AI Power-Play: India Data Center Capacity Could Surge to 19-23 GW by 2030

Power Ecosystem: A Hidden way to Play AI in India

Table of Contents

AI Power-Play: India DC Capacity Could Surge to 19-23 GW by 2030

POWER-PLAY IN AI

AI, cloud computing, and data localization are driving unprecedented demand for digital infrastructure in India:

- **19-23 GW** capacity required by 2030
- **~160 TWh** annual power demand by 2030 (~6% of total consumption and 8% by 2035)
- 11-12 diversified sub-sectors in DC ecosystem to invest & rotate capital
- ~230+ listed companies in the Power Ecosystem itself

Power Ecosystem: A Hidden way to Play AI in India

The typical narrative is that Indian stock market does not have any potential AI investments. AI models and applications run on the AI cloud infrastructure inside a data centre. With India generating nearly ~15-20% of the global data there is opportunity for developing Indian AI models and applications. The training and inferencing of these models would require significant data centre capacity. They are becoming the physical backbone of the AI economy.

"Such investments will strengthen India's position as a global hub for cloud computing and AI, while generating employment opportunities, supporting local supply chains and accelerating innovation-led growth. It is clear that the future of the world's digital economy is increasingly being shaped in India!" - **PM Narendra Modi**

This report estimates that **India requires 19-23 GW of Data Centre capacity by 2030**, implying one of the largest digital infrastructure buildouts in the country's history. This translates to a **power demand of ~160 TWh annually** (**~6% of India's projected electricity consumption in 2030**, may rise to **~8% by 2035**).

The investment opportunity spans across **11-12 infrastructure sub-sectors**. The power ecosystem forms the backbone of AI (~230 listed companies), while cooling, HVAC, networking, security and other facility infrastructure further broaden the investment universe. With ~**US$160-200 Bn** of investments already announced, we believe the power ecosystem is likely to be one of the largest beneficiaries of AI buildout.

"In the digital economy, data may be the new oil- but data centres are the refineries."

What is driving Data Centers Demand

The whole world is experiencing an unprecedented data generation boom.

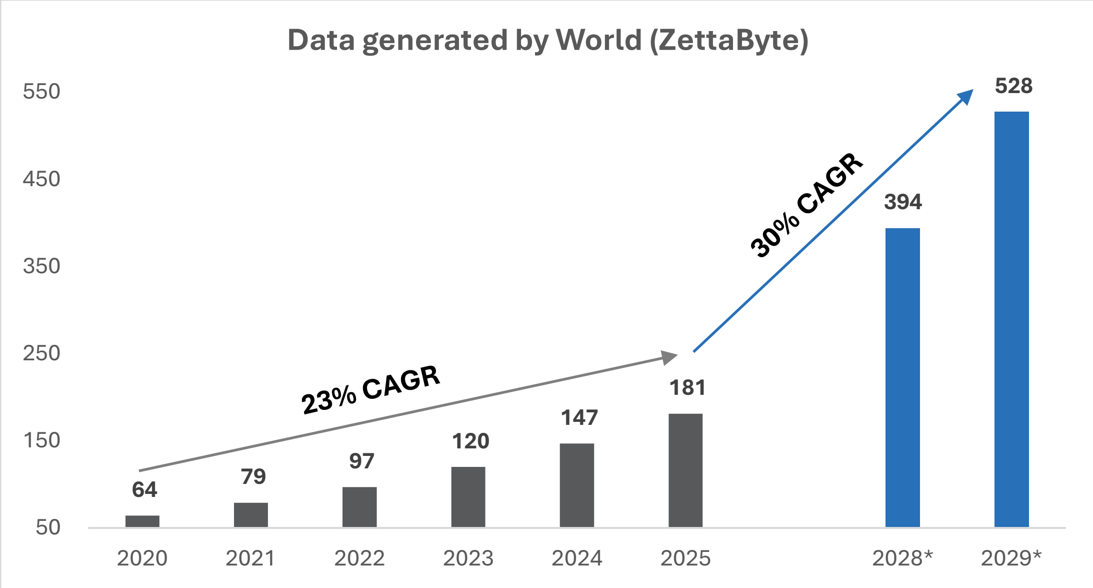

Exhibit 1: Data generation growth from 2020 to 2025 and forecast by IDC
Source: IDC, Red-gate

India also sits at the intersection of multiple structural megatrends including artificial intelligence, cloud computing, digital public infrastructure, digitization, and data localization requirements.

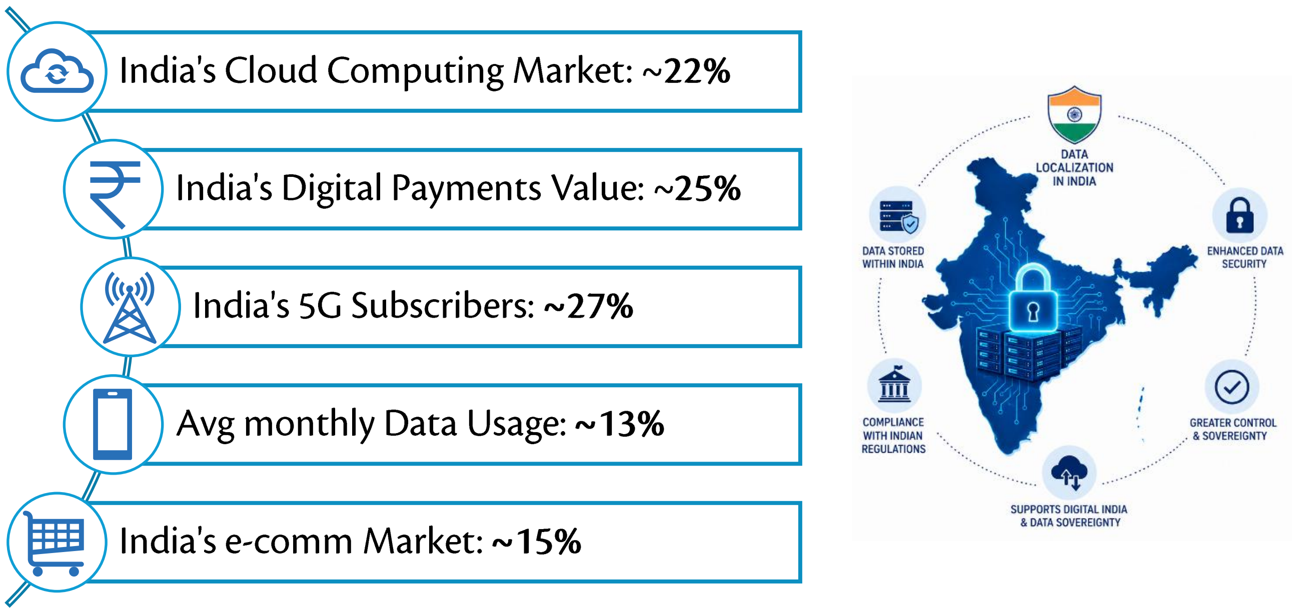

Exhibit 2: Projected CAGR (2025-2030) for the megatrends shaping DC demand, CEEW

AI boosts Everything

The AI revolution is transforming compute power into one of the most valuable resources of the 21st century. Supporting advanced AI systems requires vast networks of data centres operating continuously, equipped with high-performance processors, memory, storage, and power infrastructure. As AI workloads expand across industries, the demand for compute capacity is growing faster than ever, driving a global race to build the next generation of digital infrastructure.

For example:

Average power densities have more than doubled in just two years, to 17 kilowatts (kW) per rack, from 8 kW, and are expected to rise to as high as 30 kW by 2027 as AI workloads increase. Training models like ChatGPT can consume more than 80 kW per rack, while Nvidia's latest chip, the GB200, combined with its servers, may require rack densities of up to 120 kW. (McKinsey)

According to McKinsey report, by 2030, data centers are projected to require **$6.7 trillion** worldwide to keep pace with the demand for compute power.

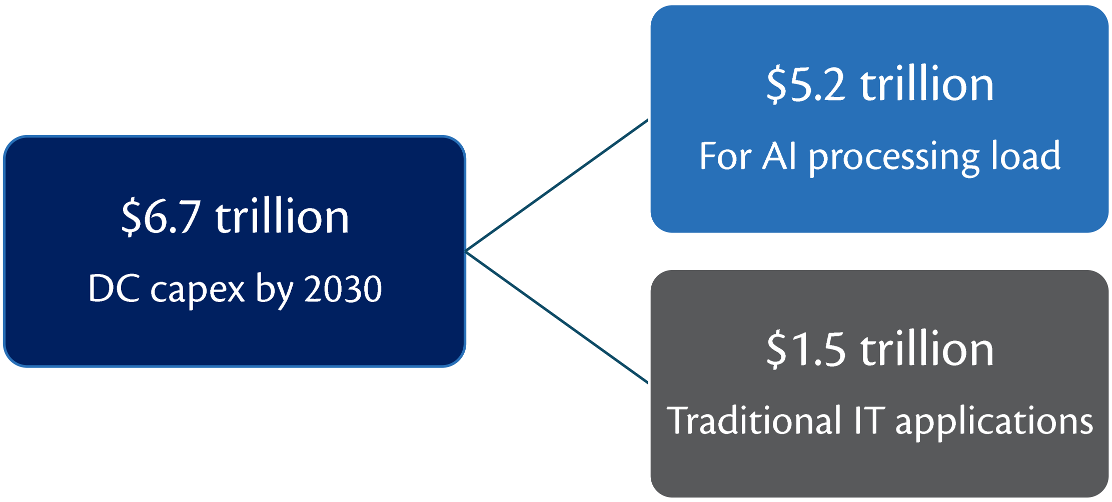

Exhibit 3: Capital investments needed for Data Center expansion for AI processing and traditional IT applications, McKinsey

The Five Layers of AI

When most people think about AI, they think about large language models and chatbots. However, AI is actually built on a multi-layered technology stack. As NVIDIA CEO Jensen Huang describes it (Video), AI is a five-layer system that starts with energy and computing infrastructure, progresses through cloud platforms and AI models, and ultimately delivers value through applications. Understanding this stack is critical because every layer depends on the layers below it, driving one of the largest infrastructure buildouts in human history.

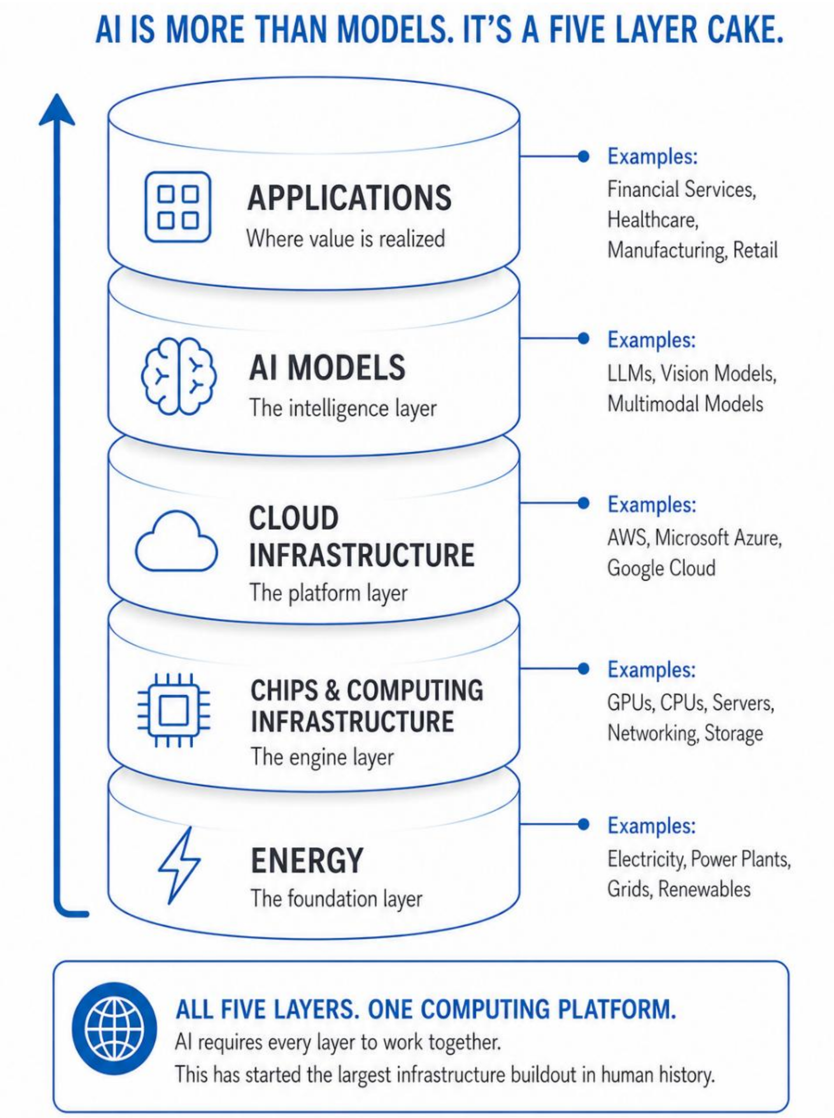

Exhibit 4: The AI Stack- More Than Just Models (Jensen Huang)

"AI, because it's processed in real time and it generates intelligence in real time; it needs energy to do so. Energy is the first layer." - Jensen Huang

Beyond GPUs: The Physical Infrastructure Powering AI

AI is often associated with GPUs, models, and software. However, behind every AI workload lies a vast ecosystem of physical infrastructure that delivers power, cooling, connectivity, security, and operational resilience.

From transformers and backup generators to chilled water systems, cable networks, and facility controls, these foundational systems ensure that AI data centers can operate reliably at scale. As AI adoption accelerates, these non-compute components are becoming strategic enablers of the entire AI economy.

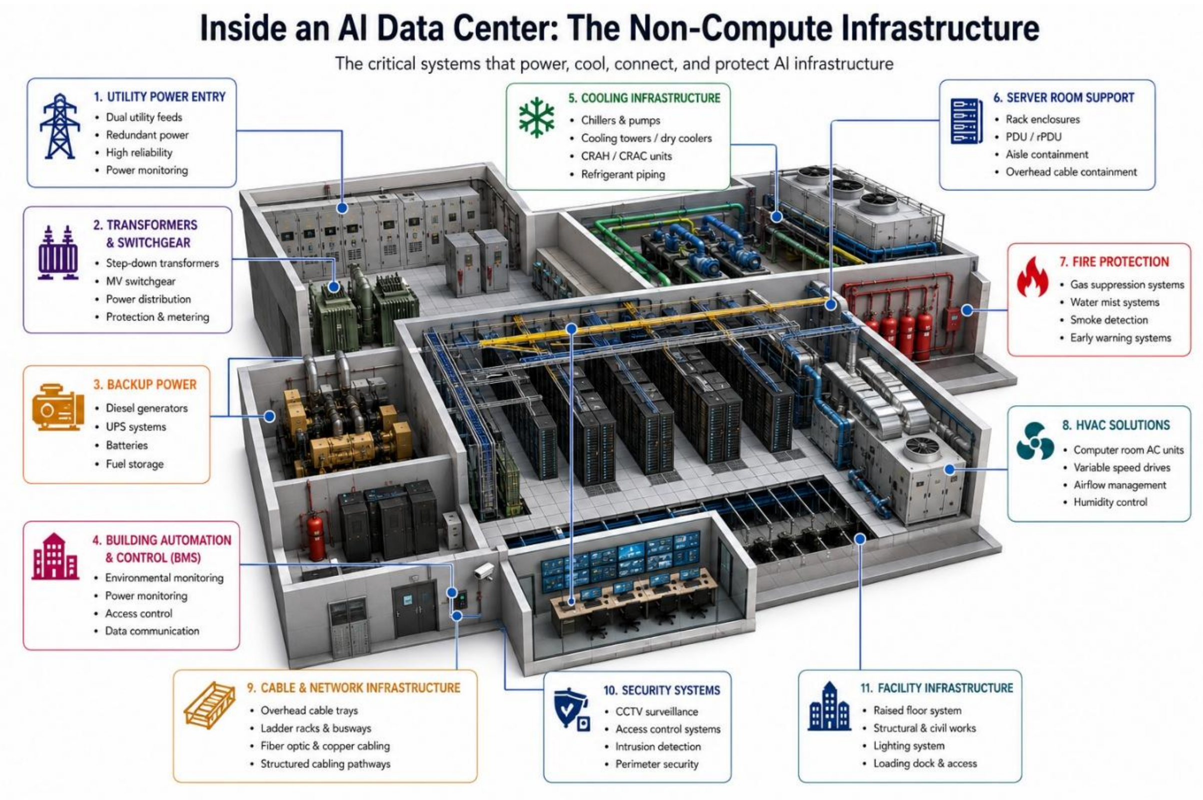

Exhibit 5: Non-compute infrastructure- including power, cooling, networking, security, and facility systems- forms the physical backbone of modern AI data centers.

A Three-Step Framework for Forecasting Future Infrastructure Demand

We estimate future data centre capacity using a three-step framework, with each component driven by a distinct underlying demand driver rather than a single top-down assumption.

First, we estimate Traditional Data Centre demand by analysing historical improvements in data centre efficiency (ZB/GW) and projecting future infrastructure requirements based on global data generation. **Second,** we estimate AI Training demand by assessing the expected growth in frontier AI models, the expansion of frontier AI developers and improvements in training efficiency. **Finally,** we estimate AI Inference demand using technology adoption principles, supported by historical digital adoption cycles and the rapid growth in AI token consumption.

Each of these components is first estimated at the global level using independent demand drivers. We then translate the global estimates into India's expected requirement using India's share of global digital data generation, together with adjustments for India's likely adoption of mature hardware and continued improvements in AI efficiency.

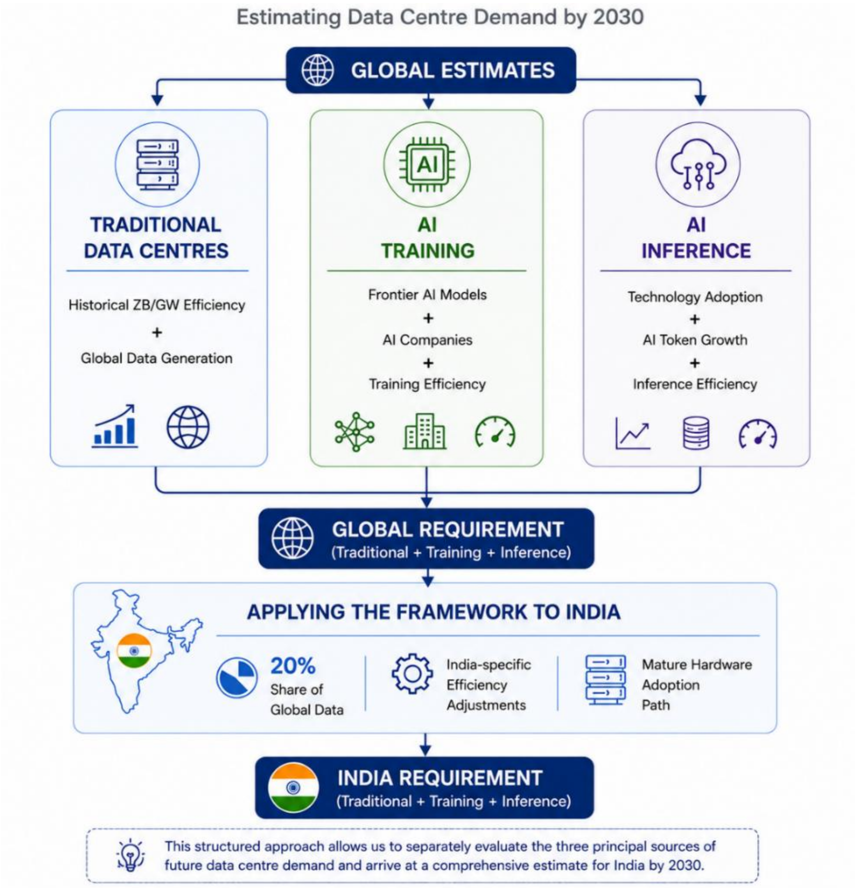

Exhibit 6: Framework used to estimate Global & India Data Centre Capacity Requirement by 2030

Q1. How much conventional data centre capacity will the world require by 2030?

Continuous Improvement in Traditional Data Centres Efficiency

Historically, the amount of digital data that could be processed per GW of installed data centre capacity has improved consistently due to advances in semiconductors, storage technologies, networking equipment, virtualization and cooling systems. This relationship can be measured using the ratio of **annual data generated (ZB) per GW of installed data centre capacity (ZB/GW)**, which serves as a useful proxy for infrastructure efficiency.

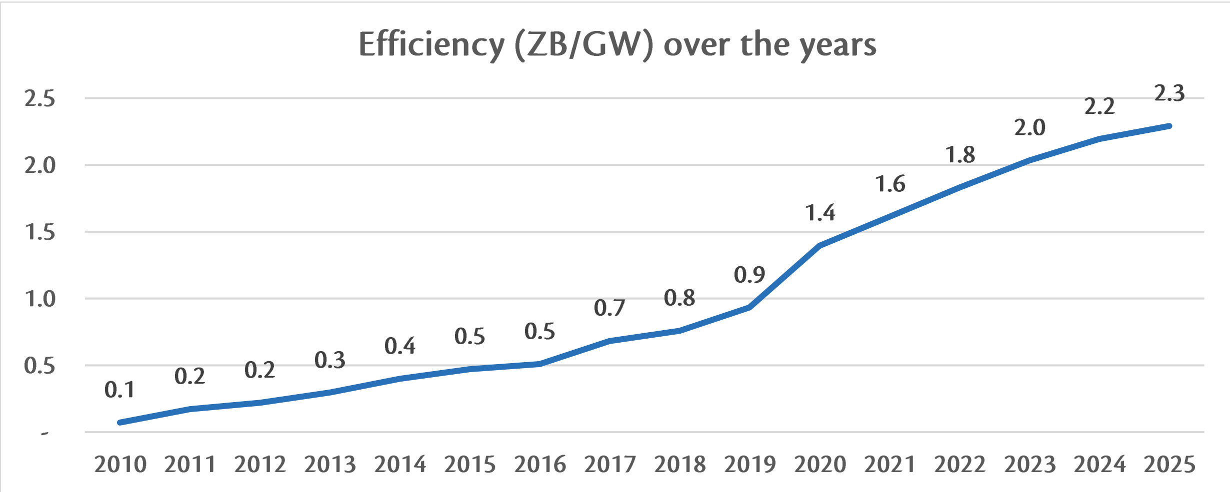

Exhibit 7: The global ratio of annual data generation per GW of installed data centre capacity
Data Sources: Data Generation, DC Capacity

Between **2010 and 2020**, before large-scale AI infrastructure emerged, global data centre efficiency improved rapidly. The ZB/GW ratio increased from just **0.07** in 2010 to **1.40** in 2020, implying an efficiency improvement of approximately **25% CAGR** over the 3, 5 and 7-year period.

AI Masks Continued Improvements in Traditional Data Centre Efficiency post 2020

The divergence after 2020 is particularly important. Had historical efficiency trends continued uninterrupted, the global data centre industry would have been capable of processing approximately **4.25 ZB per GW** by 2025. Instead, the observed ratio stands at only **2.3 ZB/GW**.

This apparent slowdown does not indicate declining technological efficiency. Rather, it reflects the emergence of AI workloads, which are substantially more compute-intensive than conventional cloud, enterprise and consumer applications. AI training clusters require significantly higher compute density, resulting in a lower aggregate ZB/GW ratio despite continued improvements in underlying hardware efficiency.

To validate this observation, we estimate the composition of global installed capacity in 2025. The United States accounts for approximately **40 GW** of installed capacity, with nearly **70% AI-dedicated**, while China contributes around **24 GW** with approximately **40% AI-dedicated** capacity. The rest of the world contributes approximately **15–16 GW**, of which only **~5%** is AI-oriented. This implies that of the world's **~80 GW** installed capacity, roughly **40 GW** continues to serve traditional digital workloads.

Applying global data generation of **181 ZB** to this estimated **40 GW** of traditional capacity yields an implied efficiency of approximately **4.45 ZB/GW**, remarkably close to the **4.25 ZB/GW** projected using historical pre-AI efficiency trends. This provides independent validation that the underlying efficiency improvements have continued but are masked by the rapid expansion of AI infrastructure.

Rapid advances in GPU architectures, networking, software optimization and low-precision training continue to improve training efficiency. Recent MLPerf benchmarks show NVIDIA Blackwell delivering **2× higher training performance per dollar** and **over 2-3× faster training performance** than the previous Hopper generation. Continued improvements in hardware, memory systems, distributed training techniques and algorithmic optimization are expected throughout the remainder of the decade.

***NVIDIA Blackwell Enables 3x Faster Training and Nearly 2x Training Performance Per Dollar than Previous-Gen Architecture** (NVIDIA)*

The World Could Require 60-75 GW of Traditional Data Centre Capacity by 2030

While AI infrastructure will continue to expand separately, traditional workloads should benefit from the technological progress that has characterized the industry over the past decade.

Assuming a conservative efficiency improvement of 15-20% annually over the next five years, traditional data centre **efficiency could improve to approximately 9-11 ZB/GW by 2030.**

At the same time, IDC projects global data generation to exceed 660 ZB annually by 2030. Applying the projected efficiency range implies that the world would require approximately 60-75 GW of traditional data centre capacity to support conventional digital workloads alone.

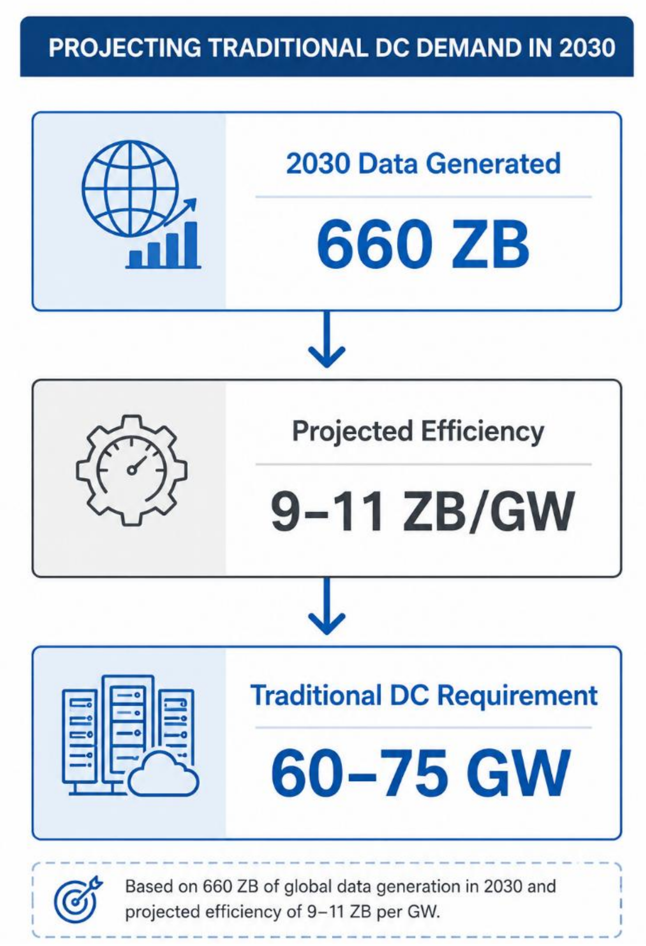

Exhibit 8: Global traditional DC capacity requirement by 2030

Q2. How Much Infrastructure Will the World Need to Train Future AI Models?

Depends on How Fast Will Frontier AI Development Expand by 2030?

Unlike traditional digital workloads, AI training demand is driven not by data generation, but by the number of frontier models being developed. Consequently, estimating AI training infrastructure requires projecting how the global ecosystem of frontier AI developers is likely to evolve over the remainder of the decade.

At present, we estimate that there are approximately 18-22 organizations globally developing frontier-scale foundation models, including OpenAI, Anthropic, Google DeepMind, Meta, xAI, Amazon, Microsoft, Alibaba, Baidu, Tencent and a handful of specialized AI laboratories.

A recent Epoch AI study estimates that by **2030**, approximately **239 frontier AI models** will exceed the **10^{26} FLOP** training threshold used within the U.S. AI Diffusion Framework. Under more aggressive scenarios, this number could exceed **500 models**. (EPOCH)

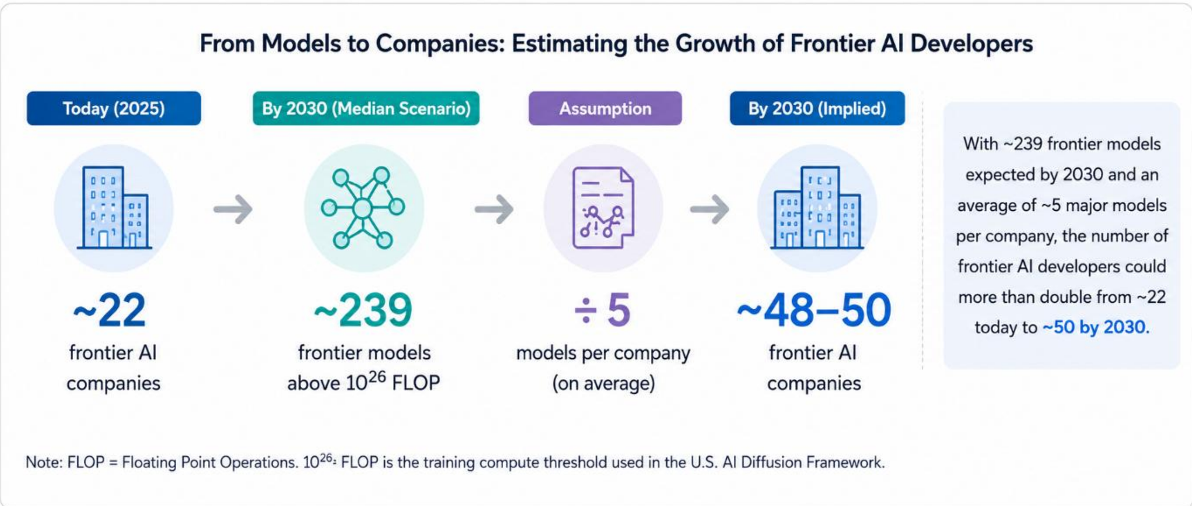

Exhibit 9: Estimated growth in number of AI models and companies by 2030

The number of companies needs to be doubled by 2030 along with efficiency improvements, hence

We Estimate Global AI Training Capacity Will Double to ~50 GW by 2030

As the number of frontier AI developers may increase by roughly **2-3×**, we estimate **global AI training capacity to increase from approximately 24 GW today to nearly 50 GW by 2030.**

Having estimated the infrastructure required to develop AI models, we next estimate the much larger opportunity arising from AI inference as AI adoption scales globally.

Q3. How Much Infrastructure Will the World Need to Run AI at Scale?

Depends on How Fast Could AI Adoption Scale by 2030?

Unlike AI training, which is concentrated among a relatively small number of frontier AI developers, inference is driven by widespread adoption across consumers, enterprises and governments. Every AI interaction- from search and coding assistants to AI agents and enterprise copilots- requires inference compute. **Inference demand is primarily a function of technology adoption rather than model development.**

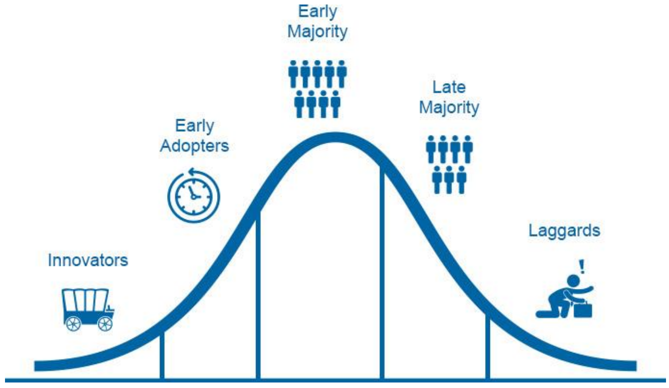

Exhibit 10:
Technology Adoption Lifecycle

So, Where Are We in the AI Adoption Lifecycle Today?

We believe AI today is still largely in the **Innovators** phase. Current adoption is primarily concentrated among technology companies, early enterprise adopters, developers and digitally sophisticated users. However, as models become more capable, inference costs decline and AI is increasingly embedded into everyday applications, adoption is expected to expand rapidly into the much larger **Early Adopters** and **Early Majority** segments. Historically, this transition has marked the period of fastest infrastructure growth for transformative technologies.

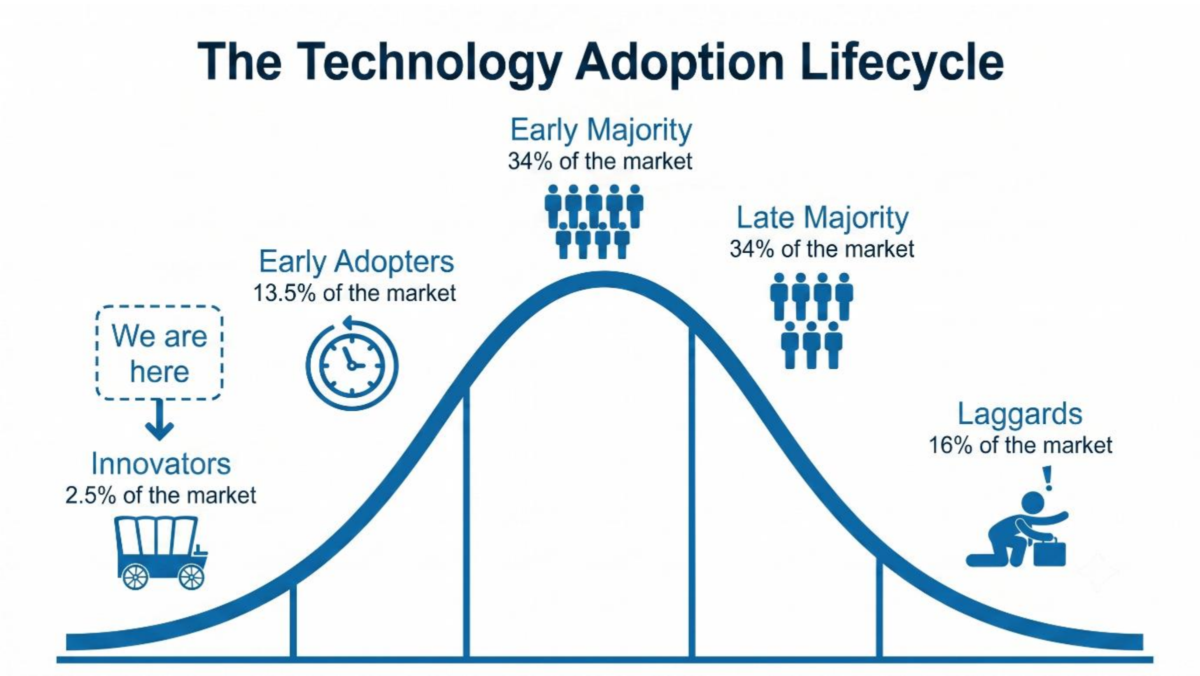

Exhibit 11: Inferencing accelerates as AI transitions from experimentation to mass adoption

According to **Goldman Sachs Research**, AI agents could drive approximately a **24-fold increase in token consumption by 2030** as enterprises and consumers increasingly automate complex workflows. This represents one of the strongest indicators of future inference demand, as every generated token requires inference compute. (Goldman)

Similarly, Google recently disclosed that its products now process **more than 3.2 quadrillion tokens every month**, representing approximately **7× year-over-year growth**, highlighting the unprecedented acceleration already underway.

History also Suggests Technologies Scale Exponentially

History provides several useful analogues for how rapidly digital infrastructure demand can scale once technologies reach mass adoption.

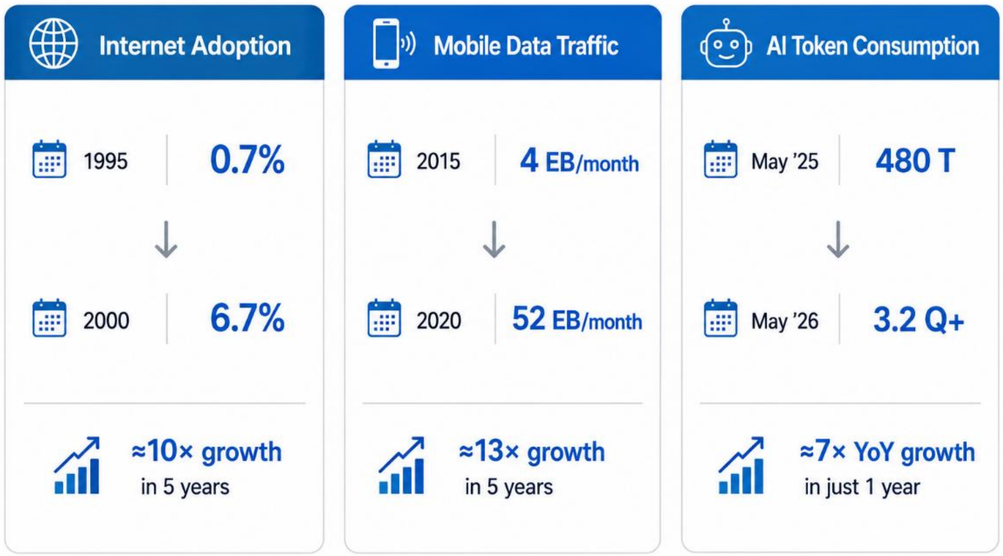

Exhibit 12: Historical growth rates during Technology Adoption Lifecycle at Early Majority stage

So, Global AI Inference Capacity Could Increase Nearly Eightfold to ~120 GW by 2030

Although token demand may increase by well over an order of magnitude during the remainder of the decade, improvements in GPUs, inference chips, model architectures, quantization techniques and software optimization will continue reducing the compute required per token.

Accordingly, rather than assuming inference infrastructure scales directly with token growth, we apply a more conservative **50% annual growth rate** to global inference capacity through 2030.

This implies global AI inference infrastructure expanding from approximately 16 GW today to nearly 120 GW by 2030.

Total Global Data Centre Capacity Could Triple from 80 GW to ~230-245 GW by 2030

Bringing together our three independent demand models suggests that global data centre capacity could increase from approximately **80 GW today to nearly 230-245 GW by 2030**. While traditional digital infrastructure will continue to underpin enterprise and cloud computing, AI is expected to become the dominant driver of incremental capacity additions. In particular, AI inference is likely to account for the largest share of future demand as AI adoption expands beyond innovators into mainstream consumer and enterprise use cases.

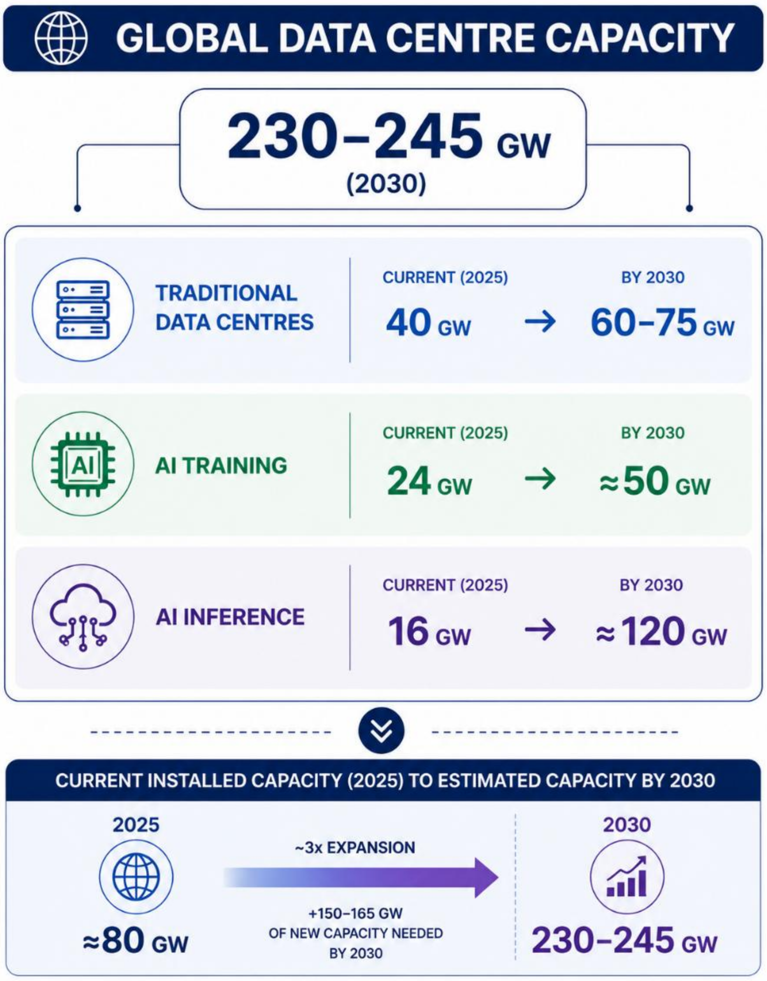

Exhibit 13: Estimated total global Data Centre capacity requirement by 2030

How Much Data Centre Capacity Could India Require by 2030?

India: 15%-20% Global Data & ~2.0% DC Capacity

India has emerged as one of the world's largest generators of digital data. According to industry estimates (Sources: Velloni, Press Insider), India accounted for approximately **15-20% of global data generation** in 2025, equivalent to roughly 27–36 zettabytes (ZB) of data annually. Despite this, India's installed data center capacity stands at only ~1.6 GW, representing merely **~2.0% of the global installed capacity of ~80 GW.** India is currently a **"Digital Giant, but Infrastructure Lightweight"** that requires much more operational data centre capacity as against installed capacity of only ~1.6 GW, highlighting a substantial infrastructure gap.

*The infrastructure for this amount of data will eventually be required as India moves towards greater **data localization, AI adoption, cloud penetration, and domestic digital sovereignty.***

we translate our global estimates using India's growing share of worldwide digital data generation while incorporating country-specific considerations around AI adoption, hardware deployment and infrastructure efficiency.

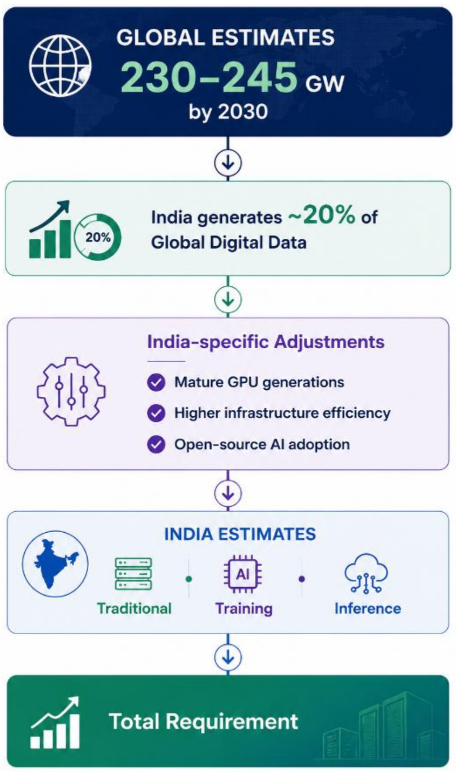

Exhibit 14: Translating global estimates to India proportional to data generation and adjustments

India will require 19-23 GW total DC capacity by 2030

India currently generates approximately 20% of the world's digital data. As a starting point, we allocate India's infrastructure requirement in proportion to its share of global data generation.

For traditional data centres, this proportional allocation is directly applicable because conventional digital infrastructure primarily scales with enterprise computing, cloud services and digital data generation. Accordingly, we estimate India's **traditional data centre requirement at approximately 12–15 GW by 2030.**

However, AI infrastructure is expected to follow a different deployment path. Domestic AI developers are expected to rely on **increasingly efficient open-source models, mature GPU generations and optimized training and inference techniques.** Consequently, the proportional allocation derived from global estimates is adjusted downward to reflect India's likely technology adoption profile.

Applying these efficiency adjustments results in an estimated **2-3 GW of AI training capacity and 5-6 GW of AI inference capacity by 2030.**

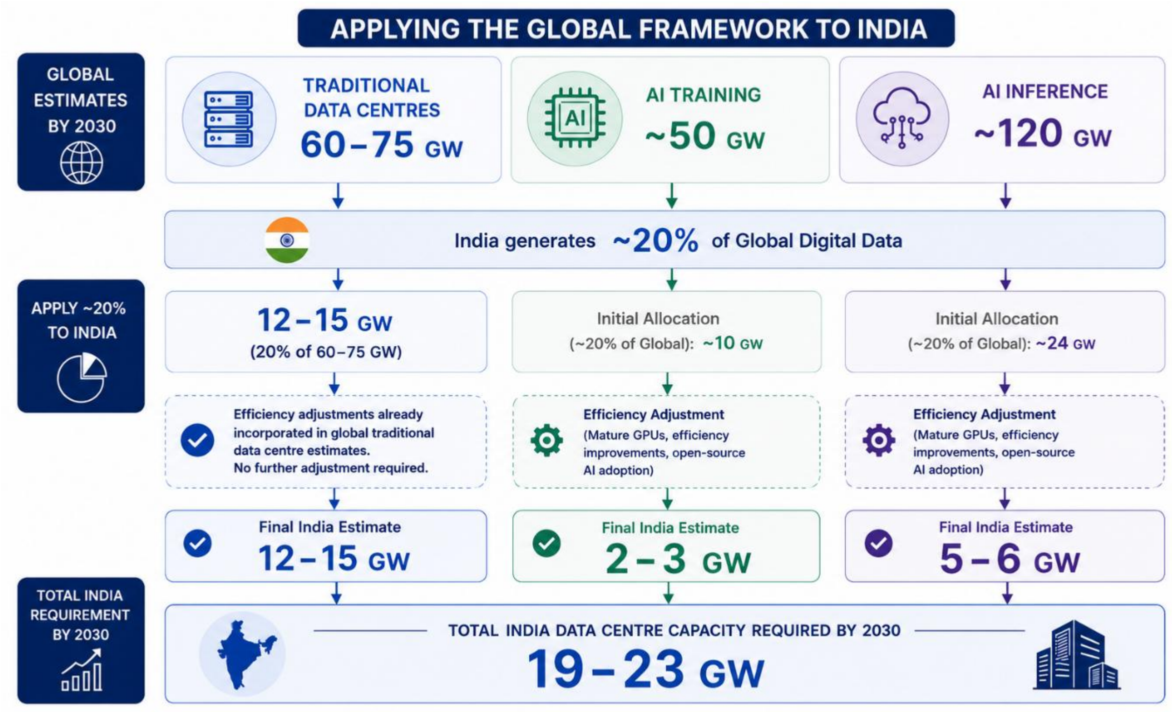

Exhibit 15: Estimated India total DC capacity requirement by 2030

Combining these three components, we estimate that **India could require approximately 19-23 GW of operational data centre capacity by 2030**, forming the foundation of the country's digital and AI infrastructure over the remainder of the decade.

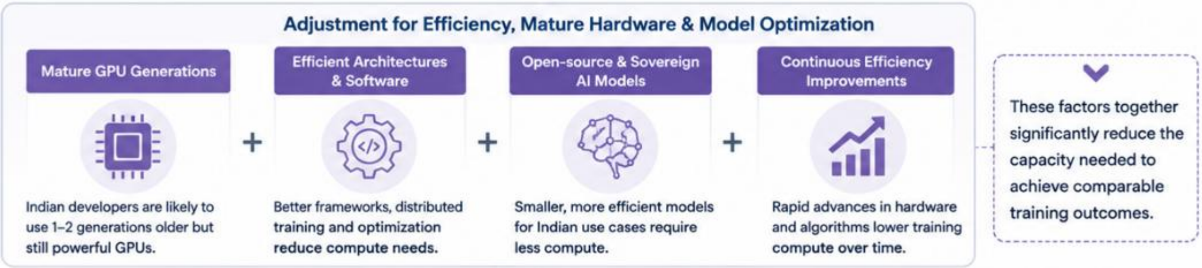

Exhibit 16: Factors Driving Lower AI Infrastructure Requirements in India

AI's Physical Layer Comes with nearly Tn-$ Price Tag

The AI era is rapidly redefining infrastructure requirements. While conventional data centers primarily require investment in power, cooling and physical infrastructure, **AI-dedicated facilities are significantly more capital intensive due to the high cost of GPUs, servers, storage and networking equipment,** which can account for nearly three-fourths of total project costs. As a result, AI data centers can require 6-7x higher capital per MW compared to traditional facilities.

Traditional Data Centers	Cost per MW stated
CBRE (Link)	~Rs 67 Cr per MW
SIFY DRHP (Link)	~Rs 40 Cr per MW
Avendus (Link)	~Rs 79 Cr per MW
PL Capital (Link)	~Rs 60-70 Cr per MW
Average	**~$6-7 Bn per GW**

AI-Dedicated Data Centers	Cost per MW stated
Alpha-Matica (Link)	~$45 Bn per GW
JLL (Link)	~$40-50 Bn per GW
McKinsey (Link)	~$42 Bn per GW
IBM CEO Arvind Krishna (Link)	~$80 Bn per GW
Average	**~$45 Bn per GW**

Exhibit 17: AI Data Centers are Significantly More Capital Intensive than Traditional Data Centers

Building India's AI Infrastructure Could Require $350-435 Billion by 2030

we estimate an average capital intensity of approximately US$6.5 billion per GW for traditional data centres and US$40 billion per GW for AI-dedicated facilities.

Applying these benchmarks to our estimated capacity requirements suggests that **India could require US$350-435 billion of cumulative investment by 2030.** Traditional data centres account for approximately US$77-96 billion, while AI infrastructure- including both training and inference clusters- represents the remaining US$270-340 billion, highlighting how AI significantly increases the capital intensity of future data centre investments.

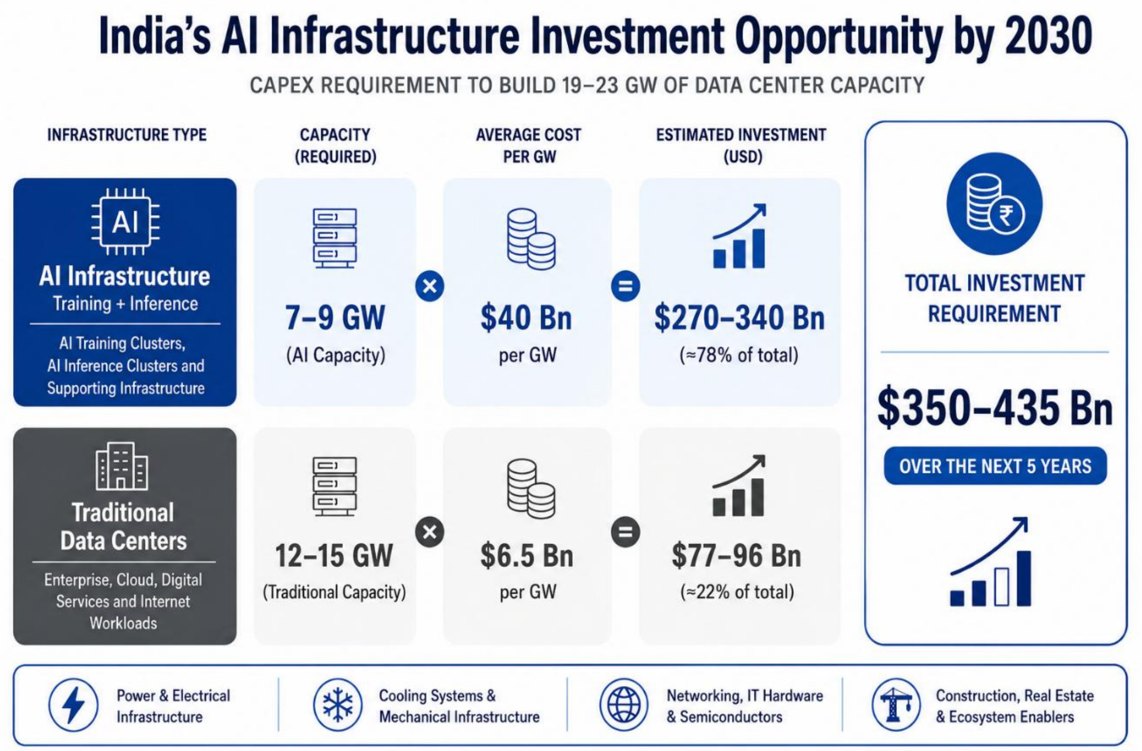

Exhibit 18: India's Data Center Infrastructure Capex opportunity by 2030

The opportunity extends well beyond data centre developers. It encompasses power generation, transmission infrastructure, transformers, switchgear, electrical equipment, cooling systems, construction, networking equipment, semiconductors and the broader digital infrastructure ecosystem.

Investment Momentum Underway: $200 Bn Pipeline Could Be the First Leg of a $300 Bn+ Opportunity

Investment announcements over the last few years have accelerated dramatically.

- Investments of nearly **USD 70 billion** are already underway in India's data centre sector, with an additional **USD 90 billion** in announced projects, totalling **~$160 Bn** and highlighting the scale of expansion. (PIB) - Feb 2026

- Over the past few years, total committed investments have reached approximately USD 95 billion and are expected to exceed **USD 100 billion by 2027.** (CEEW) - Feb 2026

- Some estimates place the broader investment opportunity at **US $200 billion** by 2030. (Fortune) - Feb 2026

This makes data centres one of the largest emerging infrastructure investment themes in India.

Policy Tailwinds and Bottlenecks

India's data centre industry sits at an inflection point. On one hand, a combination of supportive government policies, AI initiatives, infrastructure incentives, and state-level programs is creating a highly attractive investment environment. On the other, the industry faces significant challenges related to power availability, transmission infrastructure, land acquisition, water availability, and project execution.

Policy Tailwinds: Building the Foundation for Growth

- **Infrastructure Status:** to data centres with more than 5 MW capacity giving them benefits of access to long-tenure financing at lower cost of capital.
- **Tax holiday:** Till 2047 for eligible foreign cloud service providers operating through India-based data centre infrastructure.
- **State-wise policies:** State policies are a major driver of data centre investments and capacity build-out in India. An analysis of **15 state data centre policies across 39 policy features** reveals the most widely adopted incentives and support measures offered by state governments.

Policy Feature	Related to	Of 15 States
Single-window clearance	Approvals, govern & regulatory facilitation	14
Electricity duty exemption	Power-related Incentives	12
Concessional land allotment	Land, zoning, and infrastructure	12
Stamp duty exemption	Land, zoning, and infrastructure	12
Dedicated data centre parks	Land, zoning, and infrastructure	12

Assured 24x7 infrastructure provisioning	Land, zoning, and infrastructure	12
Assured 24x7 water supply	Water-use governance and cooling strategies	11
Dedicated feeders/substations	Power-related Incentives	10
FAR/zoning relaxations	Land, zoning, and infrastructure	10
Capital subsidy/FCI subsidy	Fiscal and tax Incentives	10
Essential services classification	Approvals, govern & regulatory facilitation	10

State policies have largely converged around a common set of incentives aimed at accelerating data centre investments. The most widely adopted measures include single-window clearances, electricity duty exemptions, concessional land allotment, stamp duty waivers, dedicated data centre parks, and assured infrastructure provisioning.

Additionally, states are increasingly adopting park-based development models that offer plug-and-play infrastructure and faster deployment timelines.

A notable trend is the **growing emphasis on power-related incentives**, reflecting the industry's most critical operating requirement.

Policy Feature	Related to	Of 15 States
Captive power facilitation	Power-related Incentives	8
Industrial tariff classification	Power-related Incentives	7
Renewable-linked power incentives	Power-related Incentives	7
EPF reimbursement	Fiscal and tax Incentives	7
Anchor unit incentive	Fiscal and tax Incentives	7

The Missing Gigawatts: Potential Bottlenecks

Despite strong investment-focused policies, several gaps persist. Most states **lack robust provisions around water management, sustainability standards, and climate resilience.** Only a handful of states have introduced measures such as water recycling, rainwater harvesting, green building standards, or advanced cooling frameworks. Similarly, performance-based metrics such as PUE, WUE, and carbon intensity targets remain largely absent.

Policy Feature	Related to	Of 15 States
Training subsidy	Fiscal and tax Incentives	6
Fast-track mega project approvals	Approvals, governance and regulatory facilitation	6
Transmission & wheeling charge waiver	Power-related Incentives	5
SGST reimbursement	Fiscal and tax Incentives	5
Interest subsidy	Fiscal and tax Incentives	5

Energy efficiency framing	Sustainability and green data centre promotion	5
Registration fee reimbursement	Land, zoning, and infrastructure	4
Infrastructure sector status	Approvals, governance and regulatory facilitation	4
Green/clean tech subsidy	Sustainability and green data centre promotion	4
Land diversion fee exemption	Land, zoning, and infrastructure	3
Rental subsidy	Fiscal and tax Incentives	3
Deemed distribution/franchisee model	Approvals, governance and regulatory facilitation	3
Green building standards	Sustainability and green data centre promotion	2
Water recycling	Water-use governance, cooling strategies	1
ZLD requirement	Water-use governance, cooling strategies	1
Rainwater harvesting & recharge	Water-use governance, cooling strategies	1
Smart water management systems	Water-use governance, cooling strategies	1
Seawater cooling promotion	Water-use governance, cooling strategies	1
Dedicated sustainability architect	Sustainability & green DC promotion	1

- **Transmission Infrastructure, Power Grid challenge:** According to World Economic Forum (WEF), Investment in AI data centres is growing faster than power grids can keep up, making grid connectivity a constraint.

  In many regions, connecting a new facility to the power grid can take 4-10 years, while AI data centres are typically planned and built within two to three. This misalignment increasingly determines which projects advance and which stall.

India has largely solved the challenge of attracting data centre investments through supportive policies and incentives. The next phase of growth, however, will depend on its ability to strengthen power infrastructure, transmission networks, water management, and sustainability frameworks.

The Power Ecosystem Driving AI

The growth of data centres is not merely a infrastructure story- it is increasingly becoming an energy story. Unlike conventional commercial buildings, data centres operate 24 hours a day, 365 days a year, requiring uninterrupted power for servers, networking equipment, cooling systems, and backup infrastructure. As AI workloads become more prevalent, power intensity is expected to increase further due to the deployment of high-density GPU clusters.

In India, data centre occupancy levels were about 75-80% during 2023 (Source). So, taking 21 GW as operational capacity in 2030 and 80%-90% utilisation gives us electricity consumption of about **~150-170 TWh** annually in 2030.

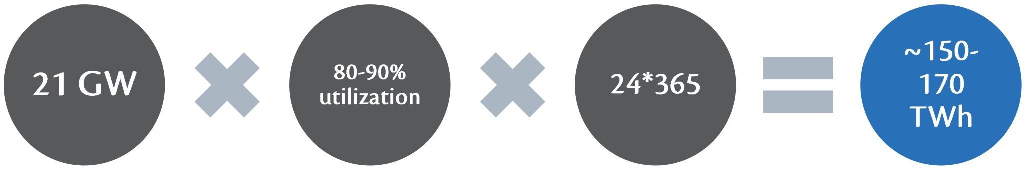

Exhibit 19: Expected electricity consumption by Indian Data Centers annually in 2030

> This translates to roughly **~6% of India's projected electricity demand in 2030** (total ~2,600 TWh) and could rise to **~8% by 2035** (total ~4,050 TWh). For context, the estimated data centre power requirement in 2030 is already equivalent to **~9% of India's 2025 electricity consumption** (1,707 TWh), illustrating the scale of incremental power demand that digital infrastructure could generate over the coming decade.

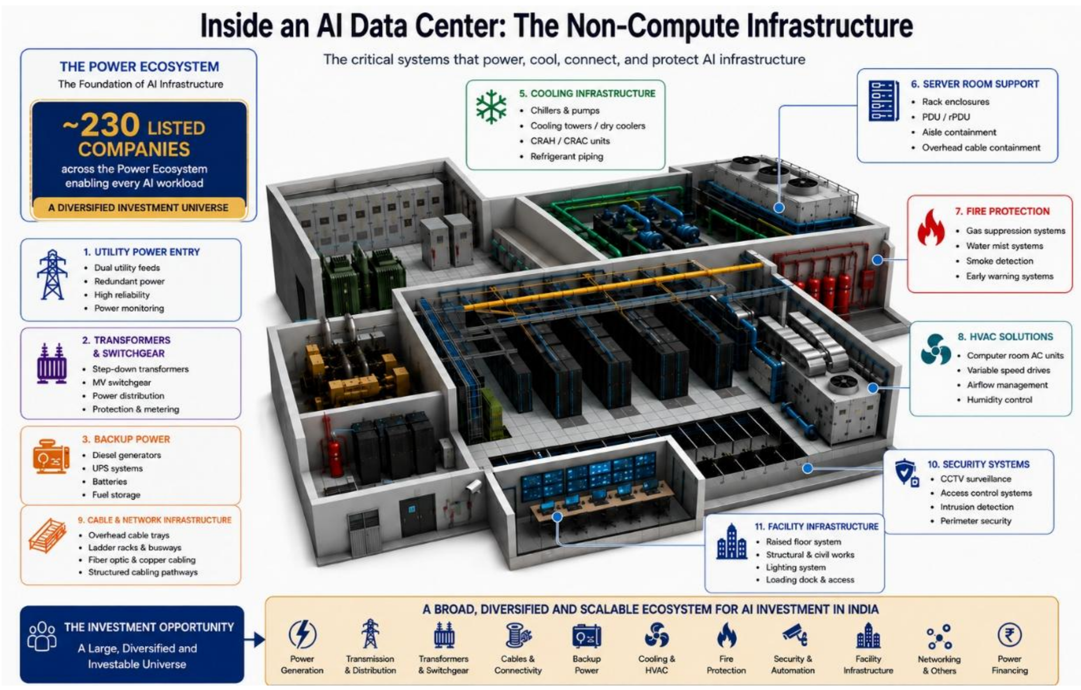

Supporting India's estimated 19-23 GW data Center requirement will necessitate substantial investments across the broader power ecosystem, extending well beyond generation capacity. From our earlier report on Power (Report), we estimated a Rs 65-70 trillion investment opportunity by 2035, with nearly **Rs 54 trillion directed towards capacity additions** and another **Rs 13 trillion towards transmission infrastructure.**

Importantly, this opportunity is highly diversified, spanning nearly **230 listed companies** across utilities, power generation, transmission, transformers, switchgear, electrical equipment, industrial cables, backup power and related infrastructure. Beyond the power ecosystem, the broader AI data centre value chain also includes cooling systems, HVAC solutions, fire protection, security systems, networking, building automation and facility infrastructure, creating an even wider universe of investment opportunities. Overall, the AI infrastructure ecosystem spans **11-12 distinct sub-sectors**, offering investors a broad platform to allocate and rotate capital as different segments emerge at varying stages of the infrastructure buildout.

While the opportunity spans the entire AI infrastructure ecosystem, we believe the power ecosystem currently offers one of the most compelling risk-reward opportunities, given its critical role in enabling AI infrastructure and the scale of investments likely to be required over the coming decade.

Risks to the Thesis

While the AI data centre is a high growth opportunity it is exposed to the economical, technological and policy uncertainties. The ultimate pace of expansion will depend on India's ability to overcome critical bottlenecks related to power availability, transmission infrastructure, water management, and execution capability among other factors.

Conclusion: India's Next Infrastructure Megatrend

India's digital economy is entering a compute-intensive era driven by AI, cloud adoption, digital public infrastructure, and data localization. While current installed data centre capacity stands at just ~1.6 GW, the country's long-term requirements could potentially reach **19-23 GW by 2030**, implying one of the largest digital infrastructure buildouts globally, which is estimated to consume **~160 TWh annually by 2030 (~6% of total consumption, ~8% by 2035).**

While data centres represent the visible face of AI infrastructure, the broader power ecosystem- including transmission, transformers, switchgear, backup power, cables, cooling and electrical equipment- may ultimately become one of the largest and most enduring beneficiaries. As of Feb 2026, the announced investment pipeline is already at **$160-200 Bn**. It is clear from the strong policy support that the government is visualising India as not only having its own AI Data Centre infrastructure but also as a global hub for international AI infrastructure.

Key References:

https://www.ceew.in/sites/default/files/ceew-data-centre-study-web-ready-final.pdf

https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers#/

https://fortune.com/2026/02/17/india-cobbles-together-200-billion-plus-for-data-center-investment/

https://www.pib.gov.in/PressReleasePage.aspx?PRID=2227953&reg=3&lang=2

https://pressinsider.com/news/india-generates-20-of-global-data-but-hosts-just-3-of-data-centres/

https://velloni.com/database-management-services-in-india/

https://www.cushmanwakefield.com/en/insights/americas-data-center-update

https://www.morganstanley.com/content/dam/msdotcom/en/indiaofferdocuments/pdfs/Project_Green_DRHP.pdf

https://www.bain.com/insights/ai-data-center-forecast-from-scramble-to-strategy-snap-chart/#

https://explodingtopics.com/blog/data-generated-per-day

https://www.visualcapitalist.com/charted-the-growth-of-global-data-center-capacity-2005-2025/

Disclaimers

This document is released by OmniScience Insights Labs and it is independent of any registration-related activities of Omniscience Capital Advisors Private Limited. It does not constitute a "research report" under SEBI (Research Analyst) Regulations, 2014, or "investment advice" under SEBI (Investment Adviser) Regulations, 2013. It is also not a solicitation or offer under SEBI (Portfolio Managers) Regulations, 2020, which govern discretionary, non-discretionary, and advisory PMS services. The content herein is intended solely for educational and knowledge-building purposes, designed to provoke thought and share insights.

Omniscience Investment Adviser is the advisory division of Omniscience Capital Advisors Pvt. Ltd. which is registered as a non-individual Investment Adviser with SEBI with a valid registration No. INA000007623, BSE Enlistment No.: 1197 and CIN U93000MH2017PTC290053. Principal Officer is Vikas V Gupta (Contact No. 9987681967, Email: vikas.gupta@omnisciencecapital.com) and Compliance Officer is Chanchal Manglunia (Contact No. 9320816319, Email: chanchal.manglunia@omnisciencecapital.com). Local office address of Securities and Exchange Board of India is SEBI Bhavan, C4-A, G Block, BKC, Mumbai - 400 051.

Omniscience Research is the research analyst division of Omniscience Capital Advisors Pvt. Ltd. which is registered as a non- individual Research Analyst with SEBI with a valid registration No. INH000020077, BSE Enlistment No.: 6517 and CIN U93000MH2017PTC290053. Principal Officer is Varun Sood (Contact No. 8879055519, Email: varun.sood@omnisciencecapital.com) and Compliance Officer is Chanchal Manglunia (Contact No. 9320816319, Email: chanchal.manglunia@omnisciencecapital.com). Local office address of Securities and Exchange Board of India is SEBI Bhavan, C4-A, G Block, BKC, Mumbai - 400 051.

Omniscience Asset Management, asset management division of Omniscience Capital Advisors Pvt. Ltd. which is registered as a Portfolio Manager with SEBI with a valid registration No. INP000009597 and CIN U93000MH2017PTC290053. Principal Officer is Ashwini Shami (Contact No. 9892140540, Email: ashwini.shami@omnisciencecapital.com) and Compliance Officer is Chanchal Manglunia (Contact No. 9320816319, Email: chanchal.manglunia@omnisciencecapital.com). Local office address of Securities and Exchange Board of India is SEBI Bhavan, C4-A, G Block, BKC, Mumbai – 400 051.

Investments in securities markets are subject to market risks. Read all related documents carefully before investing. SEBI registration, BSE membership, or NISM certification do not guarantee returns or the reliability of any intermediary. Past performance is not indicative of future outcomes, and no investment is risk-free. Investors are advised to carefully evaluate their financial goals, risk tolerance, and associated costs, and consult a qualified financial advisor before making investment decisions.

This communication does not imply guaranteed returns, a minimum yield, a specific target performance, or risk-free investment outcomes. Any securities mentioned may or may not be part of an investment strategy offered by the firm and are presented for illustrative purposes only. All views expressed are for informational purposes and do not constitute personalized recommendations or solicitations.

We, or our associates, may hold long or short positions in the securities of the companies mentioned in our reports/documents. This disclosure is made to ensure transparency and should be considered when evaluating the information provided. Omniscience Capital Advisors Private Limited disclaims any liability arising from investment decisions based on this material. For detailed disclosures regarding financial interests and potential conflicts of interest, kindly visit: https://www.omnisciencecapital.com/disclosure/