Sovereign Minds, Quantum Eyes

A strategy, plan and roadmap for a dedicated armed-forces AI model, AI data centre and quantum-sensing upgrade of India's radars and communications

Six chapters, 24 figures and over 40 tables on what India has under way, what a sovereign AI stack needs, what quantum sensors can and cannot do for radar, and what it costs. 40 pages, evidence graded.

DELTA DOCTRINES

October 2026 | Open-source analysis. Reported figures are sourced; model results and judgements are labelled.
Not intelligence reporting. Contains no classified material. Cost inputs are Delta Doctrines assumptions.

CONTENTS

Contents and how to read this paper

This paper answers one question: how can India build a sovereign AI model and AI data centre dedicated to its armed forces, and use quantum sensors to modernise radars, navigation and communications, and what will it cost? It gives an approach, a strategy, a plan, a ten-year roadmap and a cost estimate that can be checked line by line.

Part	Content	Page
Summary	Executive summary; ten recommendations; how evidence is graded	3-4
1. Strategic case	Why a separate sovereign model; what India already has under way; funding landscape	5-7
2. Sarvam and the sovereign base	What Sarvam has built; what an armed-forces model needs beyond it; compute landscape	8-10
3. Architecture	Reference architecture; model tiers; data, security, governance; decision speed; human control	11-15
4. Quantum sensing	Navigation and timing; maturity; the quantum-radar question; radar and communication upgrades; QKD	16-20
5. Sizing and cost	Workloads, inputs, capex, annual spend, power, own versus rent, Monte Carlo, tornado, scenarios	21-30
6. Roadmap and delivery	Roadmap; phase plans; governance; funding and affordability; risks; procurement; first 100 days	31-38
Annexes	Limits of the evidence; source register	39-40

How evidence is graded. Defence reporting mixes primary facts, press accounts and analyst opinion. Every figure in this paper carries one of five labels, so the reader can see how far to trust it.

Label	Meaning	Example in this paper
Stated	Official or primary document, or a named official on the record	FY2026-27 capital outlay of Rs 2.19 lakh crore [S4]
Reported	Secondary press or trade reporting; not independently confirmed	Sarvam-HQ IDS five-year MoU [S1]; defence AI layer [S2]
Model	Delta Doctrines calculation from stated inputs; reproducible	Base case Rs 24,631 crore over ten years
Judgement	Delta Doctrines assessment with no empirical scale behind it	Maturity scores; priority matrix; risk map
Unverified	Claim found but not confirmed from a primary source	Indicative GPU rental rates above the Rs 65 subsidised rate

READ THIS FIRST

All costs are in constant 2026 rupees at Rs 90 to the US dollar, with no escalation. The cost model, the Monte Carlo ranges, the three scenarios and every unit price are Delta Doctrines assumptions. They are not Ministry of Defence figures and no Indian programme price has been published. The reported MoU and the reported defence AI layer are used as context, not as costed baselines.

EXECUTIVE SUMMARY

India can afford a sovereign armed-forces AI stack; the hard part is delivery

India's armed forces are adopting AI faster than they are building the foundations for it. The Army reported about 23 AI applications in Operation Sindoor [S21]. The Defence Secretary reported on 23 October 2026 that the Ministry and the IndiaAI Mission are developing a dedicated AI layer for the armed forces [S2]. Sarvam is reported to have signed a five-year agreement with HQ Integrated Defence Staff to build AI applications for all three Services on sovereign infrastructure [S1]. Each step is real but reported in pieces, with no published cost, architecture or timeline.

This paper proposes the missing frame: one sovereign model family, trained and run on a dedicated, air-gapped, tri-service compute estate, evaluated by an independent body, and linked to a quantum-sensing programme that hardens navigation, timing and radio-frequency sensing. The quantum programme is deliberately not built around microwave quantum radar, which the published engineering analyses do not support for early warning [S18][S19][S20].

BOTTOM LINE

A ten-year programme of about Rs 24,631 crore (USD 2.74 billion at the modal case; Rs 26,554 crore at the Monte Carlo median) would build a 8,192-GPU sovereign cluster, 600 edge nodes, a three-tier model family and four classes of quantum sensor. That is 11% of one year's defence capital outlay, spread over ten years. Money is not the binding constraint. Data access, accreditation, talent and supply of advanced chips are.

Item	Base case	Evidence
Ten-year total, 2027-2036 (modal inputs)	Rs 24,631 crore (about USD 2.74 billion)	Model
Monte Carlo median; P10 to P90	Rs 26,554 crore; Rs 24,900 to 28,324 crore	Model
AI track (data centres, edge, model, evaluation)	Rs 18,281 crore (74%)	Model
Quantum track (R and D, units, industry, QKD)	Rs 6,350 crore (26%)	Model
First five years (2027-2031)	Rs 15,072 crore (Monte Carlo median Rs 16,065 crore)	Model
Peak year	2030: Rs 4,614 crore, 2.1% of FY27 capital outlay	Model
Dedicated cluster at full build	8,192 Blackwell-class GPUs; 14.7 MW IT load; 19.2 MW at the meter	Model
Scale range (modal)	Lean Rs 9,844 crore; Base Rs 24,631 crore; Ambitious Rs 57,451 crore	Model
Defence capital outlay FY2026-27	Rs 2,19,306 crore; DRDO Rs 29,100 crore	Stated [S4]

Caveat. The base-case cluster is five times the training cluster reported for Sarvam 105B, but it is a planning number, not a requirement stated by any Indian authority. The model is a tool for choosing between scales, not a forecast of what the Ministry will pay.

EXECUTIVE SUMMARY

Ten recommendations, with owners, timing and cost

The recommendations follow the order in which dependencies bind: governance and data first, compute and models second, fielding third. Owners are proposals by Delta Doctrines, not announced assignments. The cost column is the base-case ten-year figure from the model in Chapter 5 and the lines sum to the total.

#	Recommendation	Owner (proposed)	When	Base-case cost, Rs crore
1	Create a Defence AI and Quantum Mission with one accountable head and a joint-staff charter	MoD with HQ IDS	Q1 2027	Within programme
2	Build a dedicated, air-gapped tri-service AI data centre in three geographically separate sites	DRDO, DDP, C-DAC style delivery agency	2027-2031	12,617 (GPU, facility, power, support)
3	Use civil IndiaAI compute as a bridge for unclassified work while the estate is built	IndiaAI, MoD	2027-2028	120
4	Train a three-tier sovereign model family: strategic, operational and tactical edge	Sarvam and other empanelled labs under MoD control	2027 onward	2,547
5	Deploy 600 ruggedised edge nodes with theatre commands	Services, with industry	2029-2033	1,200
6	Create an independent evaluation and red-team body that can stop a release	Outside Service chains of command	2028	340
7	Build the data fabric and cross-domain guards before scaling models	HQ IDS, Services, DRDO	2027-2031	1,200
8	Start quantum sensing with clocks, navigation and wideband receivers, not radar	DRDO QTRC, NQM hubs, industry	2027 onward	2,150
9	Buy quantum units only after field trials; scale from 2030	Services	2030-2036	3,700
10	Fund talent and the supplier base as part of the programme, not afterwards	MoD, MeitY, DST	2027 onward	756

WHY THE ORDER MATTERS

A cluster without accredited data produces a model that cannot be used. A model without an evaluation body produces a tool nobody can trust in a crisis. A sensor without a field trial produces a procurement the Services will not accept. The programme is staged so that each of these gates must be passed before the next tranche of money is released (Chapter 6).

What this paper does not recommend. It does not recommend a single national model for civil and military use, a build-first approach that waits for perfect data, microwave quantum radar as a near-term procurement, or reliance on foreign cloud regions for classified work. The reasons are set out in Chapters 1, 3 and 4.

1. STRATEGIC CASE

Why the armed forces need a model and a data centre of their own

Three problems make a shared civil model insufficient for defence work. The first is **classification**: operational orders, signals intelligence and target data cannot leave an accredited enclave, so they cannot be used with a model hosted elsewhere. The second is **assurance**: a model that sometimes invents facts is tolerable in customer service and dangerous in targeting. The third is **dependence**: the Defence Secretary cited technology being withheld and possible backdoors or kill switches as the risks of foreign supply [S2].

Eight years from a task force to a dedicated defence AI layer

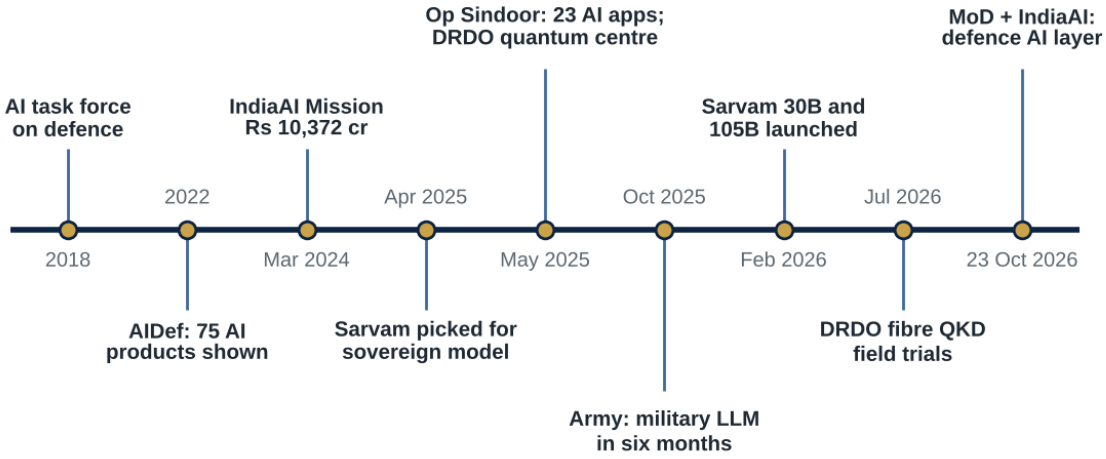


Figure 1. Eight years from a task force to a dedicated defence AI layer. Events are as reported [S1][S2][S13][S15][S21]; the 2018 and 2022 entries are carried from earlier research and are not re-verified here.

Problem	Why a shared civil model fails	What a dedicated stack changes
Classification	Data cannot cross into a commercial cloud	Model and data stay inside an accredited, air-gapped estate
Assurance	Public benchmarks do not test military tasks; hallucination is high on some reported tests [S7]	Military test sets, red teams and a release gate
Dependence	Weights, chips and updates can be withheld	Weights held in India; open-weight base reduces lock-in
Speed	Queues and rate limits in shared services	Reserved capacity and edge inference
Language and context	Training data thin on military doctrine and Indian operational context	Domain data programme and continued pre-training

Evidence status. The problems are well-argued in public reporting but the severity for India is a Delta Doctrines judgement. No published Indian evaluation compares a general model with a defence-tuned model on military tasks.

1. STRATEGIC CASE

What is already under way, and what is missing

India is not starting from zero. The table lists the public record. It shows applications, a centre, a layer and an MoU, but no architecture, no cost, no independent evaluation and no scale for the compute. That gap is where this paper aims.

Programme or event	Date	What is reported	Label	Source
Op Sindoor AI use	Oct 2025 briefing	About 23 applications for a common operational, intelligence and logistics picture; 72-hour weather forecasts; an electronic-intelligence application claimed over 90% accuracy in tracking sensors and units	Stated (briefing); claim unverified	[S21]
Army military LLM	Oct 2025	DG EME said the military LLM would be fully functional after testing in six months	Stated; delivery not confirmed	[S21]
Rs 300 crore defence AI centre	22 Apr 2026	Ministry planning a Centre of Excellence; advanced talks with Sarvam and other domestic labs; funder not stated	Reported	[S3]
DRDO Quantum Technology Research Centre	27 May 2025	Opened at Metcalfe House, Delhi: QKD platforms, a chip-scale atomic clock, an optically pumped magnetometer, vapour cells	Stated (PTI)	[S15]
Defence AI layer	23 Oct 2026	MoD and IndiaAI developing a dedicated AI layer for decision support, target acquisition and data fusion; compute already procured; cost not stated	Reported	[S2]
Sarvam-HQ IDS MoU	Reported 27 Oct 2026	Five-year engagement to design and deploy AI applications for all three Services; shared foundation models and agentic infrastructure; funding not disclosed	Reported	[S1]
Trust framework	22 Oct 2026	A Member of Parliament referred to a domestic framework for assessing trustworthiness of defence AI: reliability, safety, transparency	Reported	[S2]

WHAT THE RECORD DOES NOT SAY

None of the reported items states the number of GPUs, the classification levels covered, who evaluates the models, how the Services share data, or what the programme costs. The 23 October statement said compute has been procured but gave no size [S2]. The MoU funding is not disclosed [S1]. This paper therefore treats the reported items as signs of direction and sizes the stack independently.

Implication for design. If the Ministry's layer and Sarvam's MoU proceed, this paper's architecture can be read as the common frame for them: one model family, one data fabric, one evaluation body. If they proceed separately by Service, the risk is three incompatible stacks. Recommendation 1 (a single mission head) is aimed at that risk.

1. STRATEGIC CASE

The funding landscape: the programme is large next to AI missions, small next to defence capital

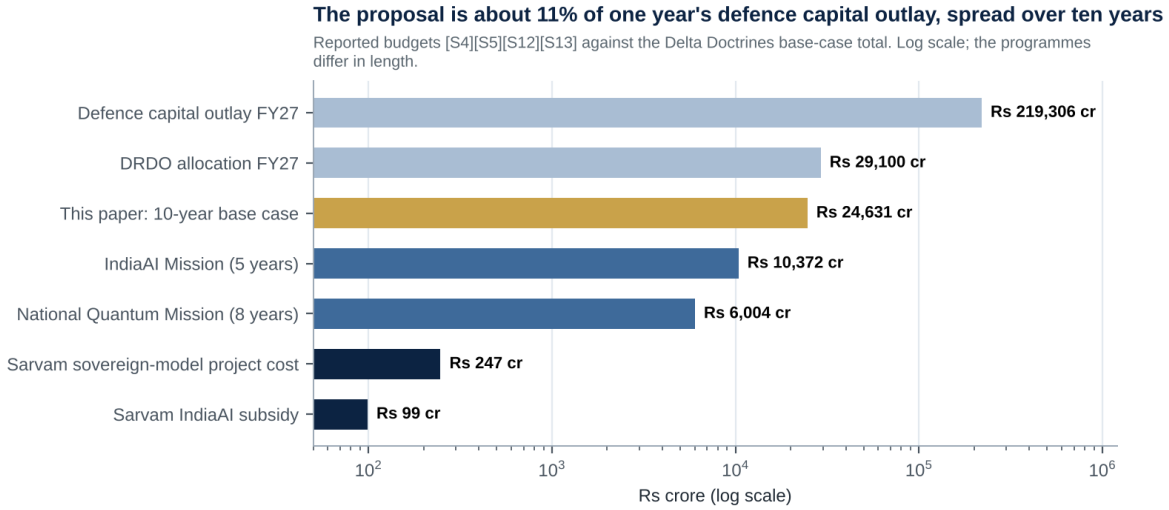


Figure 2. Reported budgets against the Delta Doctrines base-case total. Log scale. The programmes differ in length and are not directly comparable.

Budget line	Amount, Rs crore	Period	Label	Source
Defence budget FY2026-27 (total)	7,85,000 (Rs 7.85 lakh crore)	1 year	Stated	[S4]
Defence capital outlay	2,19,306	1 year	Stated	[S4]
DRDO allocation	29,100	1 year	Stated	[S4]
IndiaAI Mission outlay	10,372	5 years	Stated	[S2][S5]
National Quantum Mission	6,004	8 years	Stated	[S12]
Sarvam sovereign-model project cost	247	Project	Stated	[S5]
Of which IndiaAI subsidy	99	Project	Stated	[S5]
This paper, ten-year base case	24,631	10 years	Model	-

The proposal costs about 2.4 times the entire IndiaAI Mission and about four times the National Quantum Mission, but only 11% of one year of defence capital outlay. On an annual basis it never exceeds 2.1% of capital outlay (Figure 22, Chapter 6). A fair reading is that the question for the Ministry is not whether it can afford the programme, but whether it will release money in tranches against gates, which is what the roadmap proposes.

> WHERE THE MONEY WOULD COME FROM
>
> The proposal assumes the Defence budget carries most of the cost, with civil missions (IndiaAI compute, NQM hubs) seeding the bridge compute and the quantum research base. A funding split is proposed in Chapter 6. It is a Delta Doctrines proposal and not an official allocation.

Exchange rate. Rs 90 to the dollar is an assumption. A 10% rupee move shifts the dollar-priced share of cost (GPUs, imported components) by about the same amount, so the base total is sensitive to it; the Monte Carlo ranges in Chapter 5 do not include currency risk.

2. SARVAM AND THE SOVEREIGN BASE

What Sarvam has built, and what independent testing says

Sarvam is the sovereign base most advanced on the public record: two open-weight models trained in India under the IndiaAI Mission, and, reportedly, a five-year MoU with HQ IDS [S1]. Sarvam received 4,096 H100 GPUs from Yotta, the largest single allocation under the mission, with a subsidy of Rs 98.68 crore against a project cost of Rs 246.71 crore [S5].

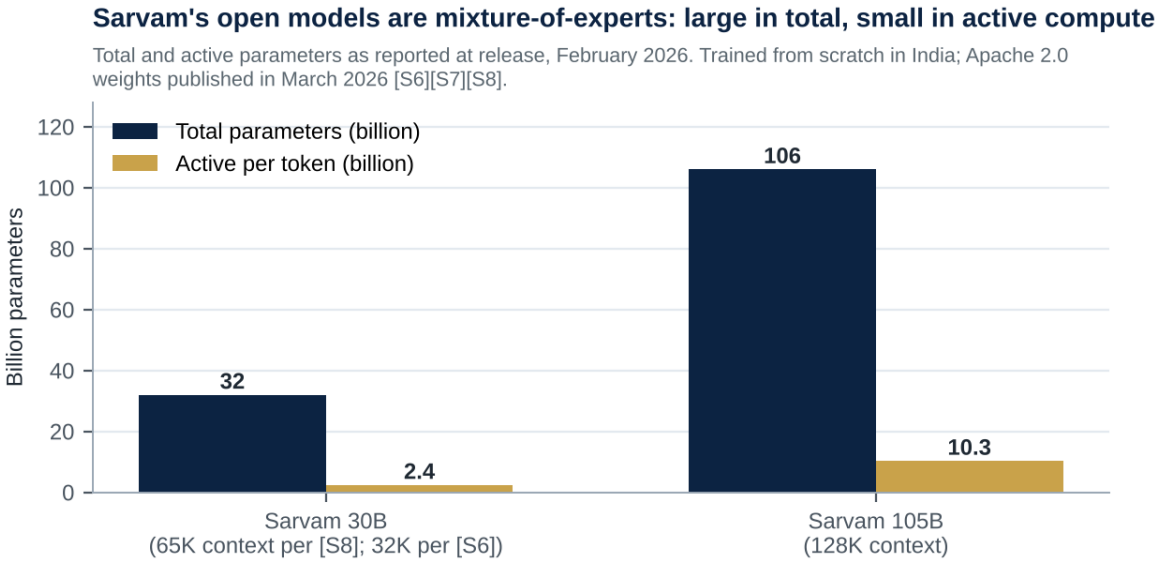


Figure 3. Sarvam's open models: total and active parameters as reported. Active parameters of the 105B model are reported as 9 to 10.3 billion in different sources; the chart uses 10.3.

Attribute	Sarvam 30B	Sarvam 105B	Label	Source
Architecture	Mixture of experts	Mixture of experts	Reported	[S6][S7]
Total parameters	32 billion	106 billion	Reported	[S7]
Active parameters per token	About 2.4 billion	About 10 billion (9 to 10.3 billion in different reports)	Reported; conflicting	[S6][S7]
Context length	65K tokens ([S8]); 32K in [S6]	128K tokens	Reported; conflicting	[S6][S7][S8]
Licence	Apache 2.0	Apache 2.0	Reported	[S7]
Availability	Sarvam API, Hugging Face, AIKosh	Same	Reported	[S7]
Modality	Text in, text out	Text in, text out	Reported	[S7]
Artificial Analysis Intelligence Index	12	18	Reported (independent)	[S7]
Hallucination (AA-Omniscience index)	-72	-60	Reported (independent)	[S7]
TerminalBench Hard	2.3%	1.5%	Reported (independent)	[S7]
Training cluster	-	About 1,000 H100s reported	Reported	[S5]

Reading the independent scores. Artificial Analysis rates both models behind peers of similar size on overall intelligence, with relatively stronger agentic scores for the 105B, and flags high hallucination rates: both models attempt far more questions than they abstain from [S7]. Sarvam's own benchmark claims are not independently verified in this paper.

2. SARVAM AND THE SOVEREIGN BASE

What an armed-forces model needs that a civil model does not

The honest finding is that Sarvam's open models are a credible starting point for a sovereign defence model family, and no more than that. Their reported scores trail peers, they are text-only, and their calibration is weak. A defence programme should treat them as a base to be continued, tested and, where necessary, replaced, rather than as a finished product.

Requirement of an armed-forces model	Sarvam today (public record)	Gap	Closing action
Handles classified data in an air-gapped estate	Open weights can be run offline [S7]	Accreditation, hardware and key management	Dedicated estate; accredit by tier
Military language and doctrine	Indic-language focus; general text model [S7]	Doctrine, order of battle, signals, imagery vocabulary	Domain data programme; continued pre-training
Multimodal fusion: imagery, signals, geospatial	Text-only models [S7]; MoU mentions geospatial and perception systems [S1]	Vision, SAR, ELINT models not demonstrated publicly	Operational-tier multimodal models
Low hallucination, calibrated refusal	Negative AA-Omniscience scores [S7]	Must abstain when unsure	Calibration training; verifier models; red team
Edge inference at 1 to 8 billion parameters	30B active about 2.4B [S7]	Quantised builds, ruggedised hardware	Tactical tier; distillation
Agentic workflows across systems	Agentic index 25 for 105B [S7]; MoU mentions agentic infrastructure [S1]	Tool permissions, audit, human gates	Policy-gated agent layer
Independent assurance	None public	Entire function absent	Evaluation body (Recommendation 6)

DESIGN CHOICE: BASE MODEL, NOT VENDOR

The programme should hold the weights, the training data, the evaluation sets and the right to retrain. Sarvam can be a lead partner under the MoU [S1], but the Ministry should be able to continue with another empanelled lab if performance or terms fail. Four labs were funded in the first IndiaAI phase [S9]: Sarvam, Soket, Gnani and Gan. A competitive second source is a deliberate part of the plan.

Why continue from open weights. Training a 100-billion-parameter-class model from scratch needs a cluster Sarvam reportedly did not have for 105B (about 1,000 H100s [S5]). Continuing from open weights inside an air-gapped estate lets the programme add doctrine, signals and imagery skills at far lower cost than a new pre-training run. This is a Delta Doctrines judgement; the model programme budget (Rs 1,947 crore over ten years) allows either path.

2. SARVAM AND THE SOVEREIGN BASE

India's compute is large on paper; a defence cluster would be a fraction of it

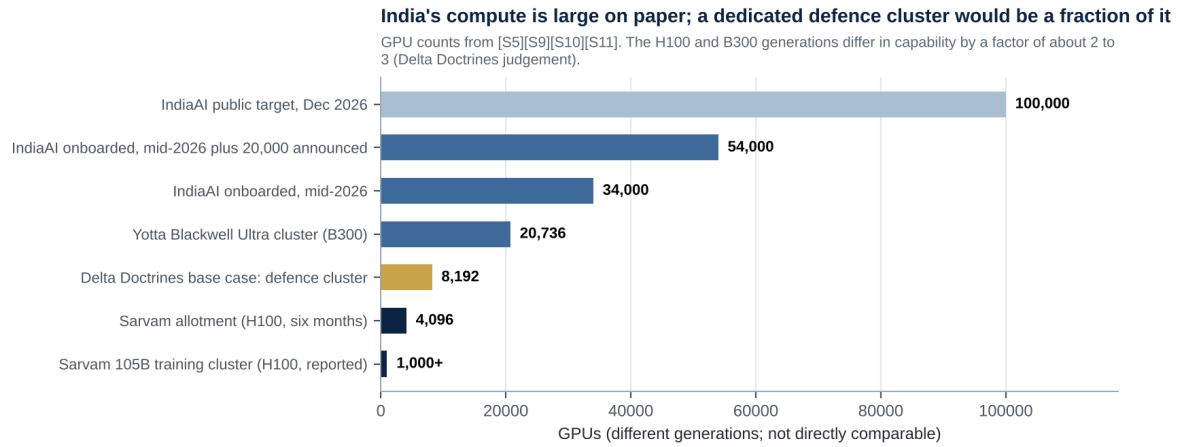


Figure 4. GPU counts across pools. Generations differ in capability by a factor of about 2 to 3; the bars are not directly comparable.

Compute pool	GPUs	Generation	Status	Source
IndiaAI public target, Dec 2026	100,000	Mixed	Target	[S5][S9]
IndiaAI onboarded plus announced additions	About 54,000	Mixed	Reported	[S9]
IndiaAI onboarded, mid-2026	34,381 (14 providers)	Mixed, H100 largely	Reported	[S9][S10]
Yotta Blackwell Ultra cluster	20,736	B300	Planned, Noida, over USD 2 billion	[S11]
Delta Doctrines base-case defence cluster	8,192	Blackwell-class	Proposal	Model
Sarvam allotment	4,096	H100	Allotted	[S5]
Sarvam 105B training	About 1,000	H100	Reported	[S5]

The subsidised IndiaAI rate averages Rs 65 per GPU-hour, with H100 at Rs 92 [S9]. That is cheap, but it is civil capacity on shared, commercial premises. It is suitable for unclassified training of the base model and for the bridge period, but it cannot hold classified data. The proposal therefore splits the work: civil compute for the bridge, a dedicated estate for everything else.

> WHY NOT SIMPLY BUY MORE CIVIL GPU-HOURS?
> At 70% average utilisation, owning a cluster costs about Rs 147 per H100-equivalent GPU-hour (Delta Doctrines model, Chapter 5), above the Rs 65 subsidised rate. Sovereignty, not price, is what a dedicated cluster buys. If the Ministry only needed unclassified compute, renting would be cheaper. This paper's argument rests on the classified-data case.

Power as a limiter. Yotta's 60 MW Noida data centre houses the 20,736-GPU cluster [S11]. The proposed defence cluster needs about 19 MW at the meter, about a third of that site. Siting across three locations (Chapter 3) raises the cost by an estimated 15% (the multisite factor in the model).

3. ARCHITECTURE

Reference architecture: a sovereign core, theatre nodes and edge

The architecture has three rings. A sovereign core trains and hosts the largest models and holds the data fabric. Theatre nodes put operational models next to the commands that use them. Edge nodes run small models where connectivity cannot be assumed. Service networks plug in through cross-domain guards rather than direct links.

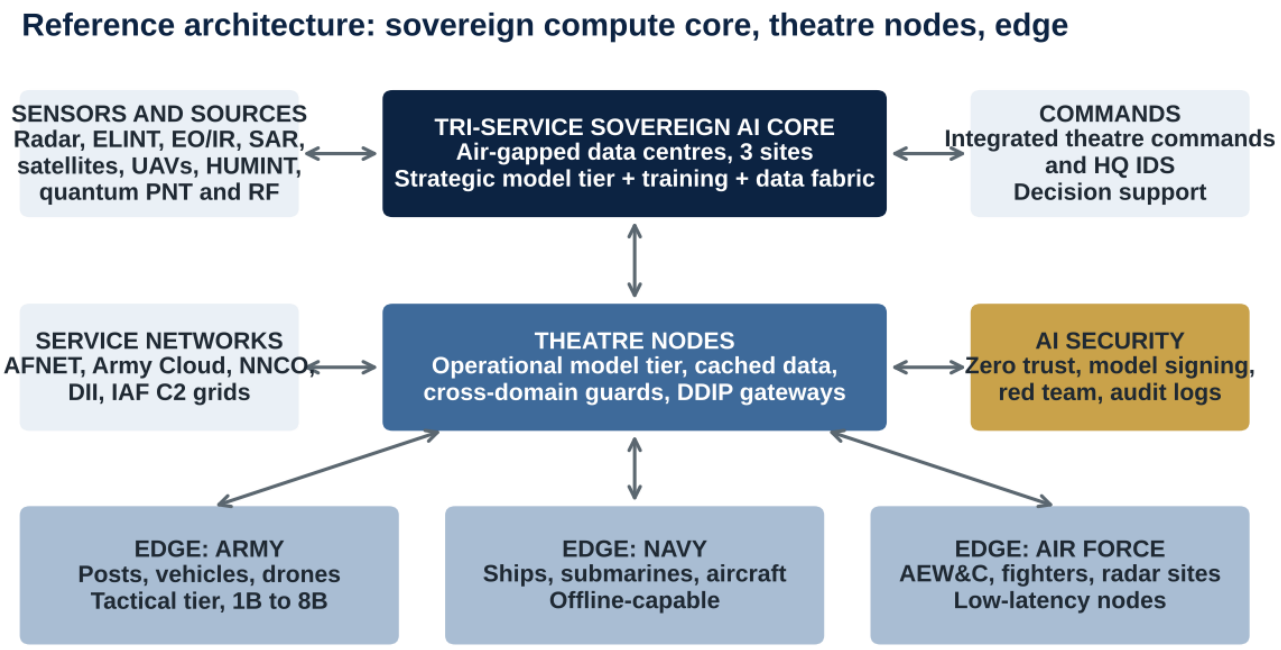


Figure 5. Reference architecture. Network names (AFNET, Army Cloud, NNCO, DII) are shown as examples of Service networks; connections are Delta Doctrines design, not a description of existing links.

Layer	What it is	Scale in base case	Security level
Sovereign AI core	Three air-gapped, geographically separate data centres: training, strategic-tier inference, data fabric	8,192 Blackwell-class GPUs; 14.7 MW IT load; 19.2 MW at the meter	Highest tier; no external connection
Theatre nodes	Operational-tier inference, cached data, cross-domain guards, gateways to Service networks	Co-located with integrated theatre commands (number set by the Services)	Accredited to the commands' classification
Edge nodes	Ruggedised tactical-tier inference on ships, aircraft, vehicles, posts	600 nodes at about Rs 2.0 crore each (modal)	Per platform; offline-capable
Data fabric	Catalogue, lineage, labelling, cross-domain transfer, retention	Rs 1,200 crore integration and Rs 600 crore data programme	Mandatory access control
AI security	Zero trust, model signing, supply-chain checks, red team, audit	Within evaluation and integration budgets	Independent of operators

Why three sites. A single site is a single point of failure and a single target. Three sites allow active-active training, disaster recovery and separation of duties between Services. The model adds a 15% multisite premium and a 30% defence-hardening premium on the facility cost (both Delta Doctrines assumptions; Chapter 5).

3. ARCHITECTURE

A three-tier model family, trained once and distilled down

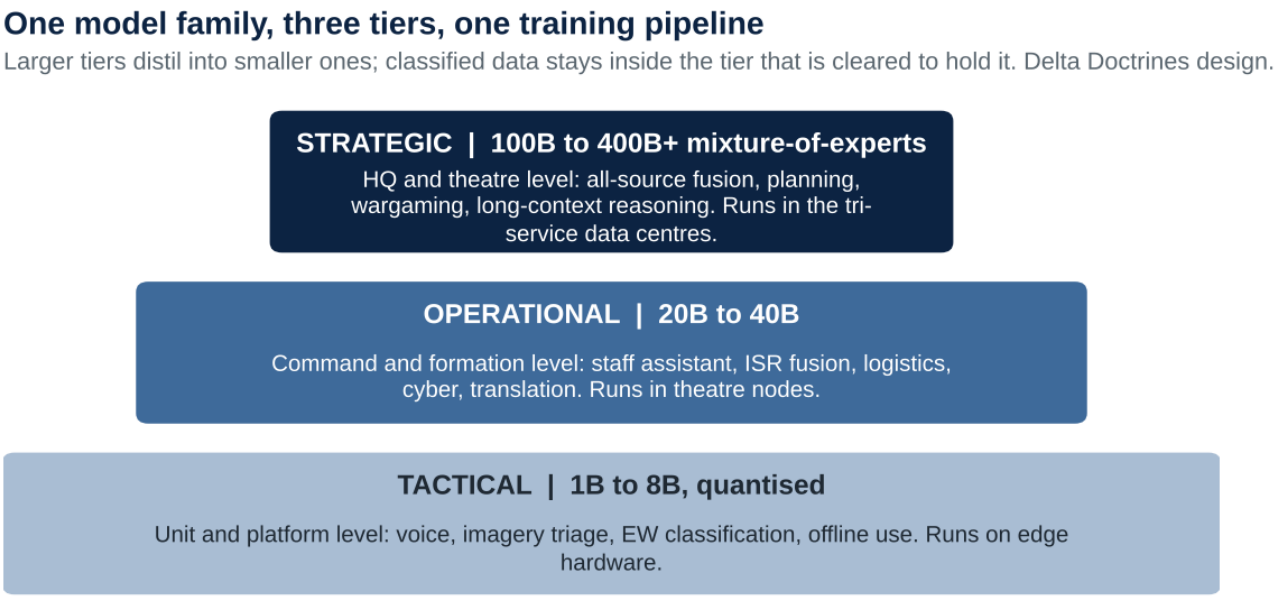


Figure 6. One model family, three tiers, one training pipeline. Larger tiers distil into smaller ones; classified data stays in the tier cleared to hold it.

Tier	Size	Where it runs	Typical tasks	Training route
Strategic	100B to 400B+ parameters, mixture of experts	Tri-service data centres	All-source fusion, planning, wargaming, long-context reasoning	Continue from open base; add doctrine, signals, imagery
Operational	20B to 40B	Theatre nodes	Staff assistant, ISR fusion, logistics, cyber, translation	Distil from strategic tier; fine-tune per command
Tactical	1B to 8B, quantised	Edge hardware	Voice, imagery triage, EW classification, offline use	Distil; quantise; test per platform

Why a family and not one model. A single very large model would be too slow and too costly to run at the edge and too sensitive to hold in every theatre. A family lets each level use what it needs. The risk is version drift; the pipeline therefore signs every release and tests each tier against the same evaluation sets.

Training stage	Compute, GPU-months (base case)	Data	Gate before next stage
Continued pre-training of base	About 8,000 to 12,000	Doctrine, manuals, unclassified open data	Benchmarks beat the base on military test set
Supervised and preference tuning	About 2,000 to 4,000	Labelled staff tasks, red-team feedback	Hallucination rate below set threshold
Multimodal adapters	About 6,000 to 10,000	Imagery, SAR, signals (classified)	Detection and recall targets met on held-out data
Distillation to smaller tiers	About 1,500 to 3,000	Synthetic and real	Edge latency and accuracy targets met

The GPU-month figures in the second table are Delta Doctrines sizing judgements consistent with the 2,048-GPU training allocation in the sizing figure. No Indian figure has been published. They indicate order of magnitude only.

3. ARCHITECTURE

Data, security and governance: the controls that decide whether the system is usable

Most defence AI programmes stall on data access and accreditation, not on compute. The 23 October statement named data fusion and decision support as aims [S2], but fusion requires owners of intelligence, operations and logistics data to share it. The architecture therefore starts with rules for data, before models.

Control area	Design rule	Why it matters	Label
Classification tiers	Train and run separate model instances per tier; never mix tiers in one set of weights without a release review	Weights can leak the data they were trained on	Judgement
Cross-domain guards	One-way and reviewed transfers between networks; automated checks plus human release for upward transfers	The largest single insider and spillage risk	Judgement
Data stewardship	Each data owner (Service, DRDO, intelligence body) signs a data-sharing charter; HQ IDS arbitrates	Interoperability is a governance, not a compute, problem	Judgement
Zero trust	Every user, device and model identified and authorised per request	Limits lateral movement after a breach	Judgement
Model supply chain	Reproducible builds; signed weights; scan training data for poisoning	A poisoned or backdoored model is a strategic vulnerability	Judgement
Hardware provenance	Accept chips and servers only from vetted suppliers; inspect firmware	Foreign dependence was the stated concern [S2]	Judgement
Logging and audit	Immutable logs of prompts, outputs, tool calls, human decisions	Accountability and after-action review	Judgement
Retention and deletion	Tier-specific retention; deletion tested	Limits long-term exposure	Judgement

THE DATA PROGRAMME IS A FIRST-CLASS WORKSTREAM

The base case budgets Rs 600 crore over ten years for data labelling, cleaning, lineage and synthetic data, and Rs 1,200 crore for integration and the data fabric, together 7.3% of the total. These are the lowest-visibility and most schedule-critical lines in the plan. A delay of one year in data access would delay the first useful operational model by at least as long, whatever the GPU count.

Open question. The Ministry has not published an accreditation scheme for AI workloads. The proposed tiers, guards and release rules are Delta Doctrines design principles drawn from general information-security practice. They need review by the cyber and intelligence agencies before adoption.

3. ARCHITECTURE

Fast decisions: where AI removes time, and where it must not

The stated aim of the defence AI layer is a sensor-to-shooter loop that supports decision-making, target acquisition and data fusion [S2][S1]. The figure shows how the time in a time-critical targeting chain might be redistributed. The stage times are illustrative; no public Indian data on them exists.

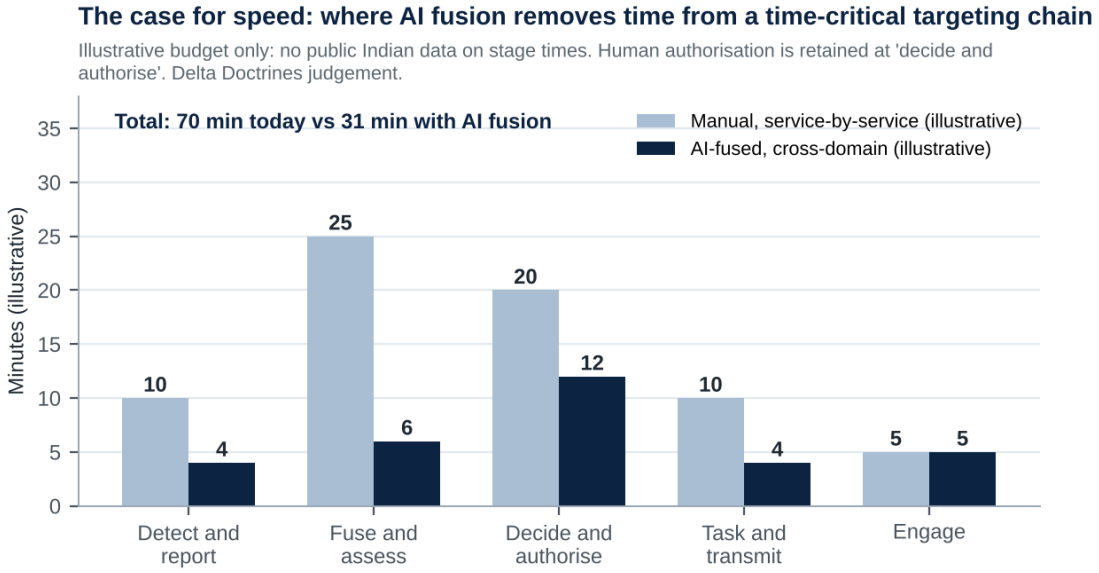


Figure 7. Where AI fusion removes time from a time-critical targeting chain. Illustrative budget (Delta Doctrines judgement); human authorisation is retained.

Stage	Manual today (min)	AI-fused (min)	What changes
Detect and report	10	4	Automatic cueing of sensor reports into a common picture
Fuse and assess	25	6	Cross-domain correlation by model; analyst reviews
Decide and authorise	20	12	Decision support; human authorises; time saved is mainly in preparing options
Task and transmit	10	4	Automatic tasking packages to platforms
Engage	5	5	Unchanged
Total	70	31	Illustrative only

What the speed-up depends on. Most of the saving is in fusion and tasking, where machines compare and route faster than people. Very little of it is in the human decision itself. This is a design intent: the time saved comes from better preparation of options, not from removing the authoriser.

REAL-TIME INTEGRATION ACROSS DOMAINS

Cross-domain integration needs common data standards, time synchronisation (a quantum-clock use case, Chapter 4), shared track identifiers and message gateways between Service networks. Without them a fast model produces fast, incompatible answers. Integration and data fabric is therefore the second-largest AI-side line after compute, at Rs 1,200 crore.

3. ARCHITECTURE

Human control and independent assurance

A trustworthy defence AI needs institutions as much as technology. The reported trust framework (reliability, safety, transparency [S2]) is a starting point. This paper proposes that it be given a body with authority and money: an independent evaluation and red-team unit, budgeted at Rs 340 crore over ten years (1.4% of the total).

Principle	Requirement	Mechanism	Owner
Meaningful human control	A named human authorises any use of force; the system cannot self-authorise	Policy-gated agent layer; approval logs	Operational commands
Independent evaluation	A body outside the Service chains of command tests every release and can block it	Military test sets; red team; release certificate	Evaluation and red-team body
Calibration	The model must say when it does not know	Abstention training; verifier models; confidence reporting	Model programme
Traceability	Every output linked to sources and model version	Retrieval with citations; versioned weights	Data fabric
Adversarial robustness	Tested against prompt injection, data poisoning, spoofed inputs	Red-team cycles each release	Evaluation body
Fail-safe operation	Degraded mode when AI is unavailable	Manual procedures drilled; fallbacks	Services
Review and sunset	Annual external review; right to retire a model	Charter and audit	Mission board

WHY INDEPENDENCE MATTERS

The independent evaluation of Sarvam's open models shows high hallucination rates on some tests [S7]. Those results came from outside the lab. A defence programme needs the same outside-in view for classified use, on test sets that the model builder cannot see. A programme that grades its own work will report success.

Release gates. Three gates are proposed: (1) a model enters pilot only if it beats the previous version on a military test set and its hallucination rate is under a set threshold; (2) it enters operational use only after a red-team campaign and a field exercise; (3) it remains in use only if monitoring shows no drift. Thresholds are for the evaluation body to set, not this paper.

4. QUANTUM SENSING

The navigation and timing problem that quantum sensors can address

ORF's May 2026 paper argues that quantum PNT, which does not rely on external signals, can resist jamming and spoofing, and that its most viable near-term use is to augment classical sensors [S14]. It reports that Russian GPS jamming cut the hit rate of US-made precision weapons from over 50% to under 10% in Ukraine [S14], a figure reported in the paper and not verified here. It also links jamming and spoofing along the India-Pakistan and India-China borders to the need for alternatives [S14].

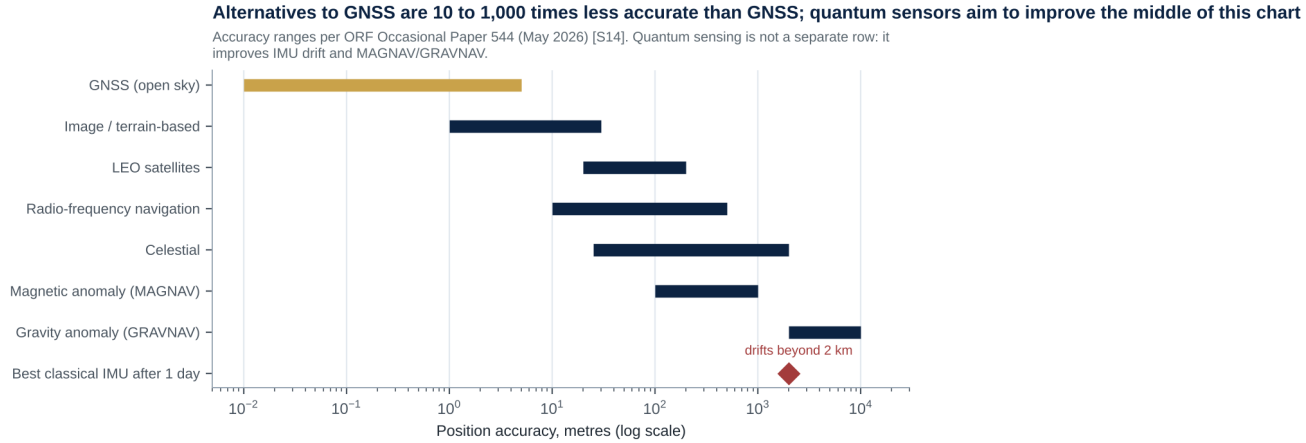


Figure 8. Accuracy of GNSS and its alternatives (log scale), per ORF Occasional Paper 544 [S14]. Quantum sensing is not a row in the source table; it improves the inertial and map-matching rows.

Alternative to GNSS	Accuracy, metres	Limit	Quantum sensing role
GNSS (open sky)	0.01 to 5	Jamming and spoofing	Clocks and receivers for holdover and anti-spoofing
Image or terrain-based	1 to 30	Needs maps and visibility	Improves inertial fusion
Radio-frequency navigation	10 to 500	Needs transmitters	Rydberg receivers
LEO satellites	20 to 200	Needs constellation	None direct
Celestial	25 to 2,000	Weather and daylight	None direct
Magnetic anomaly (MAGNAV)	100 to 1,000	Needs magnetic maps	Quantum magnetometers
Gravity anomaly (GRAVNAV)	2,000 to 10,000	Needs gravity maps	Quantum gravimeters
Best classical IMU	Drifts beyond 2 km in a day	Error grows with time	Atom-interferometer IMU

Implication. The case for quantum sensing in India is first about navigation and timing, because those are the functions that fail when satellite signals fail, and because they underpin everything else, including the cross-domain fusion in Chapter 3. Radar is a secondary and harder case (page 18).

4. QUANTUM SENSING

How mature are the sensors? Uneven, and mostly short of the field

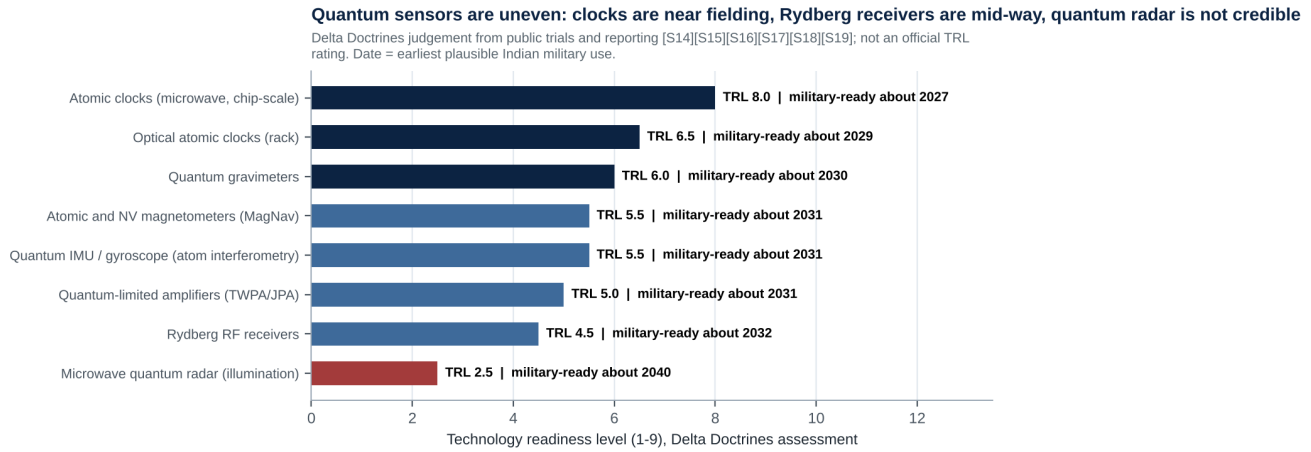


Figure 9. Maturity of quantum-sensing technologies: Delta Doctrines judgement from public trials and reporting [S14]-[S19]; not an official TRL rating.

Technology	Readiness (1-9)	Military-ready (earliest)	Indian base	Judgement note
Chip-scale atomic clocks	8.0	2027	DRDO ultra-small clock and vapour-cell work [S15]	Closest to fielding; the QTRC lists a coherent-population-trapping clock
Optical atomic clocks (rack)	6.5	2029	Academic labs; NQM hubs	Needed for radar references and timing networks
Quantum gravimeters	6.0	2030	Limited	Useful with gravity maps; low priority
Atom-interferometer IMU	5.5	2031	Early R and D	Largest payoff for GNSS-denied navigation
Atomic and NV magnetometers	5.5	2031	DRDO optically pumped magnetometer [S15]	MAGNAV; ASW; mine detection
Quantum-limited amplifiers	5.0	2031	Limited; cryogenic dependence	Strategic radar receivers
Rydberg RF receivers	4.5	2032	Limited; vapour cells at DRDO [S15]	Laboratory spectrum analyser to 20 GHz shown in the US [S16]
Microwave quantum radar	2.5	About 2040	None	Not credible for early warning (next page)

What the evidence shows. The US Army Research Laboratory demonstrated in 2021 a Rydberg-atom receiver that sampled from zero frequency to 20 GHz, with the researchers noting that significant physics and engineering work remained before a field-testable device [S16]. The Indian base is a research centre opened in May 2025 at Metcalfe House, Delhi, with QKD platforms, an ultra-small atomic clock, an optically pumped magnetometer, and vapour-cell work [S15]. The scores combine these public items with Delta Doctrines judgement; they are not measured.

PLANNING RULE

Fund clocks and navigation first; fund Rydberg and amplifier research in parallel; field-test everything before buying in quantity; and do not plan a procurement around microwave quantum radar. Unit purchases in the cost model start only in 2030 (Chapter 5).

4. QUANTUM SENSING

The quantum-radar question: the evidence does not support an early-warning radar

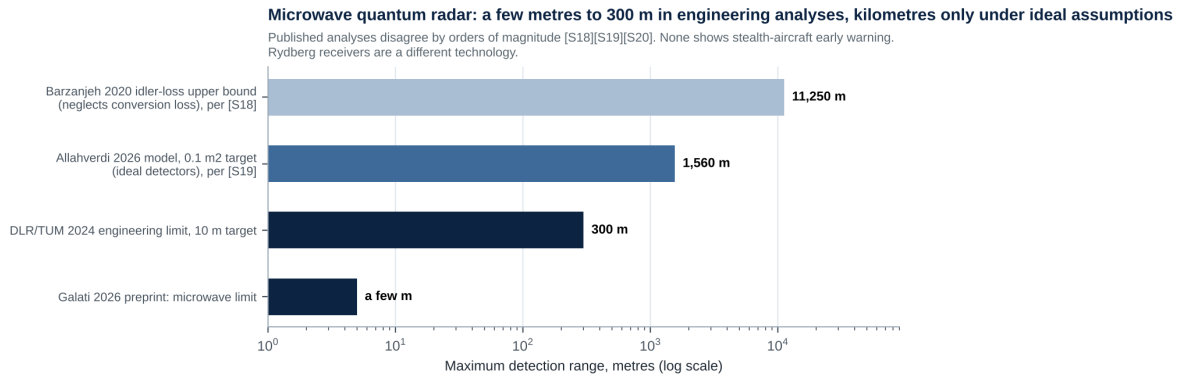


Figure 10. Maximum detection range from four published analyses. They disagree by orders of magnitude because of their assumptions about detectors and losses.

Study	Finding	Assumption	Source
Bischeltsrieder et al., DLR/TUM, 2024	Maximum feasible range about 300 m for a 10 m target; about 1 km if antenna gain is raised to 30 dBi; advantage over classical radar at most about 3 dB	Engineering constraints incl. conversion losses	[S18]
Galati, Pavan and Daum, 2026 (preprint, not peer-reviewed)	Microwave quantum radar range should not exceed a few metres; a 1 km detection at X-band needs about 479 hours of integration	X-band, 1 m2 target, 1 GHz bandwidth	[S20]
Allahverdi and Motazedifard, 2026 (arXiv)	About 1.2 to 1.6 km for a 0.1 m2 target and up to about 2 km with longer integration; falls to 4.8 to 8.5 m with classical antennas; not adequate for stealth-aircraft early warning	Ideal detectors and transducers	[S19]
Barzanjeh et al., 2020, as cited in [S18]	Idler-loss upper bound 11.25 km	Neglects microwave-optical conversion losses	[S18]

Finding. The better-quality engineering analyses place microwave quantum radar at metres to hundreds of metres, and the most optimistic model reaches about a kilometre or two under ideal assumptions and concludes it is inadequate for stealth-aircraft early warning [S18][S19][S20]. The theoretical quantum advantage over a classical radar is about 3 dB at best, and is outweighed by waveform penalties in one analysis [S20]. Claims of 100 km or more circulate in the press; this paper found no published data supporting them and does not use them.

WHAT THIS MEANS FOR INDIA'S RADAR UPGRADE

A programme that promised stealth detection from quantum radar would very likely fail. The credible quantum contributions to radar are different: (1) low-phase-noise optical clock references that improve clutter rejection and coherent integration; (2) quantum-limited amplifiers that lower receiver noise in strategic radars; (3) Rydberg receivers for wideband electronic-support measures. The next page ranks them.

Note on the Sudarshan Chakra reporting. A study-site claim of quantum-enhanced radar numbers for the Sudarshan Chakra air-defence concept was seen in earlier research and could not be confirmed from a primary source. It is excluded from the analysis.

4. QUANTUM SENSING

Radar and communication upgrades that make sense, in priority order

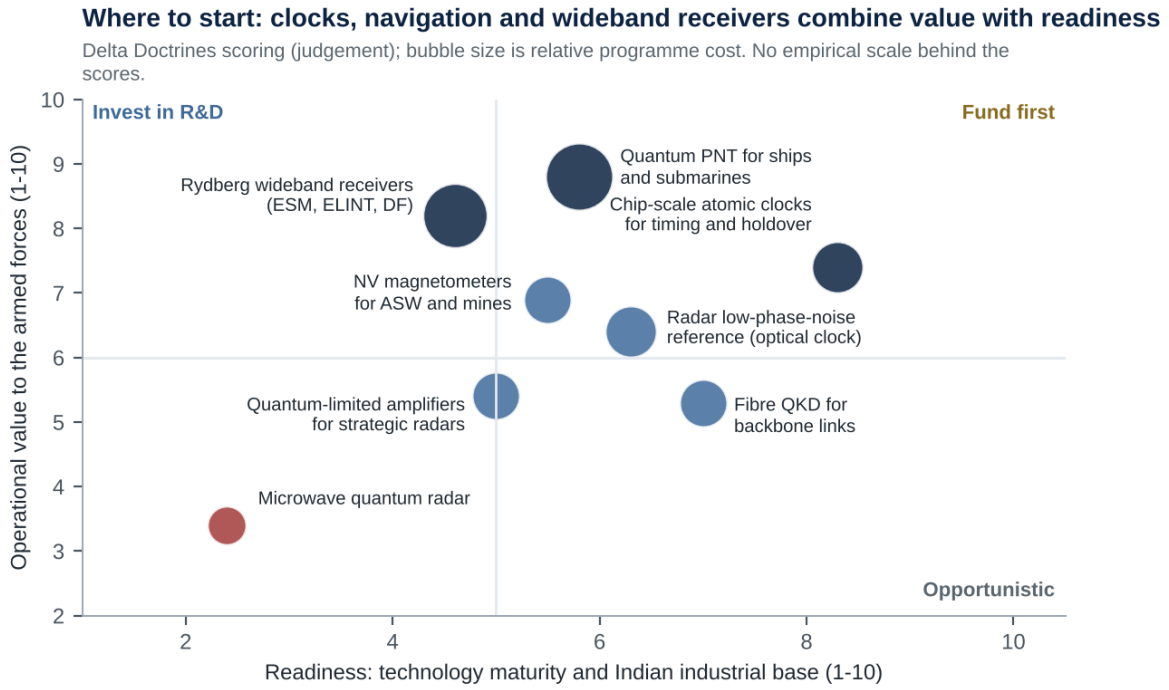


Figure 11. Priority matrix: operational value against readiness. Delta Doctrines scoring (judgement); bubble size indicates relative programme cost.

Upgrade	Applies to	Mechanism	Expected benefit	Label	Timing
Optical or chip-scale clock reference	Ground and AEW&C; radars; timing networks	Lower phase noise; holdover when GNSS fails	Better coherent integration; time sync across nodes	Judgement	2028-2031
Quantum-limited amplifiers (TWPA, JPA)	Strategic and ballistic-missile-defence radars; SIGINT receivers	Lower noise floor	Longer detection range at the same power	Judgement	2030-2034
Rydberg wideband receivers	ESM, ELINT, direction finding	Atom-based detection over wide bandwidth	Broad-band reception; calibration-free fields; immunity to some EMP	Judgement	2031-2035
Atom-interferometer IMU, quantum PNT	Ships, submarines, long-range missiles, UAVs	Low drift inertial navigation	Navigate without GNSS for hours	Judgement	2030-2036
NV and atomic magnetometers	ASW aircraft, ships, mine countermeasures	Magnetic anomaly sensing	Detect submarines and mines; MAGNAV	Judgement	2030-2036
Fibre QKD and PQC	Backbone links among data centres and HQs	Quantum key distribution plus post-quantum cryptography	Long-term confidentiality	Stated field trials [S13]	2028 onward

Why this order. Clocks and navigation score highest because they are the most ready and because other functions depend on them. Rydberg receivers score high on value but lower on readiness, so they sit in research, not procurement, until trials show field performance. Quantum-limited amplifiers matter for strategic radars but need cryogenic supply chains India does not yet have at scale. The matrix scores are not measured; they are intended to structure the debate.

4. QUANTUM SENSING

Indian quantum programmes and the communications layer

India has a national mission, a DRDO centre and a track record in QKD trials. What it lacks is a defence-specific sensor roadmap with money and field trials. The proposal adds one, sized at Rs 6,350 crore over ten years: a quarter of the programme.

Item	What is reported	Label	Source
National Quantum Mission	Outlay of Rs 6,003.65 crore over eight years; hubs in sensing, communication, computing and materials	Stated	[S12]
1,000-km quantum communication	Headline in April 2026 of a 1,000-km secure quantum communication network advance under the NQM; details not verified here	Reported (headline only)	[S13]
DRDO QKD	QKD experimental platforms at the Quantum Technology Research Centre led by the Scientific Analysis Group	Stated (PTI)	[S15]
DRDO atomic clock	Ultra-small coherent-population-trapping clock at SSPL	Stated (PTI)	[S15]
DRDO magnetometer	Optically pumped magnetometer	Stated (PTI)	[S15]
Vapour cells	Micro-fabricated alkali vapour cell characterisation (relevant to Rydberg and clock devices)	Stated (PTI)	[S15]

Quantum track line	Base-case, Rs crore	What it buys
R and D	1,300	Five sensor lines: PNT 400, RF 300, clocks 200, magnetometers 250, radar references 150 (modal, before scale factor)
Units	3,700	Field units in 2030-2036, from clocks (cheapest) to radar references (dearest per unit)
Industrial base	850	Vapour cells, photonics, magnetic and gravity maps, test ranges
Quantum-secure comms	500	QKD links and post-quantum cryptography migration
Total quantum track	6,350	26% of the programme

QUANTUM-SECURE COMMUNICATIONS

Fibre QKD is a point-to-point technology with distance limits; satellite QKD and trusted nodes extend it. For the data-centre backbone, the more important near-term step is post-quantum cryptography, which runs on ordinary hardware. The quantum comms line (Rs 500 crore) funds both. A mature QKD network is not a precondition for the AI programme.

Unit costs and counts are Delta Doctrines judgement. The model assumes 600 PNT, 400 RF, 300 clock, 100 magnetometer and 150 radar-reference units at Rs 2.5, 1.5, 1.0, 4.0 and 6.0 crore. No Indian price has been published.

5. SIZING AND COST

Sizing the cluster by workload: 2,048, 8,192 or 24,576 GPUs

The cluster is sized from the work it must do, not from a headline number. Five workloads share the estate: foundation training and adaptation; multimodal ISR inference; simulation and wargaming; chat and agent inference with redundancy; and evaluation. A spares and surge reserve covers failures and crises.

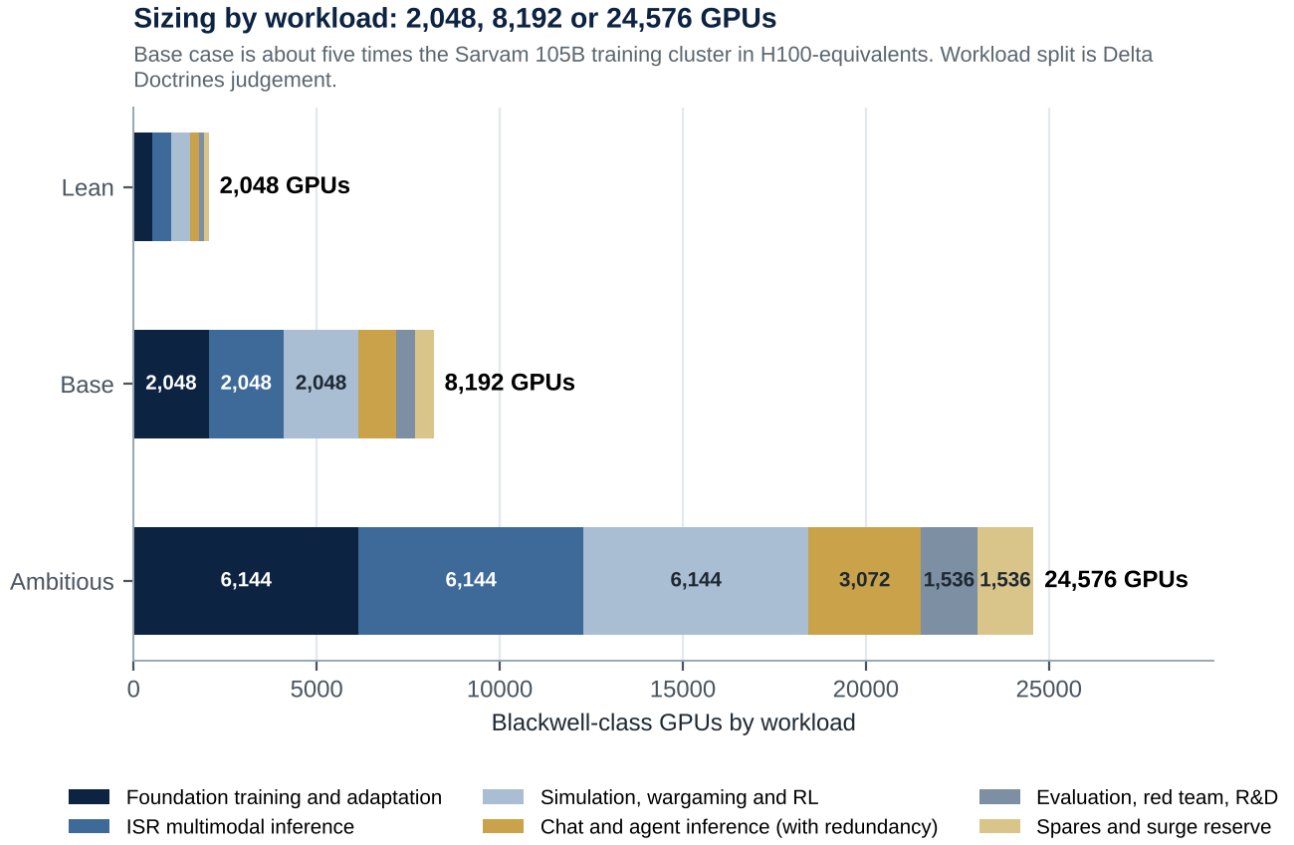


Figure 12. Sizing by workload. The base case is about five times the Sarvam 105B training cluster in H100-equivalents (a Blackwell-class GPU counted as 2.5 H100s). Workload split is a Delta Doctrines judgement.

Measure	Lean	Base	Ambitious
GPUs online at full build	2,048	8,192	24,576
IT load, MW	3.7	14.7	44.2
Facility power at meter, MW	4.8	19.2	57.5
Hardened facility cost, Rs crore	827	3,307	9,920
Edge nodes	300	600	1,500
Ten-year total (modal), Rs crore	9,844	24,631	57,451

Why 8,192. It lets the strategic tier be trained and continually adapted without queueing, runs ISR inference at theatre scale, and keeps a one-eighth reserve, while staying under 20 MW at the meter. Lean (2,048) supports adaptation of an existing model and a small operational tier but not repeated large training. Ambitious (24,576) supports frontier-scale training and heavy simulation, at about 2.3 times the cost of Base. The choice among them is a policy decision this paper informs but does not make.

GPU capex per unit in the model is Rs 56.9 lakh (USD 44,000 server price share x 1.25 system overhead x 1.15 landed cost at Rs 90 per dollar).

5. SIZING AND COST

Cost model: 18 inputs, all Delta Doctrines assumptions

The model runs a ten-year cash flow, 2027-2036, in constant 2026 rupees at Rs 90 per dollar with no escalation. It has 17 cost lines in two tracks. Seventeen inputs are drawn from triangular distributions in the Monte Carlo. Programme scope (GPUs, edge nodes, programme scales) is held fixed per scenario.

Input	Low	Mode	High	Basis
GPU server price, USD per GPU	38,000	44,000	52,000	Blackwell-class server; assumption
Network, storage, software overhead (x)	1.2	1.25	1.35	Judgement
Landed-cost factor: duty, logistics (x)	1.08	1.15	1.25	Duties and logistics; defence exemptions possible
IT kW per GPU, all-in	1.6	1.8	2.1	Judgement from vendor class
Power usage effectiveness	1.25	1.3	1.45	Air-cooled to liquid-cooled
Facility cost, Rs crore per IT MW	110	150	200	AI-optimised facility excluding GPUs; judgement
Defence hardening premium (x)	1.2	1.3	1.45	Shielding, access control, redundancy; judgement
Multi-site factor (x)	1.08	1.15	1.25	Three sites; judgement
Power tariff, Rs per kWh	4	5	9	Industrial to captive; judgement
Facility O and M, share of capex a year	0.035	0.05	0.07	Judgement
GPU support, share of capex a year	0.04	0.05	0.08	Judgement
Refresh GPU price vs original (x)	0.45	0.6	0.8	Half the fleet refreshed in 2033-2035
Edge node, Rs crore each	1.2	2	3.5	Ruggedised unit; judgement
Integration scale (x)	0.7	1	1.6	Scale on Rs 1,200 crore
Model programme scale (x)	0.75	1	1.5	Scale on Rs 220 crore a year
Quantum R and D scale (x)	0.8	1	1.5	Scale on Rs 1,300 crore
Quantum unit cost scale (x)	0.7	1	1.6	Scale on Rs 3,700 crore
Quantum industrial base scale (x)	0.8	1	1.4	Scale on Rs 850 crore

Fixed anchors (Rs crore, base case, ten years). Integration and data fabric 1,200; data programme 600; model programme 220 a year at steady state; evaluation 40 a year; talent academy 27 a year; bridge compute 120; QKD and comms 500. Quantum R and D 1,300; quantum units 3,700; quantum industrial base 850. These are programme-sizing judgements rather than quotes.

WHAT IS NOT IN THE MODEL

Land, civil works outside the data-centre shell, classified personnel costs, salaries of Service staff, inflation, currency movements, the cost of the platforms that carry quantum sensors, and any export-control or sanctions premium. The Monte Carlo covers unit-cost uncertainty only, not scope, schedule or policy risk. A risk-adjusted budget would add a contingency (Chapter 6).

5. SIZING AND COST

The AI track: Rs 18,281 crore over ten years, led by GPUs and the hardened facility

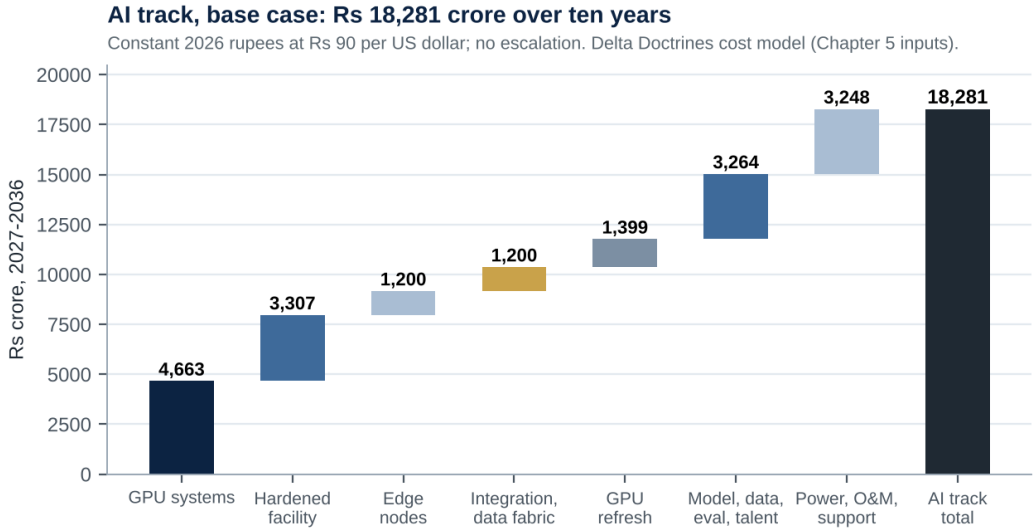


Figure 13. AI track, base case, by component. Constant 2026 rupees at Rs 90 per dollar; no escalation.

AI-track line	Rs crore	Share of AI track	What it covers
GPU systems	4,663	25.5%	Servers, network, storage; 8,192 GPUs in three tranches
GPU refresh	1,399	7.7%	Half the fleet in 2033-2035 at 60% of original price
Facility (hardened)	3,307	18.1%	Three sites, 14.7 MW IT
Edge nodes	1,200	6.6%	600 ruggedised nodes, 2029-2033
Integration & data fabric	1,200	6.6%	Gateways, guards, standards
Data programme	600	3.3%	Labelling, lineage, synthetic data
Bridge compute	120	0.7%	Civil IndiaAI compute for unclassified work
Power	426	2.3%	Energy at 70% load
Facility O&M;	1,248	6.8%	5% of capex a year
GPU support	1,574	8.6%	5% of capex a year
Model programme	1,947	10.7%	Training, tuning, distillation, labs
Evaluation & red team	340	1.9%	Independent body
Talent academy	256	1.4%	Training and retention
AI track total	18,281	100%	74% of the programme

Reading the numbers. Compute hardware, its refresh and its support is Rs 7,636 crore, or 42% of the AI track. The hardened facility is Rs 3,307 crore (18%). Model, data, evaluation and talent together are Rs 3,144 crore (17%). Recurring costs in the later years are dominated by support and power; capex falls away after 2031.

> CROSS-CHECK
> A reported commercial benchmark helps scale the order of magnitude: Yotta's planned 20,736-GPU Blackwell Ultra cluster is reported at over USD 2 billion [S11], about USD 97,000 per GPU for the whole project. The base case here spends about Rs 7,970 crore (USD 0.89 billion) on GPU systems and the hardened facility for 8,192 GPUs, about USD 108,100 per GPU. The comparison is rough because scope differs (the commercial figure may include land, power and other items), but the two are in the same range.

5. SIZING AND COST

Annual spend peaks in 2028-2031, then settles near Rs 1,500-2,300 crore

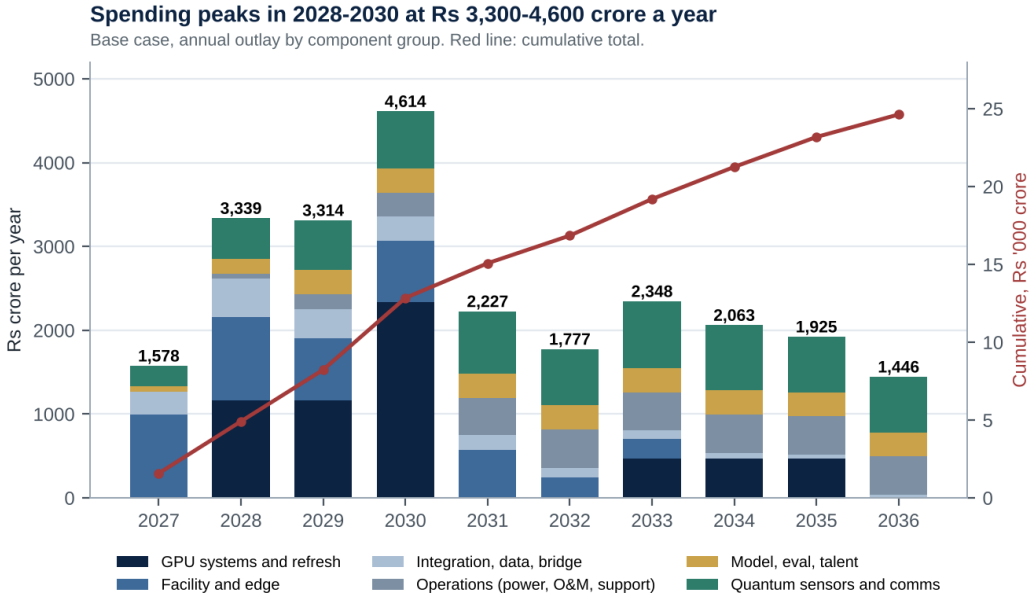


Figure 14. Base-case annual outlay by component group. Red line: cumulative total.

Year	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036
AI track	1,337	2,850	2,716	3,927	1,478	1,104	1,547	1,283	1,259	780
Quantum track	241	489	598	687	749	673	801	780	666	666
Total	1,578	3,339	3,314	4,614	2,227	1,777	2,348	2,063	1,925	1,446
Cumulative	1,578	4,916	8,231	12,845	15,072	16,850	19,197	21,260	23,184	24,631
% of FY27 capital outlay	0.7%	1.5%	1.5%	2.1%	1.0%	0.8%	1.1%	0.9%	0.9%	0.7%

Spending peaks in 2030 at Rs 4,614 crore, when the last GPU tranche, the last facility tranche, edge nodes and the first quantum unit purchases overlap. Outlay settles at Rs 1,400-2,300 crore a year by 2034-2036, dominated by support, power, the model programme and quantum unit purchases. Refresh in 2033-2035 adds about Rs 466 crore a year.

Window	Spend, Rs crore	Share of ten-year total
2027-2028 (set-up)	4,916	20%
2027-2031 (build)	15,072	61%
2032-2036 (operate and refresh)	9,558	39%

Constant 2026 rupees. Real budgets will rise with inflation and could fall if the rupee strengthens; the figures should be indexed when a budget is set.

5. SIZING AND COST

Power: 19 MW for the base case, and the tariff matters more than the scenario

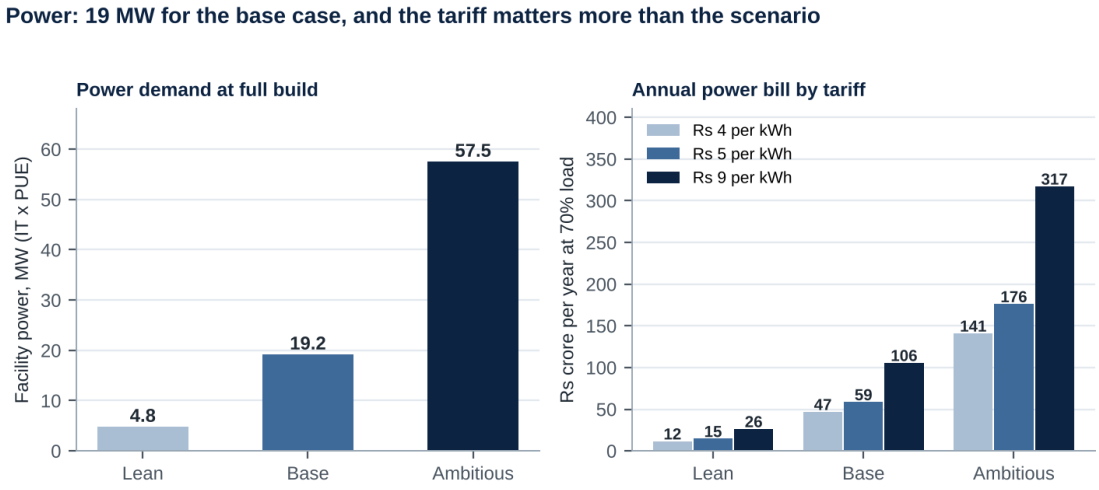


Figure 15. Facility power by scenario and annual power bill by tariff at 70% average load.

Scenario	Facility MW	Energy, GWh a year	Rs 4 per kWh	Rs 5 per kWh	Rs 9 per kWh
Lean	4.8	29	Rs 12 crore	Rs 15 crore	Rs 26 crore
Base	19.2	118	Rs 47 crore	Rs 59 crore	Rs 106 crore
Ambitious	57.5	353	Rs 141 crore	Rs 176 crore	Rs 317 crore

The base-case cluster draws about 14.7 MW of IT load and 19.2 MW at the meter (PUE 1.30). At 70% average load that is about 118 GWh a year, or Rs 59 crore at Rs 5 per kWh. Power is therefore only 1.7% of the ten-year total, because GPU hardware dominates. But siting matters for another reason: grid connection, backup generation and cooling water decide how fast the sites can be built, and the 2027-2028 facility tranche (Rs 1,984 crore) cannot start without them.

SITING CRITERIA (DELTA DOCTRINES JUDGEMENT)

Three sites in different seismic and flood zones; two independent grid feeds each; strategic depth from borders and coast; proximity to fibre trunk routes; cooling water or a closed-loop design; and the ability to accept accredited construction crews. Each site is sized for about 5-7 MW at the meter in the base case. Existing sites such as Yotta's 60 MW Noida campus [S11] are civil and are unsuitable for classified workloads, but their design and supply chain are useful references.

Renewables. Contracted solar or wind can cut the energy cost and the carbon footprint, but cannot replace on-site generation for an air-gapped estate. The model assumes grid power at the stated tariffs.

5. SIZING AND COST

Own versus rent: owning is the price of sovereignty

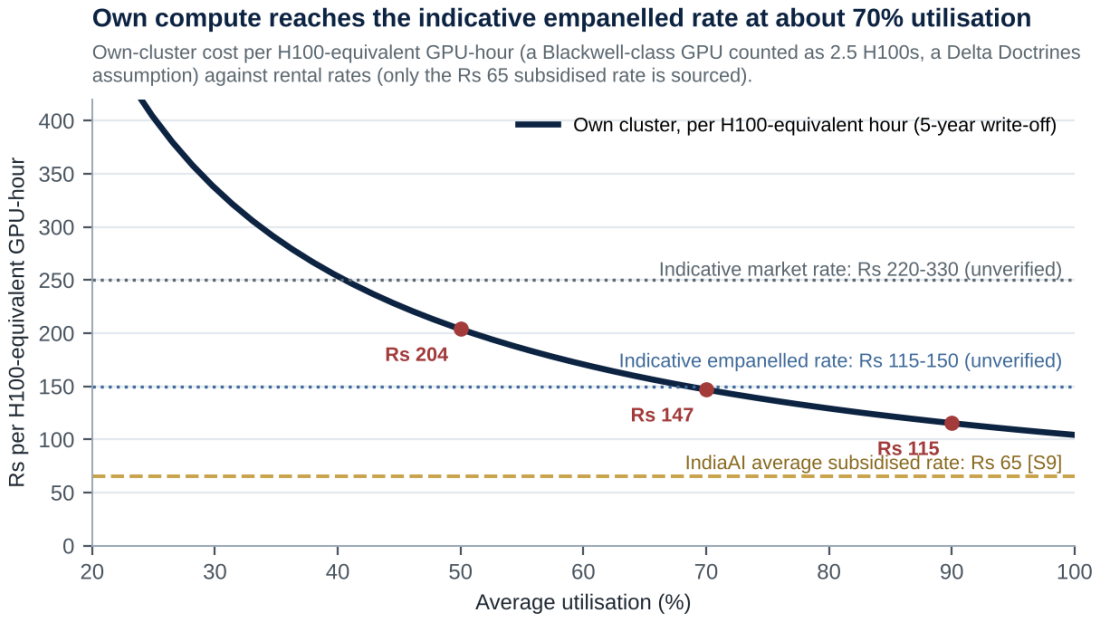


Figure 16. Own-cluster cost per H100-equivalent GPU-hour against rental rates (only the Rs 65 subsidised rate is sourced). A Blackwell-class GPU is counted as 2.5 H100s, a Delta Doctrines assumption.

Average utilisation	Own cost, Rs per H100-equivalent hour	Versus subsidised Rs 65 [S9]
50%	204	3.1 times
70%	147	2.3 times
90%	115	1.8 times

Own compute reaches the indicative empanelled rate of about Rs 115-150 (unverified) at around 70% utilisation, and exceeds the Rs 65 average subsidised rate at any utilisation [S9]. The result is intuitive: civil capacity is subsidised and pooled, while a dedicated cluster carries idle time, hardening, three sites and a security team.

DECISION RULE

Rent for anything unclassified; own for everything else. If more than about 70% of the expected workload could run on civil capacity, the owned cluster should be smaller than the base case. The sizing in sizing figure assumes the opposite: that the main workloads (ISR fusion, operational planning, classified training) cannot leave an accredited estate.

Method. Cost per hour = (GPU capex per H100-equivalent, plus a share of facility, power and support over five years) divided by hours at the stated utilisation. It uses base-case inputs; the five-year write-off is a modelling choice, not an accounting rule. Rates above the subsidised figure are indicative and unverified.

5. SIZING AND COST

The quantum track: Rs 6,350 crore, with research first and production from 2030

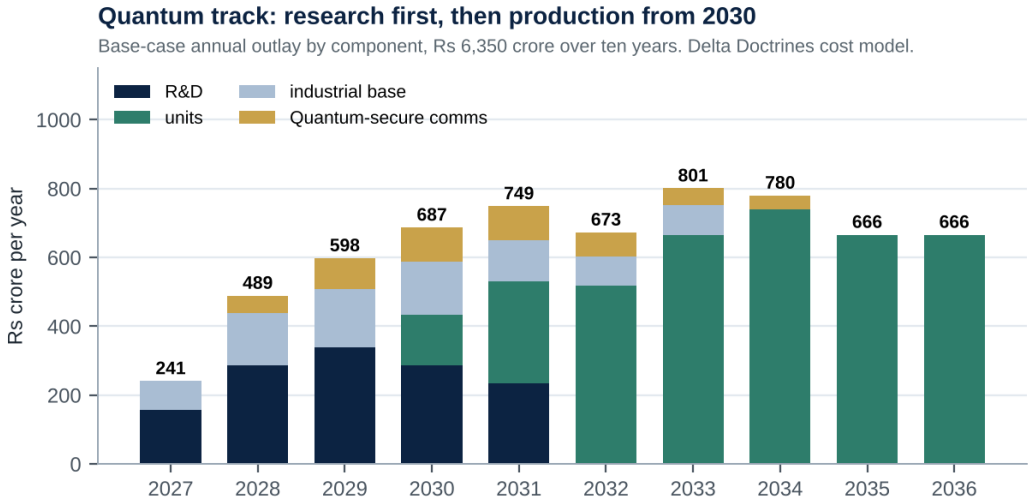


Figure 17. Quantum track, base-case annual outlay by component.

Year	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036
R&d;	156	286	338	286	234	0	0	0	0	0
Units	0	0	0	148	296	518	666	740	666	666
Industrial base	85	153	170	153	119	85	85	0	0	0
Quantum secure comms	0	50	90	100	100	70	50	40	0	0
Total	241	489	598	687	749	673	801	780	666	666

Sensor line	R and D, Rs crore	Units (number)	Unit cost, Rs crore	Unit spend, Rs crore
Quantum PNT	400	600	2.5	1,500
Rydberg RF receivers	300	400	1.5	600
Atomic clocks	200	300	1.0	300
Magnetometers	250	100	4.0	400
Radar references and amplifiers	150	150	6.0	900
Total (modal, base scale)	1,300	1,550	-	3,700

Research and the industrial base are front-loaded in 2027-2031 (Rs 2,150 crore); unit purchases start in 2030 and build to Rs 740 crore a year. The unit counts and prices are Delta Doctrines judgement and carry the single largest swing in the whole model (see the tornado chart). If trials in 2029 fail, purchases can be cut or deferred without stranding the AI track, because the two tracks are financed separately.

The quantum track is deliberately a sensor-upgrade programme, not a quantum-computing programme. It does not fund quantum computers, which are outside the scope of the request.

5. SIZING AND COST

Monte Carlo: the base case is a floor, not a midpoint

Monte Carlo on unit costs: 10-year base-case total of Rs 24,900-28,300 crore (P10-P90)

100,000 draws from triangular distributions on 17 cost inputs; programme scope held at the base case. Delta Doctrines model.

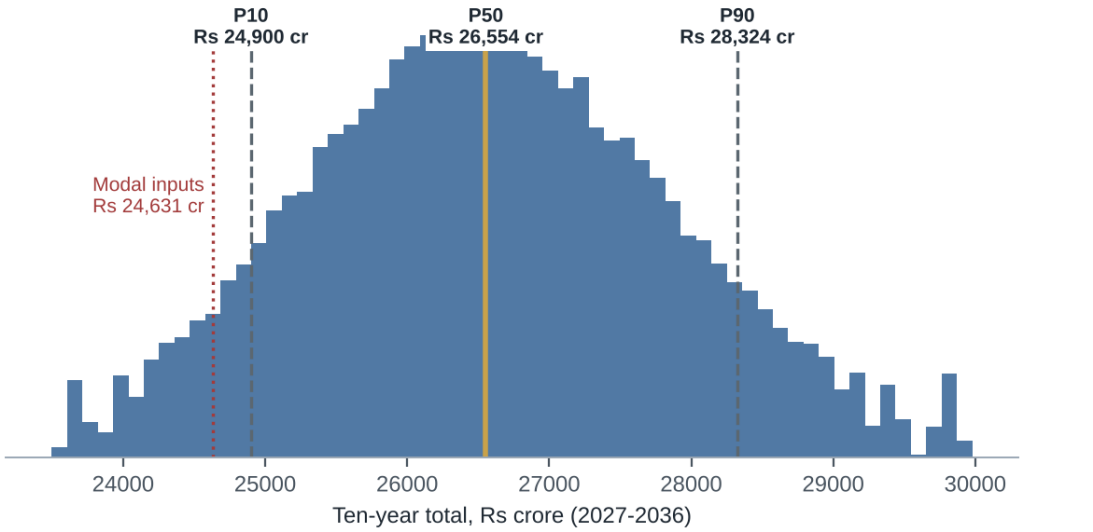


Figure 18. Distribution of the ten-year base-case total from 100,000 draws on 17 cost inputs. Programme scope is held at the base case.

Statistic	Total, Rs crore	AI track	Quantum track
Modal inputs (base case)	24,631	18,281	6,350
P5	24,453	17,896	5,798
P10	24,900	18,269	5,993
P50 (median)	26,554	19,651	6,843
Mean	26,590	19,684	6,906
P90	28,324	21,138	7,920
P95	28,832	21,572	8,185

The modal case of Rs 24,631 crore lies just below the 5th percentile (Rs 24,453 crore). That is not an error. Most inputs have right-skewed ranges (the high end is further from the mode than the low end), so the average draw costs more than the modal draw. For budgeting the paper recommends the median, Rs 26,554 crore (USD 2.95 billion), and a P90 of Rs 28,324 crore for the contingency envelope.

WHAT THE RANGE DOES AND DOES NOT MEAN

The P10-P90 spread of about Rs 3,400 crore (14% of the median) reflects only unit-cost uncertainty with a fixed scope. It does not include schedule slip, scope growth, a sanctions or export-control premium, exchange-rate moves or inflation. Real programmes of this type overrun by more than the unit-cost range suggests. A realistic contingency is therefore a policy choice, not an output of this model.

First five years (2027-2031): median Rs 16,065 crore (P10 to P90: Rs 15,075 to 17,114 crore). Core capex (GPUs, hardened facility, edge nodes, integration): median Rs 10,994 crore.

5. SIZING AND COST

What moves the cost: quantum unit prices, facility cost per MW and GPU price

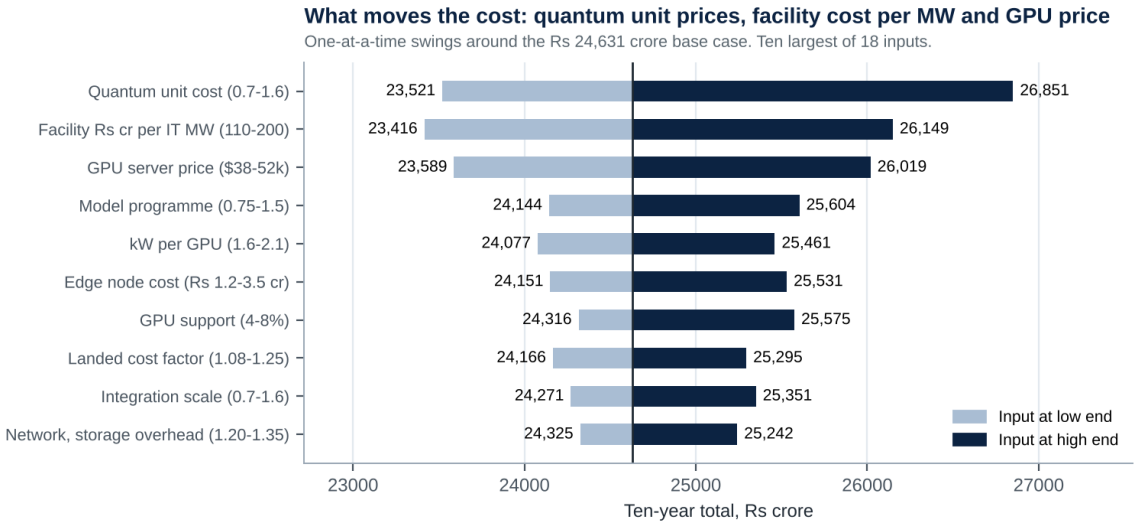


Figure 19. One-at-a-time swings around the base case. Ten largest of 18 inputs.

Input (low to high end)	Total at low end	Total at high end	Swing, Rs crore
Quantum unit cost scale (x) (0.7 to 1.6)	23,521	26,851	3,330
Facility cost, Rs crore per IT MW (110 to 200)	23,416	26,149	2,733
GPU server price, USD per GPU (38,000 to 52,000)	23,589	26,019	2,430
Model programme scale (x) (0.75 to 1.5)	24,144	25,604	1,460
IT kW per GPU, all-in (1.6 to 2.1)	24,077	25,461	1,384
Edge node, Rs crore each (1.2 to 3.5)	24,151	25,531	1,380
GPU support, share of capex a year (0.04 to 0.08)	24,316	25,575	1,259
Landed-cost factor: duty, logistics (x) (1.08 to 1.25)	24,166	25,295	1,129
Integration scale (x) (0.7 to 1.6)	24,271	25,351	1,080
Network, storage, software overhead (x) (1.2 to 1.35)	24,325	25,242	916

Three inputs account for most of the uncertainty: the quantum unit-cost scale (Rs 3,330 crore swing), the facility cost per MW (Rs 2,733 crore) and the GPU server price (Rs 2,430 crore). The first two are procurement-negotiable: unit costs fall with volume, and facility cost falls if sites reuse existing hardened buildings. The GPU price depends on the supplier market.

WHERE TO LOOK FOR SAVINGS

Reusing hardened buildings could cut facility cost per MW; a multi-year volume agreement for GPUs could lower the server price; and phasing quantum unit purchases after field trials avoids buying early at high unit cost. Conversely, power tariff and PUE matter little to the total. Effort spent negotiating electricity tariffs will return less than effort on hardened-building reuse.

A tornado shows each input alone. The Monte Carlo (previous page) shows them together; because inputs are drawn independently, it may understate the chance of correlated overruns, for example a GPU price shock that also lengthens delivery.

5. SIZING AND COST

Three scales: Rs 9,800, 24,600 and 57,500 crore over ten years

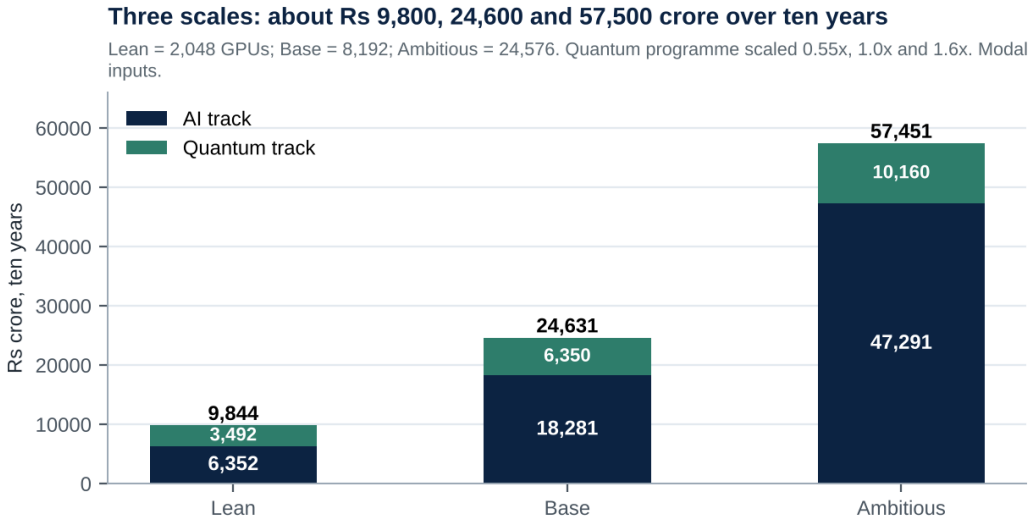


Figure 20. Three scenarios at modal inputs. Lean = 2,048 GPUs; Base = 8,192; Ambitious = 24,576. Quantum programme scaled 0.55x, 1.0x and 1.6x.

Measure	Lean	Base	Ambitious
GPUs	2,048	8,192	24,576
Edge nodes	300	600	1,500
Ten-year total, modal, Rs crore	9,844	24,631	57,451
Monte Carlo median, Rs crore	10,608	26,554	61,910
P10 to P90, Rs crore	9,922 to 11,311	24,900 to 28,324	57,588 to 66,305
AI track, modal, Rs crore	6,352	18,281	47,291
Quantum track, modal, Rs crore	3,492	6,350	10,160
Quantum scale	0.55x	1.0x	1.6x
Capability	Adapt an open model; small operational tier; pilot sensors	Strategic, operational and tactical tiers; theatre nodes; four sensor lines	Frontier-scale training; heavy simulation; wide fielding

The scenarios are not a ladder to climb automatically. Lean is a legitimate end-state for a Ministry that wants to adapt an existing model and run pilots. Base is the paper's recommendation because it supports the three-tier family and theatre-level operations. Ambitious is only justified if the Ministry decides to train frontier-scale models in India for defence, which would compete with civil missions for chips and power.

RECOMMENDED PATH

Commit to the Base design, release money for the Lean tranche first (2,048 GPUs by 2028), and make the step to 4,096 and 8,192 GPUs conditional on the gates in Chapter 6. The Lean case is therefore the first stage of the Base case, not an alternative to it, and the decision to go beyond Base can be deferred to 2031.

The Ambitious case Monte Carlo is run on a 5% sample of draws (5,000), so its percentiles are less precise than the Base case (100,000 draws).

6. ROADMAP AND DELIVERY

Roadmap 2027-2036: two tracks, three phases, three gates

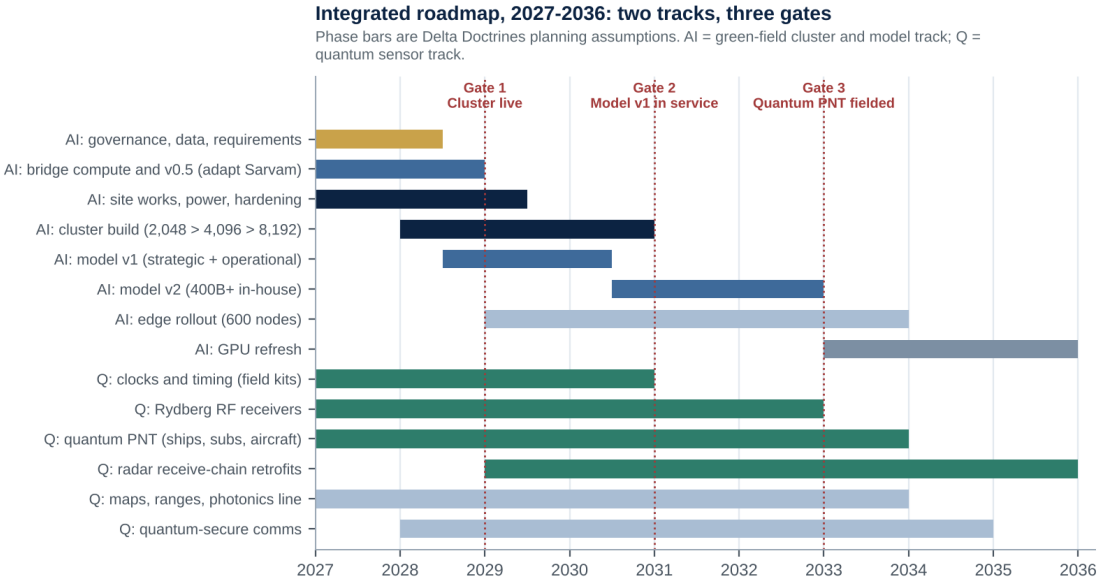


Figure 21. Integrated roadmap. Phase bars are Delta Doctrines planning assumptions: AI = cluster and model track; Q = quantum sensor track.

Phase	Years	Aim	Spend, Rs crore	Share	Exit gate
1. Foundation	2027-2028	Mission office; data charter; bridge compute; adapt an open model; start site works; quantum R and D and clock kits	4,916	20%	Gate 1: first 2,048 GPUs live; accreditation scheme approved
2. Build	2029-2031	Cluster to 8,192 GPUs; model v1 (strategic and operational); edge rollout begins; quantum PNT and RF trials	10,156	41%	Gate 2: model v1 in service after independent release test
3. Scale and sustain	2032-2036	Model v2; 600 edge nodes; quantum units fielded; GPU refresh; steady operations	9,558	39%	Gate 3: quantum PNT fielded on first ships and aircraft

Gates release money. Each gate is a decision point at which the next tranche is released only if defined tests have been met. The model's funding profile follows the gates: Rs 4,916 crore is committed for Phase 1, and the rest only after Gate 1. This keeps the Ministry's downside bounded: if the programme is stopped at Gate 1, the loss is about 20% of the total.

WHY A TEN-YEAR VIEW

GPU generations turn over in three to five years; sensors take five to ten years from lab to field; and a classified data programme takes years to mature. Planning to 2036 includes a mid-life GPU refresh (2033-2035) and the first quantum unit purchases (2030), which a five-year plan would miss.

6. ROADMAP AND DELIVERY

Phase 1 (2027-2028): foundations, bridge compute and first pilots

Phase 1 spends Rs 4,916 crore, 20% of the ten-year total, mostly on facility works, the first 2,048 GPUs and quantum R and D. It is a pilot in every sense: it tests whether data owners will share, whether the sites can be built on time, and whether an adapted open model is useful. The programme is designed so that Phase 1 yields value even if it stops.

Workstream	2027 actions	2028 actions	Deliverable by end-2028	Rs crore, 2027-28
Governance	Set up mission office; appoint head; draft charter with HQ IDS	Approve accreditation scheme; data-sharing charters signed	Charter; accreditation rules	Within programme
Data	Inventory data sources; start labelling pilots	Stand up data fabric v1; first cross-domain guard	Catalogue of Service datasets; guard in test	612
Compute	Site selection; power and hardening design; bridge compute contract	Build first site; install first 2,048 GPUs	One site live (Gate 1)	3,270
Models	Adapt Sarvam open weights on unclassified data using civil compute	Continued pre-training on first classified tier; v0.5 pilot	Pilot staff assistant; first military test set	187
Evaluation	Form body; recruit; draft test sets	First independent test of v0.5	Release certificate procedure	20
Talent	Launch academy; secondments from labs	First cohort of 200	Core team in place	40
Quantum	Award R and D lines; clock kit specification	Clock field kits to trial; Rydberg lab prototypes; start photonics line	Trial reports; test range	730

THE BRIDGE

Until the first site is live, unclassified adaptation of the open model can run on IndiaAI civil compute at the subsidised average of Rs 65 per GPU-hour [S9], budgeted at Rs 120 crore. This keeps model work moving while sites are built. Classified data stays out of the bridge.

First deliverable the Services will see: a staff assistant for unclassified documents, and a trial of quantum clock kits on one air or naval platform. Early, modest, visible results are more valuable than a late, large one.

6. ROADMAP AND DELIVERY

Phases 2 and 3 (2029–2036): scale, field and sustain

Phase	Workstream	Key actions	Gate test
2 (2029-2031)	Compute	Cluster to 4,096 then 8,192 GPUs; three sites; theatre nodes at first commands	Utilisation above 60%; uptime target met
	Models	v1 strategic and operational tiers; multimodal ISR adapters; distillation to tactical tier	Independent test beats previous version; hallucination under threshold
	Edge	First 240 nodes with Army, Navy and Air Force pilots	Field exercise passed
	Quantum	Clock kits to fleets; Rydberg and quantum-PNT trials; QKD links between sites	Trial reports; go or no-go for unit purchases
3 (2032-2036)	Compute	GPU refresh of half the fleet in 2033-2035	Cost per GPU-hour within model
	Models	v2 (400B-class) if justified; continuous evaluation	Sunset or retrain decisions annually
	Edge	Reach 600 nodes by 2033	Coverage target met
	Quantum	Fielded PNT on ships, submarines, aircraft; Rydberg receivers in ESM; amplifiers in strategic radars	Gate 3: fielded and in service

The KPIs below are proposals. They are chosen to be measurable and to test the three claims the programme makes: that decisions are faster, that the models are trustworthy and that navigation survives jamming.

KPI	Target by 2031	Target by 2036	Label
Cluster uptime	99% (core)	99.5%	Judgement
Share of staff tasks with approved model support	20% in pilot commands	60%	Judgement
Decision-chain time, targeting (illustrative budget)	-25%	-45%	Judgement
Hallucination rate on military test set	Under set threshold	Under tighter threshold	Judgement
GNSS-denied navigation hold time on fielded platforms	Trial kits	Hours at set drift	Judgement
Share of workloads on sovereign stack	50%	80%	Judgement

WHAT A STOP LOOKS LIKE

If Gate 2 fails (the model does not pass independent testing), the programme should hold the GPU tranche at 4,096, retrain or switch lab, and re-test within 12 months. It should not buy more hardware to compensate for a model that is not ready. The cost model shows the GPU tranche after 2029 is Rs 2,332 crore, so the stop option has real value.

6. ROADMAP AND DELIVERY

Governance: one mission, one accountable head, independent assurance

Governance is the point on which this kind of programme most often fails. The proposal is to concentrate accountability in one mission office with a named head, give data policy to a council the Services trust, and place assurance outside the chain that builds and uses the models. The bodies below are proposals by Delta Doctrines; none exists as described.

Body	Role	Members	Decision rights
Defence AI and Quantum Mission board	Sets strategy; releases tranches at gates	Defence Secretary (chair), CDS, three Service chiefs, Secretary DDR and D, MeitY, DST	Budget, gates, scope
Mission office	Programme management; schedule; supplier management	Mission head (joint secretary-level or above) and small team	Day-to-day delivery
Data council	Data charters; cross-domain release rules	Data owners from Services, DRDO, intelligence bodies; HQ IDS chairs	Who may share what
Evaluation and red-team body	Tests and certifies every model release	Independent experts, academia and technical staff	Can block a release
Technical design authority	Standards, interfaces, accreditation	DRDO, C-DAC-type institutions, Service technical staff	Standards and acceptance
Industry and lab consortium	Builds under contract	Sarvam [S1] and other empanelled labs [S9]; integrators; quantum start-ups	No decision rights
External review panel	Annual review of safety, ethics, value for money	Retired judges, scientists, auditors	Advisory, reports to the Defence Minister

THREE STRUCTURAL CHOICES

First, a joint, tri-service owner (HQ IDS and the mission office), not a Service owner, so that no one Service controls the common data fabric. Second, an evaluation body that reports to the board, not to the builders. Third, open interfaces and a second-source rule for models and sensors, so that no vendor, including Sarvam, can lock in the Ministry. The reported MoU has HQ IDS as the counterparty [S1], which is consistent with the first choice.

Relationship to existing bodies. The Defence AI Project Agency (DAIPA), AI Centre of Excellence and Service AI cells would remain as user and application bodies. The mission would provide the shared platform. The reported Rs 300 crore Centre of Excellence [S3] could be the application layer on top of the platform; this paper does not cost it.

6. ROADMAP AND DELIVERY

Funding and affordability: about 1.1% of one year's capital outlay a year

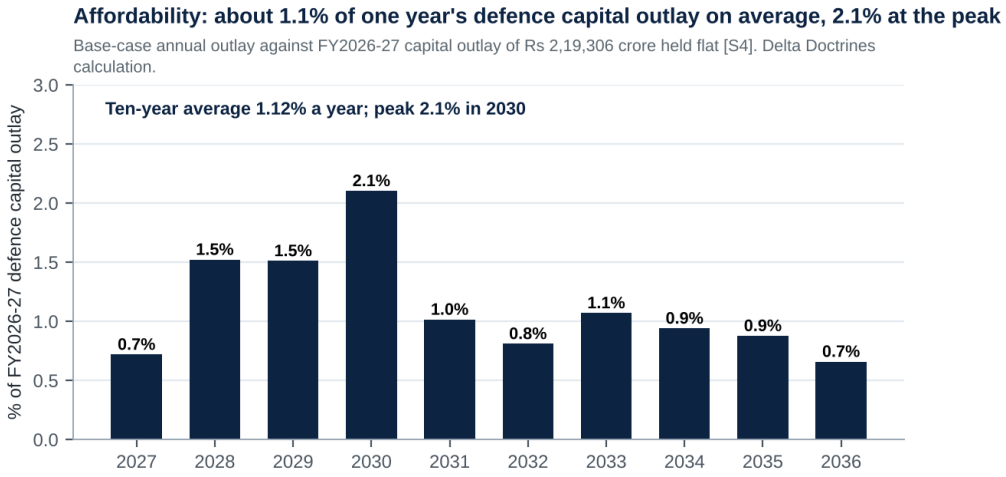


Figure 22. Base-case annual outlay as a share of FY2026-27 defence capital outlay (Rs 2,19,306 crore, held flat) [S4].

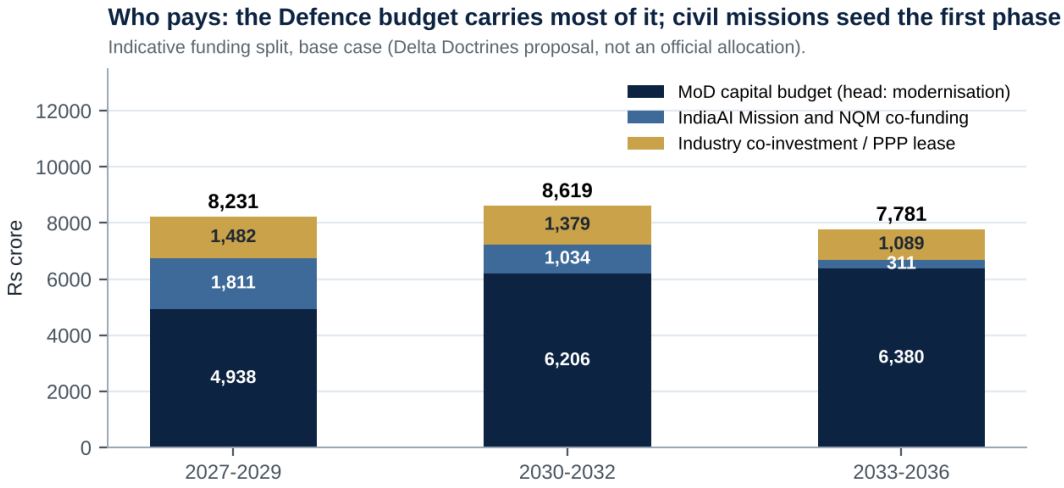


Figure 23. Indicative funding split (Delta Doctrines proposal, not an official allocation).

Window	Total, Rs crore	MoD capital budget	IndiaAI and NQM co-funding	Industry co-investment or lease
2027-2029	8,231	4,938 (60%)	1,811 (22%)	1,482 (18%)
2030-2032	8,619	6,206 (72%)	1,034 (12%)	1,379 (16%)
2033-2036	7,781	6,380 (82%)	311 (4%)	1,089 (14%)
Ten years	24,631	17,525	3,156	3,950

The programme averages Rs 2,463 crore a year, 1.1% of capital outlay, and peaks at 2.1% in 2030. Allocated to the Defence capital head it would be absorbed within a single year's growth: capital outlay rose from Rs 1.80 lakh crore to Rs 2.19 lakh crore between FY2025-26 and FY2026-27 [S4], an increase of about Rs 39,000 crore. The funding split is a proposal: civil missions (IndiaAI, NQM) seed the first phase, and industry co-invests through leases of cluster capacity or sensor production lines.

> CONTINGENCY
> A P90 envelope of Rs 28,324 crore is Rs 3,693 crore above the modal total. Adding a further 20% for scope growth and delay (a Delta Doctrines judgement, not modelled) gives an authorisation ceiling of about Rs 33,989 crore. Releases would still be gate-based.

6. ROADMAP AND DELIVERY

Risks: delay and supply are likeliest; data and model compromise are most severe

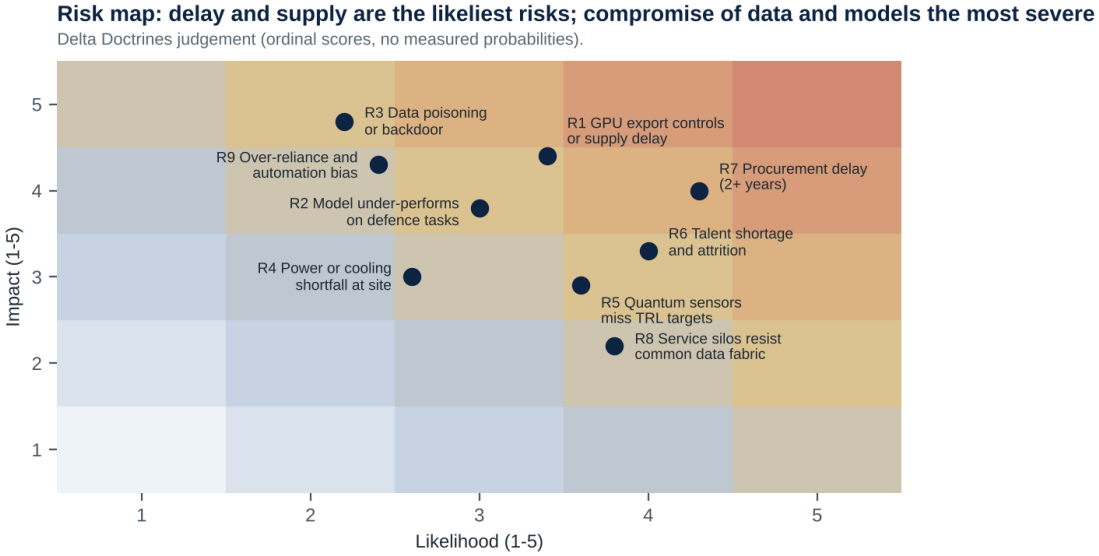


Figure 24. Risk map: likelihood against impact (1-5). Delta Doctrines judgement; ordinal scores, no measured probabilities.

Risk	L	I	Mitigation	Owner
R1 GPU export controls or supply delay	3.4	4.4	Multi-year supply agreements; dual vendors; buy tranche 1 early; domestic accelerator R and D	Mission office
R2 Model under-performs on defence tasks	3.0	4.0	Military test sets; second-source lab; gate before hardware	Evaluation body
R3 Data poisoning or backdoor	2.2	4.8	Supply-chain scanning; signed weights; air gap; red team	Security authority
R4 Power or cooling shortfall	2.6	3.4	Dual feeds; on-site generation; early works	Mission office
R5 Quantum sensors miss readiness targets	3.6	3.0	Staged buys after trials; parallel lines	DRDO
R6 Talent shortage and attrition	4.0	3.2	Academy; secondments; retention pay; foreign training	Mission office
R7 Procurement delay of two years or more	4.3	3.9	Pre-approved tranche; empanelled suppliers; fast-track rules	MoD
R8 Service silos resist a common data fabric	3.8	2.6	HQ IDS charter; incentives; early wins	HQ IDS
R9 Over-reliance and automation bias	2.6	4.2	Training; calibrated outputs; human authorisation; drills	Services

L = likelihood, I = impact on a 1-5 scale. The scores are for structuring discussion and have no empirical basis. They would be refreshed by the mission board each year.

THE RISK THAT IS NOT ON THE MAP

The largest programme risk is that the work is split across Services, labs and ministries into incompatible pieces: three stacks, three data stores, three models. That is a governance failure, not a technical one; Recommendation 1 and the structural choices in the governance section are the mitigation.

6. ROADMAP AND DELIVERY

Procurement, partners and the industrial base

A sovereign programme does not mean building everything in India. It means holding the weights, the data, the evaluation sets and the design authority, and buying components from the widest vetted set of suppliers. The proposal draws a line between what must be sovereign (data, weights, classified training, evaluation) and what can be sourced (chips, servers, test equipment).

Area	Proposal	Why	Label
Hardware	Buy GPU servers through a government-to-government or vetted-supplier route; stagger tranches; reserve an option on refresh	Export controls and supply risk; learning curve on price	Judgement
Facilities	Reuse hardened buildings where possible; build under a single design authority	Facility cost per MW is the second largest swing factor	Model
Models	Prime contract with Sarvam [S1]; second-source with another empanelled lab [S9]; Ministry owns weights	Avoid lock-in; maintain competition	Judgement
Data	Services own data; mission provides tools; HQ IDS arbitrates	Data sharing is the binding constraint	Judgement
Quantum	Fund through DRDO QTRC [S15] and NQM hubs [S12]; open competitions for clocks, PNT, Rydberg; test range	Build a supplier base before production	Judgement
Industrial base	Rs 850 crore for vapour cells, photonics, maps and ranges	Vapour cells and photonics are the shared bottleneck for clocks, magnetometers and Rydberg devices [S15]	Model
Contracting	Fixed-price for hardware; cost-plus with milestones for R and D; incentives tied to evaluation results	Match risk to contract type	Judgement
Foreign partners	Open to friendly countries for components and joint tests [S2]; none for classified model training	The Defence Secretary did not rule out foreign partnerships [S2]	Stated (reported)

CHIPS ARE THE ONE DEPENDENCY THAT CANNOT BE FIXED BY 2036

No Indian GPU is expected within the programme's horizon. The risk is managed by early purchase, supplier diversity and a plan to run on smaller or older clusters if supply is cut: the Lean case (2,048 GPUs) is a fallback that still supports adaptation of an open model. This is a Delta Doctrines judgement and the supply position should be reviewed by the mission board each year.

Export-control premium. The model contains no sanctions or export-control premium. If one applies, it would act on the GPU price (the largest hardware swing in the tornado) and on delivery time.

6. ROADMAP AND DELIVERY

The first 100 days

The reported announcements of October 2026 [S1][S2] mean that the first 100 days should begin by establishing what has already been committed, and by what authority. The plan below is designed to avoid a second parallel programme.

Day	Action	Owner	Output
1-15	Cabinet-level or Defence Minister's approval of mission and head; announce HQ IDS as owner	Defence Minister	Mandate
1-30	Confirm status of the reported MoD-IndiaAI layer [S2], the Sarvam MoU [S1] and the Rs 300 crore centre [S3]; get sizes and costs	Mission head	Baseline
15-45	Draft data charter; identify first five datasets from three Services	HQ IDS	Charter v1
15-60	Commission accreditation scheme for AI workloads and cross-domain guards	Technical design authority	Scheme draft
30-60	Select sites; start power and hardening design	Mission office, Engineers	Site shortlist
30-75	Set up evaluation body; recruit; write first military test set	Mission board	Body in place
45-90	Contract bridge compute on IndiaAI; start adaptation of open model on unclassified data	Mission office, Sarvam	Model v0.1
45-90	Quantum: call for proposals for clock kits, PNT, Rydberg; confirm test range	DRDO	Awards
60-100	Gate 0 review: cost baseline, risk register, first-tranche release	Mission board	Go or adjust

DECISIONS NEEDED FROM THE MINISTRY IN THE FIRST 30 DAYS

(1) Single owner and head. (2) Whether the reported layer and MoU are the first phase of this programme or separate. (3) The scale to plan for: Lean, Base or Ambitious. (4) Whether the evaluation body reports outside the Services. (5) Whether civil bridge compute may be used for unclassified adaptation. Each decision is a policy choice with a cost consequence shown in Chapter 5.

What would change this plan. If the 23 October layer already includes a large procured cluster [S2], the compute line could shrink; if the MoU funding is large and Service-specific, the governance proposal becomes more important than the sizing. The paper's numbers are the stack's independent price, not its reported price, which has not been published.

ANNEX

Limits of the evidence, and what to check before relying on a number

This paper is open-source analysis. It contains no classified material and no inside knowledge of Ministry plans. The table lists the main limits and what each one affects.

Area	Limitation	Effect on conclusions	How to reduce it
Costs	All unit costs, counts and phasing are Delta Doctrines assumptions; no Indian price has been published	Total could differ materially; P10-P90 range reflects unit costs only	Request quotes; refine at Gate 0
Cost basis	Constant 2026 rupees at Rs 90 per dollar; no inflation, currency or escalation	Nominal outlay would be higher	Index when budgeting
Reported items	The MoU [S1], defence AI layer [S2] and Rs 300 crore centre [S3] are press reports; funding is undisclosed	Used only as context; not costed	Confirm with MoD
Sarvam	Parameters, context and active-parameter figures differ between sources; benchmark scores are third-party index results	Choice of model base is open	Independent military test
Quantum radar	Published analyses disagree by orders of magnitude	Recommendation to avoid it is a judgement from the better-quality sources	Track new trials
Maturity scores	Judgement; not official TRL	Priority order could change	Test in trials
Latency figures	Illustrative; no public Indian data	Speed-up claim is a design intent, not a measurement	Measure in exercises
Chart-side figures	Rental rates above Rs 65 are unverified; 2018 and 2022 timeline entries not re-verified	Own-versus-rent conclusion rests on the subsidised rate only	Verify rates
Sudarshan Chakra	A study-site radar claim could not be confirmed	Excluded	-

HOW THE MODEL CAN BE AUDITED

The cost model is a ten-year cash flow with 17 lines and 18 inputs, listed on page 22. The Monte Carlo uses 100,000 draws with a fixed seed. Every number in the tables of Chapter 5 is calculated from the same code and ties to the Rs 24,631 crore base total.

Not covered. Nuclear command and control, space-based sensing architecture, offensive cyber operations, autonomous weapons policy, quantum computing, and legal and treaty questions. Each deserves separate treatment. Nothing in this paper should be read as advice on the use of lethal force.

ANNEX

Source register

Sources were read in October 2026 unless noted. Items marked 'search listing only' were seen in search results and not read in full; they support only the facts shown. URLs are given without the scheme for width.

Tag	Source	Location	Label
S1	Digital Terminal, 'Sarvam signs five-year MoU with Defence Ministry to develop AI solutions for India's armed forces', 27 Oct 2026	digitalterminal.in/amp/story/tech-companies/sarvam-signs-five-year-mou-with-defence-ministry-to-develop-ai-solutions-for-indias-armed-forces	Reported
S2	GKToday, 'India to Build Secure Military AI Layer', 24 Oct 2026 (statement of 23 Oct)	gktoday.in/india-to-build-secure-military-ai-layer/	Reported
S3	Whalesbook, 'India to Build Rs 300 Cr Defence AI Centre with Sarvam AI', 22 Apr 2026	whalesbook.com (aerospace-defense, article 69e891afbca97ee1069f2e86)	Reported
S4	Business Today, 'Budget 2026: Defence gets all-time high of Rs 7.85 lakh crore', 1 Feb 2026 (article has unit errors in the capital figure)	businesstoday.in/union-budget/story/union-budget-2026-defence-gets-all-time-high-...-514048-2026-02-01	Stated
S5	Outlook Business, 'Over 17,000 GPUs successfully installed under IndiaAI Mission' (Sarvam: 4,096 H100s; Rs 98.68 crore subsidy; Rs 246.71 crore cost)	outlookbusiness.com/deeptech/artificial-intelligence/over-17000-gpus-successfully-installed-under-govts-indiaai-mission	Reported
S6	Storyboard18, 'Sarvam AI releases India-built 30B and 105B open-source AI models'	storyboard18.com/amp/digital/sarvam-ai-releases-india-built-30b-and-105b-open-source-ai-models-91544.htm	Reported
S7	Artificial Analysis, 'Sarvam 105B and Sarvam 30B: everything you need to know', 2 Apr 2026	artificialanalysis.ai/articles/sarvam-105b-Sarvam-30b-everything-you-need-to-know	Reported (independent)
S8	APXML, Sarvam 30B model page (65K context)	apxml.com/models/sarvam-30b	Reported (search listing only)
S9	SiliconIndia, 'IndiaAI Mission powers ahead with subsidised GPUs, indigenous models and skill push' (Rs 65 average; 34,381 GPUs)	siliconindia.com/news/general/indiaai-mission-powers-ahead-with-subsidised-gpus-indigenous-models-skill-push-nid-237061-cid-1.html	Reported
S10	Morung Express, 'IndiaAI Mission: 34,381 GPUs onboarded from 14 empanelled service providers'	morungexpress.com/indiaai-mission-34381-gpus-onboarded-from-14-empanelled-service-providers	Reported (search listing only)
S11	Data Center Dynamics, 'Yotta Data Services to invest USD 2bn in 20,000 Nvidia Blackwell deployment in Noida' (20,736 B300; 60 MW site)	datacenterdynamics.com/en/news/yotta-data-services-to-invest-2bn-in-20000-nvidia-blackwell-deployment-in-noida-india/	Reported
S12	Drishti IAS, 'National Quantum Mission', 18 Mar 2026 (outlay Rs 6,003.65 crore; not re-read in full)	drishtiias.com/daily-updates/daily-news-analysis/national-quantum-mission-3	Stated (secondary)
S13	Drishti IAS, 'India's 1,000-km Quantum Communication Breakthrough', 13 Apr 2026 (headline only)	drishtiias.com/daily-updates/daily-news-analysis/indias-1000-km-quantum-communication-breakthrough	Reported (headline only)
S14	Prateek Tripathi, 'Quantum sensing and its applications in Indian military navigation systems', ORF Occasional Paper 544, 8 May 2026	orfonline.org/research/quantum-sensing-and-its-applications-in-indian-military-navigation-systems	Stated (think-tank)
S15	ThePrint (PTI), 'DRDO's Quantum Technology Research Centre inaugurated in Delhi', 27 May 2025	theprint.in/india/drdos-quantum-technology-research-centre-inaugurated-in-delhi/2639119/	Stated (PTI)
S16	US Army, 'Army researchers detect broadest frequencies ever with quantum receiver', 4 Feb 2021	army.mil/article-amp/242980/army_researchers_detect_broadest_frequencies_ever_with_quantum_receiver	Stated
S17	QED-C and MITRE material on quantum PNT (search listings; not read in full; not relied on for any figure)	mitre.org (PR-23-0577); thequantuminsider.com (28 Oct 2024)	Context only
S18	Bischeltsrieder et al., 'Engineering constraints and application regimes of quantum radar', IEEE Trans. Radar Systems, 2024 (DLR)	elib.dlr.de/202555/1/quantum_radar_bischeltsrieder.pdf	Stated (peer-reviewed)
S19	Allahverdi and Motazedifard, 'Analytical performance evaluation of quantum radar architectures', arXiv 2606.10436, 10 Jun 2026	arxiv.org/pdf/2606.10436	Preprint
S20	Galati, Pavan and Daum, 'Quantum radar and noise radar: overview and performance evaluation', preprints.org 202605.0752, May 2026	preprints.org/manuscript/202605.0752/v1	Preprint, not peer-reviewed
S21	Raksha Anirveda, 'Home-grown AI tools widely used in Operation Sindoor; military LLM to be functional in six months: DG EME', 8 Oct 2025	raksha-anirveda.com/home-grown-ai-tools-widely-used-in-operation-sindoor-military-llm-to-be-functional-in-six-months-dg-eme/	Stated (briefing)