomorrow's. Source renewable power. Reuse treated water. Publish the results.

PART B | CHAPTER 7

What it costs and how long it takes

Plan with ranges, not single numbers, and spend most of your attention on power and electrical systems.

The figures below are planning ranges published by IMARC Engineering in its 2026 guide for India. They exclude land unless stated, and will move with specification, location and market conditions. Get current quotations before you commit.

COST PER MEGAWATT OF IT LOAD

Facility type	Cost per MW	Approx. in US\$
Tier III colocation, 5 to 10 kW per rack	₹50 to 65 crore	6 to 8 million
Tier III hyperscale, optimized design	₹55 to 70 crore	n/a
Tier IV mission-critical	₹70 to 90 crore	n/a
AI-ready, 20 to 50+ kW per rack, liquid cooling	₹70 to 95 crore	n/a

Whole project, including land	Total cost
50 MW hyperscale facility	₹3,000 to 4,500 crore
100 MW campus	₹6,000 to 9,000 crore
Large multi-phase campus	₹15,000 crore and above

WHERE THE MONEY GOES




Figure 4. Typical share of total capital cost by category. Source: IMARC Engineering, 2026.

Electrical systems are the largest line at 25 to 35% of capital cost, followed by mechanical and cooling at 18 to 28%. Together they are more than half the budget, which is why equipment choices, local manufacturing and early orders for long-lead items such as transformers, switchgear and generators have more effect on cost and schedule than almost anything else.

HOW LONG IT TAKES

Phase	Typical duration
Concept and feasibility	8 to 16 weeks
Site acquisition and approvals	6 to 12 months
Detailed engineering	4 to 8 months
Construction (EPC)	18 to 30 months
Testing and certification	3 to 6 months
Overall for a hyperscale facility	About 30 to 48 months, because phases overlap

BUILD, LEASE OR COLOCATE?

Option	Best for	Trade-off
Build your own	Large, stable, specialized needs; full control	Highest capital, longest schedule
Wholesale lease	Hyperscale or large enterprise needs	Long commitments; less control
Retail colocation	Enterprise and government systems of moderate size	Shared facility; pay per rack and power
Public cloud	Variable or fast-growing workloads	Running cost can exceed owning at steady scale; data residency to check

WORKED EXAMPLE: A 10 MW TIER III COLOCATION FACILITY

At ₹50 to 65 crore per MW, the build is about ₹500 to 650 crore. Electricity from the example in chapter 4 adds ₹50.6 to 78.2 crore a year, or about ₹253 to 391 crore over five years. That is between about 39% and 78% of the build cost, depending on tariff and build cost, which shows why power strategy belongs in the business case from the start.

LEVERS THAT CUT COST AND RISK

Order long-lead electrical equipment early. Standardize halls and modules. Prefer proven local suppliers where they meet the specification. Keep contingency of 8 to 12%, as IMARC suggests. Commission in stages and test the full system before handover.

C

PART C | CHAPTERS 8 TO 10

Intelligence

AI and IoT operations, AI-ready facilities, and edge data centers for cities.



Make & Made in India Data Centers | AJEEVI Technologies

PART C | CHAPTER 8

Intelligent operations

A data center produces more data about itself than almost any other building. Most of it is thrown away.

Running a facility well is a data problem. Heat, power, vibration and movement can all be measured, and the measurements can tell you what will fail, what is wasting energy and who is where. This is the layer where AJEEVI works: AI, IoT and software for urban infrastructure, applied to the buildings that hold a city's data.



Figure 5. The operations loop: sense, connect, understand, decide and act.

WHAT TO SENSE

What	Typical measurements	Why it matters
Rack inlet air	Temperature and humidity, at bottom, middle and top	Finds hot spots before servers throttle or fail
Power	Metering at feeds, branches and rack PDUs	Shows load, waste and capacity left
Cooling plant	Supply and return temperature, flow, pressure, valve and fan status	Lets you tune plant for efficiency
UPS and batteries	Voltage, temperature, internal resistance	Warns of weak batteries before an outage
Rotating equipment	Vibration and motor current on chillers, pumps and fans	Enables predictive maintenance
Leaks and environment	Water, smoke, door, particulates	Protects equipment and people
Security	Cameras, access logs, intrusion	Verifies who is in each zone

A common practice is to place three inlet sensors on every second or third rack. For a 1,000-rack site, three sensors on every third rack is about 1,000 temperature points.

WHERE AI EARNS ITS PLACE

Use case	How it works	Benefit
Predictive maintenance	Models learn normal vibration, temperature and current, and flag drift	Repairs planned before failure; fewer outages
Cooling optimization	Models adjust setpoints, fan speeds and chiller sequence from live load and weather	Lower PUE without risking temperature limits
Anomaly detection	Spots odd power draws, thermal patterns or access events	Earlier warning, fewer false alarms
Capacity planning	Forecasts space, power and cooling headroom	Better timing of expansion capital
Video analytics	Detects intrusion, tailgating, vehicles and safety gear at gates and halls	Stronger security with fewer guards on watch

Results differ from site to site. Measure a baseline for a few months first, then compare.

DIGITAL TWIN AND 5D SIMULATION

A digital twin is a live 3D model of the facility that is updated by its sensors. Operators can ask "what if" before they act: what happens to inlet temperatures if a cooling unit fails, or if 200 more racks go into Hall 2? A 5D simulation adds time and cost to the 3D model, so the construction team can test the sequence of work, find clashes between pipes, trays and ducts, and see the effect of a delay on schedule and budget before a single wall is built. Both cut surprises, and both are cheapest to start at design stage.

ONE COMMAND CENTER FOR MANY SITES

Operators with several facilities, including city edge sites, benefit from a single command center that brings alarms, video, tickets and field teams into one view. It is the same idea as an integrated city command center, applied to infrastructure that has to stay up.

HOW AJEEVI APPROACHES THE DIGITAL LAYER

- **Own the stack.** AI, IoT devices and software are developed in-house, so updates and support do not depend on an overseas vendor.
- **Rugged devices, designed and built in India.** Sensors and gateways connect over LPWAN, cellular and eSIM, and work in harsh conditions.
- **Intelligence that stays in the country.** Data, models and operations can run inside India, which supports sovereign AI and compliance.

GETTING STARTED IN FIVE STEPS

1. List the five failures that would hurt most, and the data that would have warned you.
2. Install meters and sensors where that data comes from, starting with power, inlet air and cooling plant.
3. Bring the data into one platform with clear ownership and retention rules.
4. Run alerts and dashboards for a few months to build a baseline.
5. Add models for prediction and optimization, and measure savings against the baseline.

Where AJEEVI fits | Products: DCIM, Video Analytics, Helpdesk | IoT devices: GyanYantram

PART C | CHAPTER 9

Building for AI

AI changes the numbers: more power in each rack, more heat in each row and new supply chains behind both.

WHAT CHANGES WITH AI RACKS

A classic enterprise rack draws 5 to 10 kW. AI racks that hold GPU servers draw 20 to 50 kW or more, and the newest designs go higher. As a rough rule, an eight-GPU server draws about 10 kW or more, so a rack with several of them quickly passes what air cooling can handle.

	Classic hall	AI-ready hall
Rack density	5 to 10 kW	20 to 50+ kW
Cooling	Air with containment	Liquid cooling, often with air for the rest
Capital cost per MW (IMARC)	₹50 to 65 crore	₹70 to 95 crore
Network	Standard switching	High-bandwidth, low-latency fabric
Floor and structure	Standard loading	Higher loading, pipe and cable space
Schedule	Well understood	Longer, because of liquid loops and testing

WORKED EXAMPLE: SIZING A 1,024-GPU CLUSTER

$1,024 \text{ GPUs} \div 8 = 128 \text{ servers}$. $128 \times 10 \text{ kW} = 1.28 \text{ MW}$. Add 15% for network and storage = 1.47 MW of IT load. At PUE 1.3, facility power = $1.47 \times 1.3 = \text{about } 1.9 \text{ MW}$.

Illustrative: assumes about 10 kW per eight-GPU server. Use your vendor's figures for a real design.

LIQUID COOLING IN PRACTICE

- **Direct-to-chip loops** carry coolant to cold plates on the chips. Coolant distribution units (CDUs) transfer the heat to a facility water loop.
- **Hybrid halls** use liquid for the hottest components and air for the rest, so plan both.
- **Leak detection and isolation valves** are part of the design, not an add-on.
- **Skills.** Operations teams need training in fluid handling, filtration and commissioning of liquid loops.

INDIA'S AI COMPUTE PUSH

The IndiaAI Mission, approved in March 2024 with an outlay of about ₹10,300 crore, is building shared national compute. On 19 February 2026, at the India AI Impact Summit in New Delhi, the Union Minister for Electronics and IT said the country's existing base of 38,000 GPUs would be expanded by another 20,000, to about 58,000. JLL notes Pune as an AI compute hub among established markets, and Cushman & Wakefield highlights Visakhapatnam as an emerging destination for AI-focused investment.

SOVEREIGN AI NEEDS SOVEREIGN INFRASTRUCTURE

Sovereign AI means that the data, compute, models and operations behind an AI service sit under national control. A data center made in India is where that becomes physical.

Pillar	Question to ask
Data	Does training and inference data stay in India, under Indian law and your keys?
Compute	Is the capacity in Indian facilities, and can you secure the accelerators you need?
Models	Are models trained, hosted and updated in India, with weights you control?
Operations	Are administrators, support and keys under Indian control, with audit logs kept in India?

DESIGN FOR THE NEXT TWO HARDWARE GENERATIONS

A building lasts 25 years or more; a GPU generation lasts two or three. Leave space for liquid pipes in risers and ceilings, size electrical busways and structure for higher density, keep spare conduits, and avoid designs that tie you to one vendor's equipment. Most leading AI accelerators are still imported, so plan the supply chain as carefully as the power.

PART C | CHAPTER 10

Edge data centers for cities

Cities create data at the street: cameras, meters, sensors and vehicles. Not all of it should travel to a distant core.

A modern city runs on connected systems: command centers, traffic signals, water networks, street lights, public safety and utilities. They share three needs. They need answers fast, because a signal or a valve cannot wait for a round trip to a faraway region. They produce data in volumes that are costly to move. And they must keep working when a link or a site fails. Edge data centers, placed close to the city, meet those needs.

Core and hyperscale Cloud regions, large AI clusters, national platforms	20 to 100+ MW
Regional and colocation Enterprise, government, DR sites, state data centers	2 to 20 MW
Urban edge City command center, video analytics, traffic and utilities	100 kW to 2 MW
Micro and roadside Cabinet at a junction, pumping station or substation	10 to 100 kW

Figure 6. Four sizes of facility, from the core to the roadside cabinet.

SIZING CITY VIDEO

Video is usually the largest data stream in a city system. The example below uses the same assumptions as the video calculator on the AJEEVI website (www.ajeevi.com/calculators.html).

Item	Value
Cameras	1,000 at 2 MP, H.265 compression, about 2 Mbps each
Network if all stream live at once	2 Gbps (1,000 × 2 Mbps)
Storage per camera per day	21.6 GB for continuous recording
30 days, continuous, with 20% redundancy	777.6 TB
30 days, AI event-based at 40% recording, same redundancy	About 311 TB

Event-based recording depends on AI that decides what is worth keeping, and on rules about evidence and retention. Agree those rules with the agency that owns the footage.

DESIGN PRINCIPLES FOR EDGE SITES

- **Small, standard and repeatable.** One design copied many times is cheaper to build and easier to support than many custom sites.
- **Managed remotely.** Plan for few visits: remote power cycling, environmental alarms and video of the room.
- **Hardened for the street.** Dust, heat, humidity and monsoon rain are the norm. Use sealed, rated enclosures and filtered cooling.
- **Backed up for your outages.** Size batteries and generators to the outage patterns you actually see, and plan a second network path, for example fibre plus cellular.
- **Locked down.** Physical locks, tamper alarms, access logs and encrypted links, since many sites are unattended.

WHERE THE PROCESSING HAPPENS

A sound pattern is to analyze at the edge, store at the regional site and learn at the core. Edge equipment runs video analytics and filters data, so only events and summaries travel upward. Raw video stays at the edge or regional site for its retention period. The core trains models on selected data and sends improved models back to the edge. Disaster recovery sits in a separate region, so that one flood or power failure cannot take down both.

GOVERNANCE MATTERS AS MUCH AS HARDWARE

Decide who owns the data from each camera and sensor, how long it is kept, who may see it and how access is logged. Logs kept within India for 180 days are a CERT-In requirement for covered providers (chapter 3), so build the logging in from the start.

Where AJEEVI fits | Video analytics: **GyanAkshi** | Command center: **GyanKendram** | IoT devices: **GyanYantram**

D

PART D | CHAPTERS 11 AND 12

Deliver

Security and resilience, and a practical roadmap with checklists.



Make & Made in India Data Centers | AJEEVI Technologies

PART D | CHAPTER 11

Security and resilience

A data center must keep out intruders, keep running through failures and recover when something goes wrong.

Security and resilience are one discipline seen from two sides: one asks "who can get in," the other asks "what happens when something fails." Plan both in layers, and test them.

Layer	What to do	How to check
Physical	Perimeter, controlled entry, mantraps, zones with escalating access, video analytics at gates and aisles	Walk-through audits, access log review
Network and IT	Segmentation, zero-trust access, encryption, hardened builds, patching routines	Penetration tests, patch reports
Building systems (OT)	Keep power, cooling and BMS networks separate from the corporate network; restrict remote access	Network diagrams, firewall rule reviews
Supply chain	Trusted sources, software bill of materials, signed firmware, source-code escrow where needed	Vendor audits, scorecard (chapter 2)
People and process	Background checks, role-based access, change control, incident runbooks	Drills and post-incident reviews
Resilience	Redundant paths, tested failover, separate disaster recovery site in a different seismic and flood zone	Scheduled failover tests

PRINCIPLES THAT HOLD UP

- **Assume breach.** Limit what any one account, device or site can reach, and log everything that matters.
- **Prove it works.** Test failover and recovery on a schedule, such as twice a year, with real switchovers where the business can tolerate them.
- **Keep control of keys and logs.** Encryption keys, audit logs and administrator rights should sit under your control and, for regulated data, inside India.
- **Plan for people.** Many outages begin with a mistake in a change. Use change windows, peer review and rollback plans.
- **Report fast.** With a 6-hour CERT-In reporting window, have a named owner, a contact list and a template ready before an incident.

RECOVERY OBJECTIVES IN PLAIN TERMS

Two numbers shape every resilience design. The **recovery time objective (RTO)** is how long a service may be down. The **recovery point objective (RPO)** is how much recent data you can afford to lose. Set them per service, then pay only for the protection each one needs.

Service type	Typical RTO	Typical RPO	Usual approach
Citizen payments, emergency dispatch	Minutes	Near zero	Active-active across two sites, synchronous replication
City command center, video analytics	Under 1 hour	Minutes	Warm standby site, frequent replication
Internal business systems	4 to 24 hours	1 to 4 hours	Cold or warm standby, scheduled backups
Archives and reports	Days	24 hours	Offsite backup with periodic restore tests

These are planning examples, not standards. Set values with the service owner and review them every year.

A SIMPLE INCIDENT DRILL

- 1. Detect.** Monitoring or a user report raises the alarm; the on-call lead confirms it is real.
- 2. Contain.** Isolate affected systems, preserve logs and stop the spread before fixing anything.
- 3. Report.** Notify management and, where the CERT-In directions apply, the 6-hour clock starts from detection.
- 4. Recover.** Restore from clean backups or standby systems and verify before returning to service.
- 5. Learn.** Hold a blameless review within a week and track each fix to closure.

QUESTIONS TO ASK EVERY VENDOR

Where is the source code kept, and can we escrow it? How are firmware and software updates signed and delivered? Which components come from outside India, and what is the alternative source? Where do logs and telemetry go? What is your response time, in hours, for a critical fault in our city?

PART D | CHAPTER 12

A practical roadmap

Start with decisions, not drawings. These ten steps and four gates keep a project honest.

THE FIRST 100 DAYS

- 1. Define the workload.** IT load by phase, density, availability and the customers or systems served.
- 2. Choose the model.** Build, lease, colocate or a mix, and the tier level.
- 3. Shortlist sites.** Use the scoring sheet in chapter 4 and visit the top three.
- 4. Start the power application.** Request capacity and compare open-access and utility tariffs.
- 5. Map compliance.** DPDP, CERT-In, sector rules, state policy and the approvals list in chapter 3.
- 6. Write the design brief.** Use the checklist in chapter 5, including PUE and WUE targets.
- 7. Set the Made in India criteria.** Agree the scorecard and weights before bids arrive.
- 8. Plan the digital layer.** Decide the sensors, data platform, ownership and AI use cases (chapter 8).
- 9. Build the budget.** Use the cost ranges in chapter 7 and add contingency of 8 to 12%.
- 10. Pick the team.** Designer, project manager, power and cooling specialists, and a commissioning agent from the start.

FOUR GATES

Gate	Evidence you need to pass
1. Concept	Workload defined, business case with power cost, tier chosen, shortlist of sites
2. Site and power	Land secured or option signed, power capacity letter, approvals plan, water source confirmed
3. Design freeze	Design brief signed, long-lead equipment ordered, scorecard applied, security and compliance review done
4. Handover	Integrated systems test passed, monitoring live, runbooks and training complete, baseline PUE measured

MISTAKES THAT COST THE MOST

- **Fixing the site before the power.** A beautiful plot with no grid capacity letter can sit idle for years.
- **Sizing for the final phase on day one.** Build in modules and add capacity as load arrives; stranded capacity is expensive.
- **Leaving the Made in India criteria until bids arrive.** Set the scorecard first, or it becomes a negotiation.
- **Treating monitoring as an afterthought.** Sensors, naming and data ownership are cheap to design in and costly to retrofit.
- **Skipping integrated testing.** Individual systems that pass alone can still fail together; test power, cooling and controls as one.
- **No plan for people.** Budget for training and a staffed operations team, not only for equipment.

QUESTIONS TO SETTLE BEFORE YOU START

Question	Why it matters
How many megawatts, and by when?	Drives site, power application and tier choice
Who owns the data and the keys?	Drives compliance, vendor choice and the digital layer
Which parts must be Made in India?	Drives the scorecard and procurement route
What is the power cost per unit?	Over five years, power can rival the build cost
Who runs it after handover?	Drives training, spares and service contracts

A CLOSING THOUGHT

India has the demand, the engineers and a growing industrial base. What it needs now are data centers designed and built with the same ambition as the services they carry. We hope this book helps you build one.

REFERENCE

Glossary

Plain-language definitions of the terms used in this book.

Term	Meaning
AI accelerator	A chip, usually a GPU, built for the heavy math of AI training and inference
Availability	The share of time a system works as intended, for example 99.982%
BMS / DCIM	Building management system and data center infrastructure management: software that monitors and controls power, cooling and space
CDU	Coolant distribution unit: moves heat between a rack's liquid loop and the facility loop
Colocation	Renting space, power and cooling for your own equipment in someone else's data center
CRAH	Computer room air handler: cools air using chilled water
CUE / WUE / PUE	Carbon, water and power usage effectiveness: carbon, water or total energy used per unit of IT energy
Digital twin	A live 3D model of a facility or city, kept current by sensor data
DPDP Act	Digital Personal Data Protection Act, 2023, India's data protection law
Edge data center	A small facility close to where data is produced
eSIM	An embedded SIM that lets a device connect to a mobile network without a physical card
Hyperscale	Very large facilities and campuses run by or for the largest cloud and internet companies
ICCC	Integrated command and control center for a city or facility
IoT	Internet of Things: connected sensors and devices
LPWAN	Low-power wide-area network for small devices that send little data over long distances
N+1 / 2N	Redundancy schemes: one spare unit, or a full duplicate system
PDU	Power distribution unit: delivers power to racks and equipment
PPP-MII	Public Procurement (Preference to Make in India) Order
Rack density	Electrical power drawn by one rack, in kilowatts
SLDC	State Load Despatch Centre, which manages the grid in a state
Sovereign AI	AI whose data, compute, models and operations are under national control
Tier (I to IV)	Uptime Institute levels of site infrastructure redundancy
UPS	Uninterruptible power supply: batteries that bridge outages until generators start
5D simulation	A 3D model plus time and cost, used to plan construction

REFERENCE

Sources

Where the figures in this book come from (accessed October 2026).

SOURCES USED FOR FIGURES IN THIS BOOK

- **Cushman & Wakefield.** India data centre market, H1 2026.
cushmanwakefield.com/en/india/insights/india-data-centre-update
- **JLL.** India data centre 2026 mid-year report. jll.com/en-in/insights/market-dynamics/india-data-centers
- **CEEW.** How is data centre infrastructure in India shaping power and water use?
ceew.in/publications/how-is-data-centre-infrastructure-in-india-shaping-power-and-water-use
- **IMARC Engineering.** Data Center Development in India: Costs and Approvals, 2026 guide.
imarcengineering.com/blog/data-center-development-in-india
- **Elets eGov.** India to add 20,000 GPUs to expand national AI compute capacity (19 February 2026).
egov.eletsonline.com
- **Digital Personal Data Protection Rules, 2025** (notified 14 November 2025).
en.wikipedia.org/wiki/Digital_Personal_Data_Protection_Rules,_2025
- **Vajiram & Ravi.** PLI 2.0 for IT hardware: revised scheme with ₹17,000 crore outlay; Storyboard18. Dell, Lenovo and HP among 27 companies approved under the IT hardware PLI scheme.

PRIMARY DOCUMENTS TO CONSULT BEFORE YOU RELY ON THIS BOOK

Public Procurement (Preference to Make in India) Order, 2017 and revisions; CERT-In Directions of 28 April 2022; Digital Personal Data Protection Act, 2023 and Rules, 2025; MeitY cloud empanelment guidelines; the Uptime Institute tier standards; National Building Code of India, 2016; your state's data center policy; the IndiaAI Mission cabinet approval of March 2024.

A NOTE ON SOURCES

Figures from different sources use different definitions and dates. Where sources disagree, both are shown. Items from general industry knowledge, such as tier availability figures, CERT-In directions and the structure of procurement orders, should be checked against the primary documents above.



ABOUT AJEEVI

HumanTech for a smarter world

AJEEVI Technologies builds AI, IoT and software for smarter urban infrastructure.
Made in India. Built for the world.

Our work spans smart city, smart building, smart water, smart mobility and smart energy solutions, AI and video analytics, IoT devices designed and built in India, mobile applications, and city digital twins and 5D simulation.

PLATFORMS

GyanAkshi	AI and video analytics
GyanYantram	IoT devices and sensors
GyanKendram	Integrated command center
GyanBhawan	Smart buildings and BMS
JalAJV	Water intelligence
VayuAJV	Environmental intelligence
TransAJV	Smart mobility
GyanSatark	Public safety
GyanRakshak	Citizen and operations helpdesk

TALK TO US ABOUT YOUR DATA CENTER

Website: www.ajeevi.com
Calculators: www.ajeevi.com/calculators.html
Contact: www.ajeevi.com/contact.html
Email: presales@ajeevi.com