PCSMANTRA.COM · Study Article

IndiaAI Mission: 15 Compute Service Providers, 237 Approved Users

Technology / policy | News context: Lok Sabha answer, 29 July 2026 | Published: 19 October 2026

Background and Context

Expensive computing capacity like GPUs is the biggest barrier to building AI models. One of the IndiaAI Mission's seven pillars, the compute pillar, wants to take this capacity to startups, MSMEs, researchers and students through subsidised GPUs. Lok Sabha question 1812 asked: does the government keep state-wise GPU access figures and which are the top ten states/districts; what action followed the IISc-FSID CORE finding that GPU access favours those who write better proposals; and has central funding gone to the Visakhapatnam AI hub, or is a rule being made to reserve capacity on private AI-compute investment for public research?

Why in the News

In the answer to Lok Sabha unstarred question 1812 (Shri Sribharat Mathukumilli) on 29 July

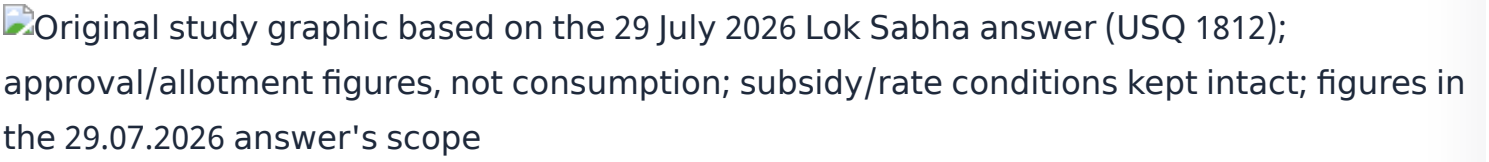
This site uses cookies to improve your experience. Read our Privacy Policy for details.

Accept Decline

allotted (sanctioned); demand aggregation and competitive price discovery secured rates significantly below market - GPUs are available to eligible users at less than 1 US dollar per hour with a subsidy of up to 40%. The compute portal operates as a single-window platform, with a ready reckoner and SOPs. A high-performance AI compute system of about 1.1 EFLOPS is being installed at the NIC Data Centre, Shastri Park, Delhi.

The Answer's Scope, as the Answer Gave It

Four cautions: (1) 237 approved users and 93.18 lakh GPU hours are allotment (sanctioned) figures - actual usage (consumed) figures are not in the answer; (2) "up to 40% subsidy" is the maximum ceiling for eligible users and "less than $1/hour" is an availability claim - neither is a universal charge; (3) state/district-wise allotment data is not maintained centrally - this does not mean absence of access in states; evaluation is project-based, not geographic; (4) the IISc-FSID CORE finding and a set-aside rule are the question's premises - the answer neither confirmed/denied the finding nor announced any set-aside rule.



Original study graphic based on the 29 July 2026 Lok Sabha answer (USQ 1812). Approval/allotment figures, not consumption; subsidy/rate conditions kept intact; figures in the 29.07.2026 answer's scope.

Key Facts and Figures

The answer's subject (as it is)	Lok Sabha unstarred question 1812, answered 29.07.2026; Ministry of Electronics and Information Technology (Minister of State Jitin Prasada); MP: Shri Sribharat Mathukumilli; two pages.
The compute pillar's goal (as it is)	A scalable, accessible and sovereign AI computing ecosystem under the IndiaAI Mission; targeted: startups, MSMEs, researchers, academia, government bodies and public-sector organisations.
The supply side (as it is)	15 compute service providers (CSPs) empaneled, in four phases, including high-end and mid-tier GPUs. Additionally, a high-

This site uses cookies to improve your experience. Read our Privacy Policy for details.

	installed at the NIC Data Centre, Shastri Park, Delhi - to build sovereign capacity and cut dependence on foreign compute.
The demand side (as it is)	237 end-users approved for subsidised compute support; about 93.18 lakh GPU hours allotted - among government bodies, startups and MSMEs, researchers, academic institutions, IndiaAI Fellows and students.
Rates and subsidy (as it is)	Through demand aggregation and competitive price discovery, compute rates were secured significantly below prevailing market prices; up to 40% subsidy for eligible users; GPU access available at less than 1 US dollar per hour (the answer's claim).
Portal and governance (as it is)	The IndiaAI Compute Portal (compute.indiaai.gov.in) is the single-window platform for discovery, requests and access; a ready reckoner and SOPs for billing/usage are published; approved projects' details (allotment, provider, subsidy) are on the portal for transparency. An End User Policy has been notified for GPU access and subsidy allotment - defining eligibility and method.
State/district figures (the answer's position)	State-wise and district-wise allotment data is not maintained centrally; per the answer the ecosystem is spread across the country and applications are evaluated project-wise, not on a geographic basis.
Visakhapatnam AI hub (as it is)	The AI hub announced in Visakhapatnam is a private investment; it is not funded by the IndiaAI Mission or any other central scheme. The government is cognizant of private-sector interest in Andhra Pradesh's AI ecosystem.

How Subsidised Compute Works

Joining supply and demand: supply was assembled by empaneling 15 CSPs in four phases; users request on the portal; through demand aggregation the government ran competitive price discovery and tied rates below market; up to 40% subsidy for eligible users - this design is the engine of "cheap compute".

This site uses cookies to improve your experience. Read our Privacy Policy for details.

What EFLOPS is: EFLOPS measures supercomputing speed (10 to the power 18 floating-point operations per second). The Shastri Park system's roughly 1.1 EFLOPS will add government-owned capacity parallel to the CSP-based model (in the answer "being installed" - not yet operational).

Reading the Answer Correctly

Approval vs use: 237 "approved for subsidized compute support" and 93.18 lakh hours "sanctioned" are the answer's own words. Reading them as active use or consumption would be wrong.

The rates' exact conditions: the subsidy is "up to 40%" - not 40% for every user; the rate "less than USD 1 per hour" is an availability claim, not a standard charge; the comparison boundary of "significantly below market" is not given in the answer.

Absence of mention is not absence of existence: not maintaining state/district-wise data centrally is a detail of the data system, not proof of lack of access in any state. The answer clarifies: evaluation is project-based.

Question premise vs answer: the IISc-FSID CORE "proposal-writing bias" finding and a set-aside rule on private investment are in the question. The answer took no direct position on them - it described the End User Policy, continuous review and outreach/handholding for new applicants.

Date scope: the figures are as of the 29.07.2026 answer; the Shastri Park system is "being installed" (in progress). No claim of the current number/status in October 2026 follows from this answer.

Which Conclusions Should Be Avoided?

"237 users are running GPUs": over-extension. 237 are approved; actual usage figures are not in the answer.

"93.18 lakh GPU hours consumed": a misreading. That many hours were allotted (sanctioned); no consumption (consumed) figure was given.

This site uses cookies to improve your experience. Read our Privacy Policy for details.

"The government accepted the bias" or "dismissed it": both unsupported. The answer does not speak directly on the CORE finding; it describes the policy system (End User Policy, review, outreach).

"Visakhapatnam hub is centrally funded": the opposite. Per the answer it is private investment; no IndiaAI Mission or central scheme money has gone to it.

"No access for smaller states/districts": conjecture. State-wise data is not maintained centrally - absence of mention is not proof of absence of access.

The Road Ahead

Our policy suggestions: (1) publish sanctioned vs consumed GPU-hour figures separately, so the pipeline-vs-consumption difference shows; (2) voluntary/aggregated publication of state/district-wise access - to measure regional imbalance without changing the geographic allotment policy; (3) a simplified application track for the proposal-writing barrier and a report of review outcomes; (4) regular publication on the portal of subsidy-eligibility and 40%-ceiling usage figures; (5) time-bound reporting of the Shastri Park EFLOPS system's progress and use.

Distinction of Concepts

Approved users are not active users; allotted (sanctioned) hours are not consumed hours; "up to 40%" is not a uniform 40%; "available below $1" is not a universal charge; absence of state data is not absence of access; a private AI hub is not a central scheme; the 29.07.2026 answer is not the October 2026 position.

Mains View

GS-III science-technology/economic development: designing AI compute as a public utility - subsidy, demand aggregation, sovereign capacity; the question of access equity in technology's democratisation. Keep in answers: the compute pillar among the seven pillars, 15 CSPs, 237/93.18 lakh (allotted), the 40%/$1 conditions, End User Policy, Shastri Park

This site uses cookies to improve your experience. Read our Privacy Policy for details.

Through subsidised GPU access, the IndiaAI compute pillar seeks to democratise the AI ecosystem. Analyse this model's structure and limits and suggest measures to ensure equal access for small/new applicants.

(Answer in 150 words) | 10 marks

Outline of the Answer

Introduction: the IndiaAI Mission, the compute pillar's place. Main body: the model (CSP empanelment, demand aggregation, portal, subsidy rates); the figures (15 CSPs, 237 approved, 93.18 lakh hours allotted); sovereign capacity (Shastri Park 1.1 EFLOPS); the access question (state-data position, End User Policy, outreach); the private-public boundary (Visakhapatnam). Conclusion: verify access and use evidence alongside empanelment and subsidy; approved users and allotted GPU hours are not synonyms for actual consumption.

Frequently Asked Questions (FAQ)

Short answer: through demand aggregation and competitive price discovery, rates were tied below market; there is up to 40% subsidy for eligible users.

Question: How do cheap-compute supply and demand join in IndiaAI?

Supply was assembled by empaneling 15 CSPs in four phases, and users request on the portal. Through demand aggregation and competitive price discovery, rates were tied below market; there is up to 40% subsidy for eligible users.

Question: Why should 237 approved users and allotted GPU hours not be taken as consumption?

Sanction vs use: 237 approvals and 93.18 lakh hours allotted show the pipeline, not consumption. The gap from sanctioned to consumed depends on time, project readiness and actual use - the answer gives no consumption figures.

Question: What does EFLOPS measure and what is the Shastri Park

This site uses cookies to improve your experience. Read our Privacy Policy for details.

EFLOPS measures supercomputing speed (10 to the power 18 floating-point operations per second). The Shastri Park system's roughly 1.1 EFLOPS will add government-owned capacity parallel to the CSP-based model (in the answer "being installed" - not operational).

Sources

- https://sansad.in/getFile/lsapps/loksabhaquestions/annex/188/AU1812_81bZkR.pdf - Lok Sabha unstarred question 1812 and answer, 29.07.2026 (two pages)

© 2026 PCSMANTRA.COM. All Rights Reserved.

This site uses cookies to improve your experience. Read our Privacy Policy for details.