MINISTRY OF NEW AND RENEWABLE ENERGY

Government of India

Atal Akshay Urja Bhawan, CGO Complex, Lodhi Road, New Delhi – 110 003

EXPRESSION OF INTEREST

for a Market Assessment Exercise on

AI-Based Retrieval Augmented Generation (RAG) Systems

Date of Publication: 20/05/2026

Disclaimer

This Expression of Interest (EoI) is not an offer, a tender, or a procurement instrument. It is a market assessment exercise issued by the Ministry of New and Renewable Energy (MNRE) to gather technical intelligence about AI-based Retrieval Augmented Generation (RAG) systems available in the Indian market.

No procurement, contract, financial transaction, or commercial relationship of any kind shall arise from this EoI or from any stage within it. Participation does not create, and shall not be construed as creating, any entitlement, preference, advantage, or eligibility implication for any future procurement process that MNRE may undertake.

If MNRE subsequently decides to procure an AI RAG system, that procurement shall be conducted as a separate, independent, open competitive process in accordance with the General Financial Rules (GFR) 2017 and applicable Government of India guidelines. Participation in this EoI shall confer no advantage in such a process.

MNRE reserves the right to accept or reject any or all responses, modify or cancel this EoI, and decide not to proceed with any further action, without assigning reasons and without incurring any liability.

No financial compensation or remuneration shall be provided to participating entities for preparation of the EoI response, provision of interface access, or participation in the market dialogue.

1. Introduction and Background

MNRE manages a large and rapidly expanding corpus of knowledge assets — scheme guidelines, policy circulars, technical standards, implementation reports, meeting records, financial data, and inter-ministerial communications — across diverse formats and sources. Current information retrieval relies substantially on manual search and keyword-based systems, which are inadequate for the scale and pace of MNRE's operational requirements.

AI-based Retrieval Augmented Generation (RAG) systems combine semantic retrieval with large language model (LLM)-based generation to provide accurate, contextual responses with source attribution. To understand the current state of the RAG technology market and the capabilities of available solutions, MNRE proposes to conduct this market assessment exercise.

This EoI invites eligible entities to participate. It is a market sounding instrument only and produces no procurement outcome.

2. Objectives

The objectives of this EoI are:

- To identify entities with demonstrated capability in designing, developing, and deploying AI-based RAG systems within India.
- To independently test, on MNRE's terms, the functional capability and technical maturity of available RAG solutions when applied to MNRE document material.
- To gather technical intelligence — including performance characteristics, scalability parameters, deployment architectures, data security approaches, and compliance posture — relevant to a Central Government Ministry use case.
- To use these findings to inform MNRE's internal deliberations and the technical specifications of any future procurement process that MNRE may undertake.

This EoI is a market assessment exercise. It does not constitute, and shall not be construed as, the initiation of a procurement process.

3. Process Overview

The EoI exercise comprises three sequential stages:

Stage	Activity	Description
Stage 1	**Written Submission and Capability Scoring**	All interested entities submit a written response. MNRE screens submissions for participation eligibility (Section 4) and scores eligible submissions against capability criteria (Section 5). Up to 15 highest-scoring entities are shortlisted for Stage 2.
Stage 2	**Independent Interface Testing**	Shortlisted entities provide MNRE with independent access to a working instance of their RAG system. MNRE's testing committee independently tests each system using documents and queries of MNRE's own choosing, with no vendor present. The committee takes internal observation notes; no formal scoring is conducted (Section 6).
Stage 3	**Market Dialogue**	Based on the testing committee's observations, MNRE may, at its discretion, invite some, all, or none of the tested entities for an informal market dialogue. The dialogue is for MNRE's deeper understanding of architecture, roadmap, and deployment considerations. No timeline, format, or minimum number of invitations is committed (Section 7).

4. Participation Eligibility

Participation is open to any entity meeting both of the following mandatory conditions. These conditions exist solely to ensure basic legal accountability under Indian law and the existence of a genuine product. Failure on either condition results in rejection at Stage 1 with no further consideration.

Ref	Condition	Requirement
PE-1	Legal Registration	The entity must be a Company or Limited Liability Partnership (LLP) registered in India under applicable laws, with a valid Certificate of Incorporation or LLP registration certificate.
PE-2	Demonstrated AI/ML Capability	The entity must have an immediately available working product capable of ingesting documents and responding to natural language queries using AI/ML or NLP-based approaches. Prior project evidence (work order, completion certificate, or deployment certificate) must be submitted to support this claim.

5. Shortlisting for Interface Testing

Submissions clearing the participation eligibility conditions (PE-1 and PE-2) are scored against the following five capability criteria. Up to 15 highest-scoring entities are shortlisted for Stage 2 (Independent Interface Testing). The cap of 15 allows MNRE flexibility to extend the shortlist where scores are close or tied.

5.1 Capability Scoring Criteria

Ref	Max	Criterion	Scoring Basis	Source Document
SC-1	**30**	Prior Experience — Nature and Scale	10 — One project, private sector client. 20 — One project, Central/State Government or PSU client. 30 — Two or more projects, at least one for Central/State Government or PSU.	Experience case study and supporting evidence

SC-2	**20**	Document Volume Capacity	5 — Up to 500 documents. 10 — 501 to 2,000 documents. 15 — 2,001 to 10,000 documents. 20 — More than 10,000 documents.	Annexure I (Section C)
SC-3	**20**	Document Format Support	5 — PDF (text-based) only. 10 — PDF + Word. 15 — PDF, Word, images, PowerPoint. 20 — All of the above plus scanned PDFs (OCR), spreadsheets, and web links.	Annexure I (Section B)
SC-4	**20**	Multilingual Capability	0 — English only. 10 — English and Hindi (production-grade). 20 — English, Hindi, and one or more additional Indian regional languages (production-grade).	Annexure I (Section D)
SC-5	**10**	Deployment Flexibility	0 — Cloud only, non-GCC. 5 — GCC-compliant cloud. 10 — On-premise or hybrid (GCC cloud + on-premise).	Annexure I (Section E)
	100	**TOTAL**		

5.2 Tie-Breaking

In the event of tied scores at the cut-off, the entity with the higher SC-1 (Prior Experience) score shall be ranked higher. If a tie persists, the entity with the higher SC-2 (Document Volume Capacity) score shall be ranked higher.

5.3 Note on Section 8 Disclosures

The disclosures required under Section 8 (Data Localisation, LLM Training, GoI Compliance) are mandatory submission requirements but do not contribute to the shortlisting score. They are read by MNRE for intelligence-gathering purposes only and inform MNRE's internal deliberations and the specifications of any future RFP.

6. Independent Interface Testing (Stage 2)

Each shortlisted entity shall provide MNRE with independent access to a working instance of its RAG system within 5 working days of intimation. MNRE's testing committee shall independently test each system.

6.1 Mode of Access

- The entity shall provide a URL with credentials, a deployable sandbox, or a comparable mechanism that enables MNRE staff to operate the system independently without entity personnel being present.
- Access shall remain available to MNRE until the testing committee concludes its assessment of all shortlisted systems. The duration is not fixed in advance.
- The system must be available with reasonable uptime during the testing period. Any planned downtime must be communicated to MNRE in advance.

6.2 Document Corpus and Query Set

- MNRE will upload between 250 and 500 documents of its own choosing. The corpus will consist of publicly available MNRE material — scheme guidelines, published circulars, press releases, annual reports, and similar.
- The corpus will not be shared with the entity in advance.

- MNRE will pose queries of its own choosing in real time. The query set will not be disclosed in advance.
- The entity shall not have access to MNRE's queries, query logs, or any record of MNRE's testing activity.

6.3 Observation Dimensions

MNRE's testing committee shall take internal observation notes across the following dimensions. These are observation categories only — no marks are awarded, no ranking is produced, and no scored outcome results from this stage.

Ref	Dimension	What MNRE Will Observe
OD-1	**Accuracy and Grounding**	Whether responses are factually correct and verifiable against source documents. Whether the system fabricates information not present in the corpus (hallucination).
OD-2	**Source Citation**	Whether responses identify the specific document, section, and page from which information is drawn.
OD-3	**Query Intent Comprehension**	Whether the system correctly interprets natural language queries, including queries with implicit meaning or domain-specific terminology.
OD-4	**Multi-Document Synthesis**	Whether the system draws upon and coherently integrates information from more than one source document in a single response.
OD-5	**Ambiguous Query Handling**	Whether the system handles underspecified queries by seeking clarification or providing structured conditional responses.
OD-6	**Response Format Adaptability**	Whether the system adapts response format — narrative, tabular, bullet — to the nature of the query without explicit instruction.
OD-7	**Multilingual Performance**	Quality of responses to queries in Hindi and any other Indian language supported by the system.
OD-8	**Multi-Format Document Handling**	Performance consistency across different document formats — PDF, Word, scanned images, spreadsheets.

6.4 Outcome of Testing

The testing committee shall prepare internal observation notes for MNRE's deliberations. No scored outcome, no ranking, and no individual feedback shall be communicated to entities. Entities will receive a standard written acknowledgement of participation at the close of the exercise.

7. Market Dialogue (Stage 3)

Based on observations from Stage 2, MNRE may at its sole discretion invite some, all, or none of the tested entities for an informal market dialogue. The purpose of the dialogue is to enable MNRE to develop a deeper understanding of system architecture, roadmap, deployment considerations, and the entity's capacity to support a Central Government Ministry use case.

- MNRE is not committed to inviting any minimum number of entities for dialogue. The decision is entirely at MNRE's discretion and shall be final.
- The format of the dialogue — individual meeting, group discussion, or other — shall be determined by MNRE on a case-by-case basis.
- No timeline for the dialogue is committed. It may be scheduled at any time after the testing committee concludes its assessment.
- The dialogue is informal and exploratory. There is no presentation, no pitch, no scoring, and no outcome ranking.
- Invitation to dialogue confers no preference, advantage, or eligibility implication of any kind in any future procurement process.

8. Mandatory Disclosures

All respondents shall provide the following disclosures as part of their EoI submission. These are mandatory submission requirements but do not contribute to the shortlisting score. They are intelligence-gathering for MNRE's internal use and inform the technical specifications of any future RFP.

Respondents are expected to be specific, accurate, and candid. Where a system does not currently meet a particular standard, respondents shall state this clearly and describe any roadmap toward compliance.

8.1 Data Localisation Plan

- State whether all data storage, processing, and inference components of your system are located within the territory of India.
- Name the specific data centre(s), infrastructure provider(s), and geographic location(s) used for each component.
- If any component — including LLM API calls, embedding model calls, vector database operations, or logging — involves data leaving India, describe which component, the reason, and whether an India-hosted alternative exists or is on your roadmap.
- Describe your data localisation roadmap, if any, including timelines for any planned migration of components to within India.

8.2 LLM Training Plan

- State whether your system uses client-provided documents, queries, or interaction data to train, fine-tune, retrain, or update the underlying LLM or any AI/ML component.
- Describe the technical mechanism that ensures client data is used for inference only. Provide an architectural diagram of the RAG pipeline indicating data flow.
- If your current system does use interaction data for any form of model improvement, describe the nature, extent, and purpose of this use.
- Describe your LLM training and update governance — how is the underlying model maintained, by whom, and on what data.

8.3 GoI Compliance — Status

- State whether your system currently complies with GIGW 3.0 as notified by MeitY. Provide a compliance statement or evidence.
- State whether your system complies with MeitY AI Guidelines as currently notified. Provide the specific notification or document reference being relied upon.
- If deployable on cloud, state whether the deployment complies with Government Cloud (GCC) requirements as prescribed by NIC/MeitY.
- Describe any known gaps in current GoI compliance and your timeline for addressing them.

8.4 GoI Compliance — Approach

- Describe how compliance is achieved across the system stack — at the application layer, the AI/ML model layer, the data storage layer, and the infrastructure layer.
- Describe access control mechanisms — how is document access restricted by user role.
- Describe data encryption in transit and at rest.
- Describe audit logging — whether user queries and system responses are logged, where logs are stored, and how they are protected.
- Describe how the system maintains compliance over time as GoI standards evolve.

9. Documents to be Submitted

Sl.	Reference	Document Required	Purpose
D-1	**Certificate of Incorporation**	Valid Certificate of Incorporation or LLP registration certificate.	PE-1 verification
D-2	**Experience Case Study**	Maximum 2 pages describing the qualifying AI/ML/NLP project. PII may be masked.	PE-2 + SC-1
D-3	**Experience Evidence**	Work order, completion certificate, or deployment certificate.	PE-2 + SC-1
D-4	**Technical Questionnaire (Annexure I)**	Completed response covering system overview, document handling, performance and scalability, multilingual capability, deployment.	SC-2 to SC-5
D-5	**Disclosure Response (Section 8)**	Response to all four disclosure heads — Data Localisation Plan, LLM Training Plan, GoI Compliance Status, GoI Compliance Approach — with architectural diagram.	Section 8
D-6	**Entity Introduction**	Maximum 1 page introduction of the entity, founding, core business, and AI/ML practice.	General

10. Important Timelines

Sl.	Event	Timeline
1	Date of Publication of EoI	20/05/2026
2	Last Date for Submission of EoI Responses	27/05/2026, 23:00 Hrs (7 working days from publication)
3	Eligibility Screening and Shortlisting	Within 7 working days of submission deadline
4	Intimation to Shortlisted Entities	Within 7 working days of submission deadline
5	Provision of Interface Access by Shortlisted Entities	5 working days from intimation
6	Independent Testing by MNRE	Open-ended; until testing committee concludes assessment
7	Market Dialogue (if any)	At MNRE's discretion; no fixed timeline
8	Standard Acknowledgement to All Respondents	At the close of the exercise

11. Submission Instructions

Completed EoI responses, with all required documents listed in Section 9, shall be submitted by email **to <pmu-rts@gov.in>** with a copy to **vikram.mnre@gov.in** , to reach MNRE on or before 23:00 Hours on **27/05/2026**.

- The subject line of the email shall read: 'EoI Response — AI RAG Market Assessment — [Entity Name]'.
- All documents shall be submitted in PDF format unless otherwise specified.
- No hard copy submission is required.
- MNRE may seek written clarifications after receipt of submissions. Entities shall respond within 3 working days.
- Responses received after the deadline shall not be considered.

Please also submit your responses in the Google Form in respect of EoI https://forms.gle/CGe5Gq79R7Lr6yU3A

MNRE reserves the right to modify timelines, seek additional information, or cancel this EoI at any stage without assigning reasons and without liability to any entity.

Annexure I — Technical Questionnaire

Provide specific and technical responses to all questions below. This questionnaire feeds shortlisting criteria SC-2 through SC-5 and provides MNRE with technical intelligence. Vague or marketing-language responses will be noted adversely.

Section A: System Overview

1. Is your system custom-built, based on an OEM product, or built on open-source components? Name all key tools, frameworks, models, and infrastructure components with version numbers.
2. Describe the end-to-end RAG pipeline: document ingestion → chunking → embedding model → vector store → retrieval mechanism → LLM for generation → response formatting. Provide an architecture diagram.
3. Which LLM(s) does the system use for generation? State hosting location and infrastructure provider for each.

Section B: Document Handling (feeds SC-3)

4. List all document formats supported. For each, state whether support is full or partial and describe any limitations.
5. Describe OCR capability for scanned PDFs. State accuracy levels and supported languages for OCR.
6. State the maximum size per individual document and the maximum total data volume the system can handle.

Section C: Performance and Scalability (feeds SC-2)

7. State the maximum number of documents your system can ingest in a production deployment.
8. State the maximum document volume your system has actually been deployed with in production.
9. Provide average end-to-end response time for a corpus of 250 documents and for 2,000 documents.
10. State the maximum number of concurrent users supported and the infrastructure required at that scale.

Section D: Multilingual Capability (feeds SC-4)

11. List all languages in which your system can accept queries. For each, state maturity: production-grade, beta, or experimental.
12. List all languages in which your system can generate responses. State maturity for each.
13. Describe any limitations in multilingual handling — e.g., cross-language synthesis, code-mixed queries, transliteration.

Section E: Deployment (feeds SC-5)

14. Is your system deployable on-premise (within MNRE's or NIC's data centre)? State hardware and software prerequisites.
15. Is your system deployable on Government Cloud (GCC)? Describe the deployment architecture and GCC compliance status.
16. Is hybrid deployment feasible? Describe constraints.

Section F: Known Limitations

17. List any known limitations of your system relevant to a Central Government Ministry deployment. Be specific.

Annexure II — Illustrative Testing Scenarios

The following scenarios illustrate the type of tasks that MNRE's testing committee may include during independent interface testing. The actual document corpus and query set will be chosen by MNRE at the time of testing and will not be shared with entities in advance.

A. Basic Retrieval and Citation

- Query a single document for a specific fact. Assess whether the response is correct and whether the source document, section, and page are cited.
- Query for a numerical parameter (subsidy amount, capacity threshold, deadline) from a scheme guideline document.
- Query for the definition of a domain-specific term (e.g., net metering, ALMM, CFA, RPO).

B. Multi-Document Synthesis

- Query requiring information from two or more documents to be combined into a single coherent answer.
- Query that surfaces a potential inconsistency between two documents — assess whether the system flags this.
- Query requesting a comparison of parameters across multiple scheme documents.

C. Ambiguous and Complex Queries

- Query with incomplete context — assess whether the system seeks clarification or provides a structured conditional response.
- Query phrased in informal language without domain terminology — assess whether intent is correctly understood.
- Multi-part query requiring the system to address several sub-questions in a single response.

D. Multilingual Queries

- The same query posed in Hindi — assess whether the response is accurate, complete, and fluent.
- A Hindi query requesting information that requires synthesis from an English-language source document.

E. Format Adaptability

- Query that naturally calls for a tabular response — assess whether the system provides one without explicit instruction.
- Query requesting a summary — assess length, accuracy, and source attribution.
- Query requesting step-by-step instructions — assess whether the system structures the response appropriately.

F. Multi-Format Document Handling

- Query against content from a scanned PDF — assess OCR quality and retrieval accuracy.
- Query against content from a spreadsheet or table embedded in a Word document.
- Query against content from a presentation file.