RFP- Building an AI-powered Cognitive Command Centre as a Decision Support System for DMRC's Operations Control Centre
Responses to Vendor Queries

Q#	RFP section	Question	Response
1	Sec 4.1.Eligibility criteria / Minimum average turnover in India in the last 2 years to be USD 1.5 Mn / INR 15 Cr	The audited balance sheet for FY 2025–26 is currently under finalisation. Please confirm whether, in lieu of the audited financials, the bidder may submit a Chartered Accountant certified letter for this period along with the bid.	A letter certified by a Chartered Accountant, accompanied by the financials, will suffice for this requirement in respect of FY 2025-26 where audited financials are not yet finalised.
2	4 - Eligibility / Legal entity in India	Please confirm whether Global parent company credentials are applicaple for uses case references .	The bidding entity must be a legal entity registered in India and must independently satisfy the turnover threshold in Sec 4.1, which is expressly framed as turnover in India. Use-case references executed by the global parent or affiliate entities may be cited in support of the technical proposal, provided the bid identifies the executing entity for each reference and confirms that the relevant capability and named experts will be available to the Indian contracting entity for this engagement.
3	4 - Eligibility / Legal history and debarment	Please provide the prescribed self-certification format and clarify the scope.	Bidders may submit a self-declaration on company letterhead, signed by CEO/ authorised signatory, confirming that as at the bid submission date the bidder is not debarred, blacklisted or suspended by JICA
4	4 / Eligibility – Conflict of interest	Is bidding as a consortium or the use of subcontractors permitted, given the conflict-of-interest single-bid rule?	Consortium bids are not permitted. Sub-contracting of the core solution/development scope is not permitted. However, Digital partner may engage specialist/expert third parties (e.g. penetration testing) with explicit approval from DMRC/Operator
5	12 - Proposal Format / Bid validity period not specified	Please confirm the required validity period of the technical and financial proposals, the period for which quoted rates must remain firm, and the process for seeking any extension of bid validity.	Proposals and quoted rates shall remain valid for 120 days from the bid submission deadline, or until completion of the selection process, whichever is later. If extension is required, it will be formally communicated in advance with no associated penalties for withdrawl
6	NA / General Query	As the proposal is to be submitted by email, we would like to confirm whether there is any limit on the total attachment size (in MB) for a single email. If such a limit applies, please confirm whether bidders are permitted to submit the proposal in multiple emails, with each email clearly indicating the page numbers and sections/criteria to be read together as one complete proposal.	There is no limit on the total attachment size. We don't expect the file sizes to exceed email limits but due to any technical challenges, proposal could be submitted in separate emails with clear instructions using the same sender email, RFP prescribed subject line on the same date/time of the day.
7	Commercial Obligations / EMD, PBG, retention, LD, and service credits not stated	Please confirm that no bid security/EMD, performance bank guarantee, retention money, liquidated damages, delay penalties, SLA service credits, or parent-company guarantee is required unless expressly included in the final addendum/contract. If any apply, please provide the amount/percentage, cap, validity, trigger, cure period, and exclusions before bid submission.	No bid security or EMD, performance bank guarantee, retention money or parent-company guarantee is required for this engagement. The consequences of delay are governed by Sec 3 of the RFP and by the delay provisions to be settled at contracting. Should any such instrument be introduced, it will be notified by addendum before the bid submission deadline, stating amount, cap, trigger, validity and exclusions, etc.

8	11.3 and Attachment 2 - Commercial Bid / USD quotation, invoicing, and taxes	Please confirm whether the Indian contracting entity will invoice and be paid in USD or INR, the exchange rate/date for contract and payment purposes, how GST will be shown and paid, treatment of TDS/withholding and tax gross-up, and whether the evaluated price will be inclusive of all taxes that are not creditable to the contracting authority.	The Digital Partner will contract with the Operator (on JICA's behalf) which will be a outside India entity. Bid amounts are quoted in USD inclusive of all applicable taxes. The precise GST treatment will be reviewed during contracting process as per advice from legal/statutory teams. For reference, in comparable prior arrangements, contracts have been executed with no GST applicable.
9	Section 9 & Section 10 — Budget, IP and Contracting / "The maximum budget for the services is USD 300,000 including all taxes and expenses." The RFP further states that the Digital Partner's contract will be with the Operator appointed by JICA, currently identified as BCG."	Which legal entity will execute the contract with the Digital Partner — JICA, DMRC, or the Operator (BCG)? And will this be a domestic Indian entity or a foreign entity? This is material to determining whether the engagement qualifies as an export of services under GST law.	
10	Sec 9 – Budget Ceiling / Scope of the USD 300,000 Budget Ceiling	We would be grateful if the department could kindly reconfirm in writing that cloud infrastructure, AI model costs, and all user licenses fall within the USD 300,000 budget ceiling, and also clarify whether this extends to the post-Go-Live/Milestone 6 costs.	Any cost incurred during PoC time duration for project delivery(including stabilization) falls within the USD 300,000 ceiling. This covers implementation and professional services, cloud consumption and managed AI/LLM usage for the PoC period, all software and tool licences including any BI or visualisation licences, OEM/API connectors fees, hardware procured by the Digital Partner, security and penetration testing, travel, stay and on-site support, training and knowledge transfer, and all applicable taxes and duties.
11	9 and Attachment 2 - Budget Ceiling / Costs included in the USD 300,000 ceiling	Please confirm whether the ceiling includes all cloud consumption, managed AI usage, software licenses, OEM/API/connector fees, travel and stay, onsite support, security testing, hardware installation, taxes, duties, training, warranty, and post-go-live support. Please specify the required license/cloud-cost period to be priced and whether taxes introduced or changed after bid submission are treated as a change in law.	
12	Sec 9 – Budget Ceiling / License Duration for Costing	We would be thankful if the department could kindly reconfirm in writing that software/tool license duration may be scoped to the PoC/tenure period itself, without the need to consider post-PoC licensing costs at this	All cloud consumption, any software and tool licence durations are to be scoped for the PoC tenure period only. Any continuation beyond the PoC will be assessed and contracted separately and bidders should indicate indicative post-PoC run rates for information only The Go-Live stabilisation period falls within the USD 300,000 ceiling.
13	9 / Budget ceiling – cloud run rate	We understand the USD 300K ceiling is all-inclusive, with cloud consumption costs to be baked into the project cost. Please confirm the run-rate period over which cloud consumption (compute, GPU / TPU, storage, backup, networking, data transfer, monitoring, STT and LLM usage) must be costed – the PoC duration only (~6 months), or an extended operating period post-PoC (e.g., 1 year, 2 years)?	

14	Sec 8 – Deliverables, Milestones and Timelines / Post Go-Live (Milestone 6) Support & Cloud Cost	We kindly request the department to confirm whether ongoing cloud consumption/LLM inference costs and AMC/support costs during the post-Go-Live stabilisation period (Milestone 6) would be included within the USD 300,000 budget ceiling, or budgeted/contracted separately.	
15	8 / Solution stabilization & handover	What post-PoC support / warranty obligations apply after Final Acceptance (T+5.5 months), and are they within the budget ceiling?	
16	Commercial Bid Evaluation (General Query) / Cloud Subscription Ownership & Cost Mechanism	We kindly request the department to clarify who would own the cloud subscription (GCP or other project), and whether cloud consumption costs should be quoted as a fixed-price component or as pass-through-at-actuals within the overall budget ceiling.	For the POC, DMRC will not provide any existing cloud subscription, solution components and landing zone. Cloud cost for the PoC is paid by the digital partner must be baked into the bid as a line item within the ceiling - i.e. absorbed, not passed through at actuals.
17	7.2.1 and 7.4 - Cloud Environment / Tenant, landing zone, connectivity, and billing	Please confirm whether JICA/DMRC will provide an existing cloud organization/project/subscription and landing zone or whether the Digital Partner must create and fund one.	Cloud provider Bill Of Material (BOM) planned for the PoC should be included in the technical solution proposal.
18	Sec 7.2.1 – AI Platform Architecture / Cloud Architecture, Hosting Model & Cost	We would be thankful if the department could kindly confirm in writing the hybrid hosting model discussed during the pre-bid meeting (application hosted on-premise on DMRC servers, with AI/model services on cloud). We would also appreciate guidance on who will provide, own and maintain the on-premise servers, and whether any related cost is to be included in our commercial bid.	The architecture planned is hybrid: the application is hosted on-premise on DMRC servers, while AI and LLM services run on cloud. A reference architecture is provided in the RFP. The Digital Partner is expected to specify the server availability/technical requirements it needs from DMRC
19	Sec 7.2.3 – System Integrations (RCW) / Hardware for RCW Integration	We would be thankful if the department could kindly confirm that, in case RCW integration is not feasible, the resulting hardware cost is to be included within the USD 300,000 budget ceiling. We would also request that RCW configuration/technical details be shared at the earliest, to help us assess integration feasibility and estimate this cost accurately.	As per our current understanding, hardware costs should be planned for a tablet like screen which will be used to display the application and dashboards for POC purpose only. The costs for hardware is within the USD 300K ceiling. It is excluded from QCBS evaluation scoring, but must be included in the commercial submission.
20	Section 7.3 ASR / STT — Technical Requirements / "STT service to be deployed on GPU/TPU hardware at DMRC OCC or on cloud"	If on-premise GPU/TPU deployment at DMRC OCC is chosen as the preferred architecture for latency or data security reasons, will DMRC provide physical rack space, power, cooling, and network connectivity at no cost to the Digital Partner? Or is procurement, provisioning, and installation of on-premise hardware entirely within the Digital Partner's scope and budget?	A reference architecture for a hybrid plan is provided in the RFP, where the application is hosted on-premise on DMRC servers, while AI and LLM services run on cloud. The Digital Partner may put forth their own recommendations ensuring delivery of the program in a cost/time efficient manner

21	7.3 - Technical requirements / STT Model	If the STT service is to be deployed on GPU/TPU infrastructure at the DMRC OCC, please clarify whether the associated hardware cost will be borne by the bidder or provided by DMRC. Please confirm whether the GPU/TPU infrastructure is configured with redundancy/high-availability provisions to ensure uninterrupted service. Kindly provide details of the containerization/orchestration platform available in the OCC environment (e.g., Docker, Kubernetes, OpenShift, etc.)	
22	Sec 7.4.1 – Performance & Scalability / Future Scale-Up Beyond Line 2 Pilot	We would be grateful if the department could kindly confirm whether the PoC solution is expected to be architected for ready scalability to all 10 lines in future, or whether infrastructure sizing for this PoC may be based only on the Line 2 numbers shared. This understanding would help us plan our infrastructure and pricing thoughtfully.	Infrastructure sizing for the PoC is to be based on the Line 2 numbers provided in the RFP. Please refer to answers on data volumes for architecture sizing Given that the architecture is planned on cloud, expected to have the design of the architecture to be scalable. Actual capacity increase to be done as per requirement depending on onboarding of additional lines. Any scope changes to be handled through change control and project governance process.
23	7.2 (3); 8 / Change control	If additional integrations are identified during BRD / TRD finalization, please confirm they will be handled through mutually agreed change control rather than automatically absorbed into the fixed bid price.	
24	Section 7.4 Non-Functional Requirements / "Concurrent LLM queries: 15" (scoped for Line 2 PoC)	If DMRC decides to expand the PoC to additional lines during the contract period (as indicated in Section 2 — Secondary Objective), will concurrency requirements, infrastructure sizing, and budget be revised through a formal change order? Or is the Digital Partner expected to design and absorb network-wide scalability within the current USD 300,000 ceiling?	
25	3 - Duration and Timeline / Note on revision of duration and delay costs	Please clarify the time and commercial relief available where delay is attributable to JICA, DMRC, the Operator, source-system OEMs, delayed data/API access, security approvals, site access, infrastructure provisioning, or delayed sign-offs. Please confirm that such delays will result in an equitable extension of milestones and, where applicable, an approved change request, and that the Digital Partner will bear additional costs only for delays directly attributable to it.	Please refer to RFP section 3.1 Contract duration.
26	3 / Duration and Timeline	Regarding additional costs of delay being borne by the Digital Partner: is this liability capped, and how are delays caused by DMRC / third-party dependencies (data access, OEM support, approvals) treated?	
27	8 / Milestone acceptance	How will milestone acceptance be handled if representative data, source-system access, OEM support or test incidents are unavailable?	
28	Sec 10 / IP and Other considerations	We request the authorities to let bidders own the source code of the solution. Any customization done will be shared with the department.	Bidders are referred to Sec 10, IP and Other Considerations, for the provisions governing intellectual property.

29	8 - Deliverables, Milestones and Timelines / Handover and environment transfer	Please confirm the required handover contents and formats, including source repository, branches/history, build instructions, CI/CD, infrastructure-as-code, deployment scripts, configuration, prompts, model artifacts/adapters, data schemas, test assets, runbooks, licenses, credentials, monitoring dashboards, cloud resources, and knowledge-transfer sessions. Please clarify whether deployment must occur in a DMRC-controlled repository and cloud tenant from inception.	
30	10 / IP and other considerations	Where models are fine-tuned on DMRC data using the bidder's pre-existing base models / accelerators, what precisely transfers to JICA / DMRC versus remains pre-existing IP (e.g., fine-tuned weights,	
31	Sec 7.2.2 – Data Sources & Systems / Volume of Historical Data for Knowledge Base	As discussed during the pre-bid meeting, we would be grateful if the department could kindly share the volume/time-range of historical audio recordings and the number/size of DMRC manuals available, at the earliest possible convenience, to help us plan our data ingestion effort.	
32	Sec 7.3.2(ii) – ASR/STT Model / STT Labelling & Annotation Volume	We would be grateful if the department could kindly share the approximate volume (in hours) of historical raw recordings that would need to be transcribed and labelled for STT fine-tuning, to help us estimate this effort accurately.	
33	Quality of Proposed Solution & Architecture / Knowledge Base Build & Ingestion Pipeline	Please provide an estimate of the volume and formats of data to be ingested (Operational Manuals, SOPs, PRA incident records, OCC audio recordings)	
34	7.1.1 - Voice Capture and Transcription / CDRS scale and historical corpus	Can sample documents, call records, transcription or any othe related artefacts be shared ?	
35	7.1.1 - Voice Capture and Transcription / CDRS scale and historical corpus	Please provide the average and peak number of PTT files per hour/day, audio hours per day, number of concurrent talk groups, historical recording period and total volume available for model adaptation, expected retention period, percentage of noisy/partial/corrupt files, and whether ground-truth transcripts or incident labels exist.	For Line 2, approximately 10,000 recordings are generated per day, each of approximately 5 -10 seconds of duration. The current system allows storage of recordings for 15 days in CDRS servers Recordings are not transcribed. No transcripts, ground-truth references or incident labels exist. Transcription, speaker labelling and annotation for fine-tuning are within the Digital Partner's scope and must be priced within the ceiling. Representative samples will be provided to the Digital Partner during project kick-off (T+7 days)
36	7.2.2 - Data Sources and Systems / Data inventory, quality, and labels	Please provide the quantity, period, format, language mix, quality profile, and accessibility of manuals, PRA records, OCC recordings, ATS, AFC, and network metadata. Please identify missing fields, known data-quality issues, existing labels/taxonomies, data dictionaries, and the volume of SME-validated examples available for training, tuning, and acceptance.	

37	7.2 (2) / Data sources – CDRS	Please confirm the volume of historical Line 2 CDRS recordings to be provided for model development, testing and acceptance (total hours and approximate storage volume), whether human-verified transcripts are available (number of hours and distribution across Hindi / English / Hinglish and noise categories), and when these will be provided.	
38	Sec 7.3.1, 7.3.2, 7.3.5 / Audio Data Ownership	Will DMRC provide historical OCC audio in raw form only, or will annotated/speaker-labelled transcripts also be provided? If annotation and speaker-labelling are expected from the Digital Partner, please confirm this is understood as in-scope effort.	
39	Sec 7.3.2, 7.4.4 / Representative Audio Samples	Can DMRC provide representative samples of historical OCC radio audio (covering normal operations, incidents, and high-interference/emergency scenarios) prior to proposal finalization, to allow bidders to benchmark STT feasibility?	
40	7.3 - Technical Requirements / Knowledge Base Preparation	Please provide the approximate volume and historical time span of CDRS audio recordings, PRA incident records, and DMRC manuals that will be available for knowledge base creation, model training, and STT fine-tuning. Will a representative sample dataset be made available to the bidders.	
41	7.1.1 (i); 7.3 (2) / Audio specifications	Please provide the CDRS .WAV file specifications – sampling rate, bit depth, channel configuration, codec, average duration and typical file size – and the transcription accuracy achieved by the existing transcription capability.	CDRS stores recordings as WAV files. The sampling rate, bit depth, channel configuration, codec, average duration and typical file size will be confirmed together with the CDRS interface specification after contract signing. DMRC does not operate a transcription capability of the recordings
42	7.1.1 (i) / Radio interference filtering	Please provide DMRC's defined audio-noise parameters, noise categories, measurement methodology and representative recordings from each category.	The OCC radio environment is noisy, with heavier interference at the train end and comparatively cleaner audio at the OCC end. Radio conversations are stored as individual radio recordings and speaker overlap doesn't exist. Recordings will be shared with the selected vendor during project set-up phase
43	7.1.1 (i) / Audio quality	What proportion of recordings is affected by radio static, clipping, weak signal, speaker overlap, incomplete transmissions or other material quality issues?	
44	7.1.1 (i) / Conversation sequencing	A single conversation generates multiple .WAV files. Beyond Radio ID and timestamp, what metadata (e.g., call/session ID, Talk Group ID) is available to sequence and group files into a single conversation thread?	Grouping of transmissions into conversation threads design should be finalized by the digital partner during design stage. Additionally the available meta data for inferring includes Unique Radio IDs, Radio ID to Rake ID mapping, Talk groups / Line / network section mapping
45	7.2 (2) / Audio-to-PRA linkage	What proportion of historical recordings is reliably mapped to the corresponding PRA incident, and what common keys or timestamp-matching mechanism is available?	Radio / Caller ID, talk group and time-stamps will be captured by CDRS. Each Radio / Caller ID is mapped to a Rake ID (unique train set identifier). ATS captures Train ID and location of the train. Mapping of Train IDs with Rake IDs would be available.

46	7.1 - Incident Resolution Guidance / 7.3 - NLP & Automated PRA Logging	To reliably correlate a radio conversation with the corresponding train/section (from ATS) and the resulting PRA incident record, please confirm: (a) the common keys available to join a CDRS audio segment to an ATS train/section and to a PRA record; and (b) whether CDRS, ATS and AFC timestamps are synchronised to a common clock (for example NTP), including how time zone, daylight-saving and clock drift are handled across these systems.	OCC recordings are not mapped to the corresponding PRA incident. At present, PRA is manually updated by OCC staff post incident occurence CDRS, ATS and AFC timestamps are synchronised to a common clock.
47	7.1.1 (i) / Radio Conversation Monitoring	Per the pre-bid call, DMRC already performs voice-to-text transcription of radio calls for call analysis. Can these existing transcriptions be re-used by the solution rather than performing duplicate transcription?	DMRC doesn't perform any voice-to-text transcription of radio calls at present - this falls within the scope of the Digital partner
48	Sec 7.3.1(i) – Document Ingestion Pipeline / Digitisation % and Responsibility	We kindly request the department to share the approximate percentage of manuals/SOPs/signalling documents/PRA records currently available in machine-readable format versus scanned or physical form, along with guidance on who would be responsible for digitising any remaining	A large part of DMRC's manuals/documentation are held in machine-readable format. A small percentage would be available as scanned documents as well. For Line 2 the document corpus comprises approximately 350 files identified. For sizing purpose maximum of 400 files should be considered with pages count range of between 2-250 pages with an average page count around 30-40 pages Document ingestion, OCR and structuring into the knowledge repository are within the Digital Partner's scope, including handling of text, technical diagrams, tabular content and handwritten annotations (very limited)
49	7.3.1 - Knowledge Base Preparation / OCR and physical-document digitization	Please provide the expected number of documents/pages, formats, languages, diagrams, tables, and handwritten annotations.	
50	7.3 (1) / Knowledge Base preparation	Please provide the approximate number, page count, file size, language and format distribution of the manuals, SOPs, circulars and other documents to be ingested, and how frequently they are updated.	
51	7.3 (1) / Document digitization	What is the approximate quantity of physical-only documents, scanned documents, handwritten annotations, technical drawings and complex tables requiring digitization or manual remediation?	
52	Section 7.3 Knowledge Base Preparation / "Document ingestion and OCR processing to produce machine-readable text. OCR engine must handle Hindi text, technical diagrams, tabular content, and handwritten annotations"	What is the approximate total volume of DMRC manuals (number of documents, total page count)? Are these manuals available in digital/searchable PDF format or primarily as scanned/physical copies? What is the total volume of (no. of docs, total page count) as scanned/physcial copies? Are the docs available in English Language or any other languages are also present?	

53	7.2 - Data Sources & Systems / 7.3 - Knowledge Base Preparation	Please indicate the approximate percentage of DMRC manuals, SOPs, signalling documents, troubleshooting guides, and PRA records available in machine-readable format versus scanned/physical form. Kindly confirm whether digitisation/OCR of non-digital documents is the responsibility of DMRC or falls within the bidder's scope.	
54	Section 7.3 Knowledge Base Preparation / "Updates to the existing versions of the documents available from DMRC need to be maintained in the knowledge base"	What is the frequency of updates to DMRC operational manuals? Is there an existing versioning/change management process at DMRC, or will the Digital Partner be required to build a detection and ingestion mechanism for manual updates? Who from DMRC will be the designated point of contact for notifying the Digital Partner of document changes?	Operating manual/ safety cirulars are updated as and when operating conditions change on adhoc basis. There is a defined process of updates and communication with-in DMRC and all concerned teams are made aware of the changes. If any changes happen during the project stage Digital partner will be made aware and informed appropriately
55	7.2 (2) / Data sources – PRA	Please confirm the number of Line 2 historical PRA records, the period covered, the expected extraction volume and the access method (database access, API, export). The approximately 1,500–2,000 records mentioned during the meeting may please be formally confirmed.	Approximately 1,500 to 2,000 historical PRA incident records are held for Line 2 on an annual basis. PRA is a structured database maintained in-house by DMRC. Each incident occupies one row, capturing – date & timestamp, line, station / section, Train ID, Train set, related department & sub department, affected equipment, equipment number, status, failure description & resolution, detention, etc, and failure and resolution remarks held as free text.
56	7.2 (2) / PRA data quality	Please provide an anonymized PRA sample, schema and data-quality profile indicating missing fields, duplicates, free-text content and the availability of incident type, severity, asset, root cause, action and resolution labels.	
57	7.1.1 (ii) / OCC-specific terminology	Will DMRC provide a glossary of OCC-specific terminology, call signs and codes for STT calibration?	Historical OCC radio recordings/ DMRC Manuals/SOPs will be provided to the Digital Partner at project kick-off for the purposes of STT model training and speech-to-text fine-tuning. Where the Digital Partner requires clarification on the content of the captured communication, such queries should be raised with DMRC, and DMRC will provide the necessary clarification through its nominated SMEs
58	Sec 7.1.6 – Front End Dashboard / Dashboard / UI Language	We kindly request the department to confirm the language(s) expected for the front-end dashboard, solution guidance display, and downloadable reports — whether English only, or Hindi and English both.	Language support required - 1. Recordings - Speech to Text - Hindi, English, Hinglish 2. Document Corpus - English, Hindi 3. Application Front End - English, Hinglish (speed-to-text transcribed) 4. Analytical dashboards, reports - English
59	7.1.1 (ii) / Multilingual transcription	Languages are listed as “including Hindi, English, Hinglish”. Please confirm the complete list of languages / dialects to be supported.	
60	Sec 7.3.7 – Storage and Audit Trail / Data Retention Period	We kindly request the department to share the specific retention period expected for call recordings, transcripts and audit logs, as this would help us plan our storage sizing and cost appropriately.	For the PoC duration, all production system recordings, transcripts and logs built should be stored and available
61	7.4 - Non-functional Requirements / Storage and Audit trail	Please share DMRC's data-retention policy for audio recordings, transcripts, incident data and system logs.	
62	8 / Programme planning & design phase	When will the Digital Partner receive access to DMRC manuals, historical PRA records and sample CDRS audio? Will samples be available during the design phase (T+1 month), given Wave 1 go-live at T+3 months?	DMRC manuals, historical PRA records, sample CDRS audio, network and OCC metadata is targeted to be made available to the Digital Partner at project kick-off (T+1 weeks). If any other data is required for solution build/ testing purpose, Digital partner should inform on

63	7.2 (3); 7.5 / Environments & test data	Please confirm the availability date of non-production environments, test credentials and representative test data for each required integration.	the same to DMRC during design stage for ensuring the readiness.
64	Section 3 Duration and Timeline / "The contract duration would be ~6 months, from September 2026 to mid-February 2027"	Wave 1 delivery (T+3 months, December 2026) requires the knowledge base, STT fine-tuning, and PRA historical data to be available from Day 1. What is DMRC's committed timeline for providing access to CDRS historical recordings, Network / OCC Metadata, PRA data exports, all operational manuals, and OCC system credentials from the contract commencement date?	
65	7.5; 7.6 / SME availability	Please confirm the number and committed availability of DMRC subject-matter experts / TCs for transcript validation, incident labelling, document validation, UAT and acceptance ground-truth preparation across the Waves.	DMRC team is committed to make this project successful and required support will be extended to the project team. Bidders are expected to share specific expertise required and estimated efforts from DMRC SMEs
66	7.3 (1); 9 / Data preparation scope	Are manual audio transcription, incident labelling, data cleansing and audio-to-PRA reconciliation included in the Digital Partner’s fixed scope?	Yes. All activities for data preparation like audio transcription, labelling, document digitisation are all within the Digital Partner's fixed scope. Data extraction from DMRC source systems (as required) will not be in Digital partner scope
67	1.1 and 7.2 - Existing DMRC Digital Environment / Existing data platform and control tower	Please provide an overview of the foundational data platform and integrated control tower established since 2023, including current cloud/technology stack, data lake/warehouse, integration patterns, reusable APIs, network connectivity, security controls, incumbent support arrangements, ownership of artifacts, and components that the selected Digital Partner is expected or permitted to reuse.	For this PoC, bidders should not assume reuse of any existing platform component
68	7.2 (3) / Integration scope	The RFP describes the integration list as indicative and non-comprehensive. Please provide the definitive integrations required in each delivery wave for commercial estimation.	The systems in scope for integration during the PoC are CDRS, ATS, AFC, PRA, RCW, Active Directory and the BI or visualisation layer, as set out in Sec 7.2.3.
69	7.2 (3) / Interface specifications	For CDRS, ATS, AFC, PRA, RCW and Active Directory, please provide the interface protocol, API or database specifications, sample payloads, authentication mechanism, expected data volume and refresh frequency.	Sample payloads cannot be shared at this stage. Data volumes for integration can be derived using the data volumes details shared across different answers.
70	7.2 - System Integrations	For CDRS, ATS, AFC, and PRA, please confirm the available integration interfaces, authentication mechanisms, and availability of non-production environments or sample datasets for PoC development and testing. Please clarify whether DMRC/OEMs will provide integration support and access enablement, or if the Digital Partner is expected to coordinate access independently. Kindly identify the owner of each interface and the expected timeline for provisioning access after project kickoff.	Yes , sample data sets and non-production environments will be made available to the digital partner during project set-up phase
71	7.1.1 (i); 7.2 (3)(i) / Radio channel monitoring / CDRS integration	Please confirm that “continuously monitor OCC radio channels” is satisfied by monitoring CDRS for new audio files (event-based trigger), rather than interfacing directly with radio channels. If direct interfacing is expected, what interface is available?	Final design is expected to confirm the approach for interfacing the CDRS audio files.

72	7.1.1 (i) / Real time / near real time capture	The pre-bid call referenced a 15–20 minute delay. Please confirm whether this applies to CDRS audio file availability, and state the expected latency from end of a PTT segment to file availability in CDRS. How does this reconcile with the real-time transcription and 10-second guidance requirements?	Push to Talk (PTT) to CDRS file availability latency is sub-second. These files are expected to be transcribed (STT) and presented on the user application in real time. Once the conversations are transcribed and AI has classified the problem, 10 second is the max response time that the LLM should take to give a solution guidance in response to the identified incident. The 15-20 minute delay mentioned in the call was regarding the refresh frequency of the AFC data with the central server.
73	7.2 (3)(i) / CDRS trigger	Does CDRS already provide a supported event trigger / API for retrieving new recordings, or must a new connector or OEM modification be developed?	Yes, the planned approach is to retrieve recordings read-only access of CDRS servers and build integration with the system. DMRC is engaging with the OEM vendor to confirm on the integration readiness and is expected to be confirmed before contracting stage. Sub-second latency from PTT release to file availability. For line 2, # of audio recordings per day is ~10k
74	Section 7.1 Voice Capture & Transcription / "The capture process must not disrupt, interrupt, or add any latency to the actual radio communication in the OCC"	Will DMRC provide read-only / passive access to CDRS audio streams or file storage, or does the Digital Partner need to implement a tap/integration at the CDRS layer?	
75	7.2 / 7.4 - AI Platform Architecture / Security	Please clarify the mechanism for making new CDRS audio files available (event/API push or polling), and the expected latency from PTT release to file availability. Please provide indicative average and peak daily PTT audio volumes for Line 2.	
76	7.2 (3)(ii) / ATS integration	Please specify the actual recording / update frequency expected for real-time train-location data.	Train location is fetched and refreshed at real time frequency using ATS. DMRC is already engaging with the OEM vendors to confirm on the integration readiness and is expected to be aligned before contracting stage. During project delivery, the digital partner is expected to lead discussions in mapping the scope and working with internal tech teams or with third party OEM providers as essential to drive integrations with source systems to ensure successful deployment of solution
77	Section 7.1 Incident Resolution Guidance / "Attributes not explicitly spoken during the conversation (e.g., train ID) must be taken as input from available data sources (e.g., ATS) without requiring TC to confirm them manually"	What is the data refresh frequency and available integration mechanism for the ATS system? Does an existing API layer exist, or will the Digital Partner need to negotiate access directly with the ATS OEM?	
78	7.2 (3)(iii) / AFC integration	Please confirm the required AFC polling frequency, and formally confirm the expected 2–15-minute source-data latency discussed during the information session.	AFC data supports station-crowding estimation as an input to traffic-management guidance incidents only. AFC source data replication is near real time in DMRC local servers and replicated to

79	Section 7.1 Traffic Management Guidance / "The system must determine the appropriate Traffic Management decision… basis decision factors such as nearest available crossover points, passenger crowding on involved stations, train timetables & bunching"	What is the current data refresh rate of the AFC system, and is an integration API already available?	central AFC servers with a delay of under 15 mins. During the solution design, the right integration approach and frequency should be aligned as per the requirements of the traffic management solution
80	Sec 7.2.3 – System Integrations / PRA Write-back Integration Ownership	We kindly request the department to share the expected process and timeline for aligning on the PRA write-back integration mechanism with DMRC's internal technical team, as this would help us plan our integration effort.	PRA is an in-house DMRC application, built and maintained by DMRC's technical team. The PRA write-back integration mechanism, and the process and timeline for aligning it, will be shared and aligned during the "Planning and Design" phase (T+1 months) The current solution approach considered is that incident write-back is event-driven, occurring when the Traffic Controller accepts a recommendation. Rejected recommendations are not written to PRA; the rejection is logged for audit and as feedback into model improvement. PRA write-back to be ensured during PoC via secure API - write access in line with DMRC IT policies and role-based authentication. Duplicate detection is planned via incident ID used as primary key Any change in solution design during the design stage should take into account the required integration approach.
81	7.1.5 and 7.3.6 - Automated PRA Logging / Trigger and human approval	Section 7.3 states that PRA write-back occurs when the TC accepts or rejects the recommendation, whereas Section 7.2 describes scheduled integration multiple times per day. Please confirm the intended real-time/batch design and whether every PRA record must be reviewed and explicitly confirmed by an authorized TC/CC before write-back, particularly for rejected recommendations.	
82	7.1.5 and 7.3.6 - PRA Integration / Interface and transaction requirements	Please provide PRA API/database/interface specifications, field dictionary and mandatory validations, authentication and authorization model, sandbox/test environment, expected transaction volumes, duplicate and idempotency rules, error/retry/rollback handling, attachment requirements, audit fields, and the support contact responsible for interface changes and production access.	
83	7.2 (3)(iv); 7.3 (6) / PRA integration	Does a secure PRA read / write API already exist and support all required operations, or must the Digital Partner design and implement it (if so, to what timeline)? Further, the RFP refers both to scheduled PRA integration and immediate write-back after TC action – please clarify whether write-back must be real-time / event-driven, scheduled, or	

84	Section 7.3 Automated Incident Logging — PRA / "Automated incident logging by writing directly to PRA… PRA write-back via secure API"	Does the PRA application currently expose a write API, or will the Digital Partner need to build a custom integration (e.g., direct database write or UI automation)? If a write API needs to be developed by DMRC's internal team, what is the committed timeline for making it available, and does this fall within the PoC period?	
85	Sections 7.1 and 7.3 - Automated PRA Logging	Please clarify the expected PRA update workflow: should the solution create/update the PRA record immediately upon TC acceptance/rejection, generate a draft for TC review, or wait until incident resolution and explicit user confirmation?	
86	Sec 7.2.3(v) – RCW Integration / Technical Details of Existing RCW System	We kindly request the department to share the technical details of the existing RCW system (make/model, available integration APIs, documentation) at the earliest opportunity, as this would help us plan our integration approach and cost.	DMRC is already engaging with the OEM vendors to confirm on the integration readiness and is expected to be aligned before contracting stage. A separate display device (tablet-like) could be required solely as the front-end interaction interface for the solution (in case RCW screen cannot be used) RCW integrations also serve requirements on train status availability. RCW configuration and technical details - OEM, product and version, supported extension mechanisms, available integration points and APIs, display constraints and security restrictions - will be shared and aligned during the project set-up phase RCW integration is a critical requirement for the PoC delivery of Wave 1
87	7.2 (3)(v) / RCW integration	Please provide the RCW OEM, product / version, supported extension mechanisms, available integration points / APIs, display constraints and security restrictions – required to assess whether the solution interface can run on the existing TC workstation, and whether the feasibility concern is connectivity only or also data exchange.	
88	7.2 (3) / Fallback hardware	If RCW integration is not feasible, please confirm the tablet-type device serves solely as the display / interaction surface for the solution front end (in lieu of embedding in RCW), and what network connectivity (DMRC intranet / OT network / internet) it will be permitted.	
89	7.2 (3); 8 / RCW feasibility decision	Who determines RCW integration feasibility, and by when in the programme, given hardware procurement lead times against the Wave deadlines?	
90	Section 7.2 System Integrations — RCW / "Radio Control Workstation (RCW): Integration to be undertaken with front end interface for TC"	The RCW is a specialised OT system. Does DMRC have an existing API or SDK from the RCW OEM that permits third-party integration? If not, will DMRC facilitate access to the OEM for integration scoping prior to contract signing, and is this integration a hard requirement for Go-Live of Wave 1 or Wave 2?	
91	Sec 7.2 (System Integrations) / RCW Integration Feasibility Assessment	Has DMRC/JICA already conducted a technical feasibility assessment of integrating the Cognitive Command Centre directly with the existing RCW workstation? If so, can findings be shared with bidders to inform the architecture proposal?	

92	7.2 (3)(vii) / AD integration	Will DMRC's Active Directory be made available for SSO / role-based access integration during the PoC?	Final solution is planned with Single Sign-on capabilities and design should consider the same. If any any technical limitations arise during the project, joint decision making will be done to evaluate the best alternatives
93	7.2 - System Integrations / 7.4 - Access Configuration & Security Requirements	Identity & Access Management (IAM) Requirements. Please confirm DMRC identity and access management requirements for the PoC, including Active Directory integration, Single Sign-On (SSO), Multi-Factor Authentication (MFA), password policies and privileged access governance. Kindly clarify whether MFA is mandatory for administrators, support personnel and privileged accounts. Please provide the expected role definitions and access hierarchy for Traffic Controllers, Chief Controllers, OCC Managers, Asset Controllers, Administrators, Auditors and Support Teams. Please confirm whether service accounts, API accounts and integration identities will be provided by DMRC or managed by the Digital Partner.	
94	7.2 (3)(vi) / BI integration	Are Power BI / Tableau licences provided by DMRC, or are they to be included in the bidder's software cost (Attachment 2)?	No BI or visualisation tool is mandated by DMRC, and no existing DMRC licence is available for reuse. The bidder may propose its own BI or visualisation stack, including open-source options. Any associated licence cost must be included in the software cost in Attachment 2 and falls within the USD 300,000 ceiling, priced for the PoC duration only.
95	7.2 (3) / OEM cooperation	Please confirm that DMRC will arrange cooperation from the CDRS, ATS, AFC and RCW OEMs, and clarify who bears OEM support, licence, interface-enablement and system-modification charges.	During project delivery, the digital partner is expected to lead discussions in mapping the scope and working with internal tech teams or with third party OEM providers as essential to drive integrations with source systems to ensure successful deployment of solution For cost on OEM integration considerations please refer to Q# 10/11
96	Sec 7.2 (System Integrations) / OEM API Access & Cost Ownership	For integrations with CDRS, ATS, AFC, and PRA (Section 7.2), will the respective OEMs provide API access and technical documentation to the Digital Partner at no cost, or should bidders assume potential OEM licensing/access fees as part of the technical risk assessment?	
97	Sec 12/13 – Proposal Format / Attachments / Network Segmentation Constraints	We kindly request the department to confirm whether any network segmentation constraints exist that could affect direct cloud connectivity from DMRC's environment, as this would help us plan our networking approach and cost.	Bidders are expected to propose the network approach based on the best practices for critical rail infrastructure with assumptions of secure and dedicated connectivity to meet the performance expectations of the solution. Required expectations from DMRC to be shared in the solution proposal and should be aligned during the project set-up phase
98	7.4 (7) / Security requirements	What connectivity currently exists between DMRC on-prem / OT systems and cloud, and what is the approval process and timeline for establishing the required connectivity (firewall changes, static IPs)?	
99	General - Integration/Connectivity	What secure connectivity option will DMRC provide for PoC: site-to-site VPN, private link, leased line, internet with static IP allow-listing, or another model? And what connectivity exists today between the OCC/DC and public cloud (leased line/VPN), what's the lead time?	

100	Sec 7.1.2 – Incident Resolution Guidance / Output Modality of Solution Guidance	We would be grateful if the department could kindly confirm that the solution guidance to the Traffic Controller is expected only as on-screen text, and that no spoken/voice (text-to-speech) output is required.	Solution guidance to the Traffic Controller is expected only as on-screen text; no spoken or text-to-speech output is required. The Traffic Controller may, however, raise queries both on chat and through voice, with voice input handled by speech-to-text.
101	Sec 7.1.6 – Front End Dashboard / Number of Dashboard/Login Users	We kindly request the department to share the final expected number of dashboard/login users across all roles (TC, Chief Controller, OCC Manager, DMRC leadership), as this would help us plan our licensing and infrastructure costs correctly.	User numbers are specified in the RFP at Sec 7.4 Non-Functional Requirements, under Performance Stability: 10 to 15 Traffic Controller users and 30 to 50 supervisory and management users, giving approximately 40 to 65 named users for the Line 2 PoC. Bidders should size licensing and infrastructure against these figures. All other details to be discussed and finalized during the "Planning & Design" phase
102	Sections 6 and 7.1 - Primary Users and Front-End Dashboard	Please confirm what view will be presented to TC, Chief Controller, OCM, asset-specific controllers, leadership, knowledge stewards, administrators and auditors from RBAC and participation point of view. What devices and channels will those roles use for the viewing?	
103	7.1.7 - Release Plan / Wave 1 functional boundary	Wave 1 includes STT, voice capture/transcription, and LLM incident assistance, while Wave 2 includes live radio transcription and LLM integration. Please clarify whether Wave 1 is to operate on prerecorded audio or text/manual input ?	Please refer the RFP - Sec 7.1 Functional Requirements, "7. Release Plan" for comprehensive details. Pre-recorded audio may be used to showcase the Speech-to-Text model for voice based conversation in Wave 1
104	Section 7.1 - Release Plan	May the bidder deliver a quick workable MVP using historical audio replay, curated manuals, AI guidance and PRA draft generation before enabling live CDRS ingestion and transactional PRA integration? Which capabilities constitute the minimum usable Wave 1 product, and which capabilities may remain in shadow or advisory mode until Wave 2?	
105	7.1.3; 7.3 (4) / Traffic Management Guidance / Optimization Models	Please confirm the scope: retrieving and presenting the applicable procedures, or generating an optimised traffic management plan (crossovers, timetables, crowding)? Section 7.3 (4) implies an optimisation model – please confirm the expected depth for the PoC.	The system must determine the appropriate traffic-management decision, for example single-line operation or short-looping, and present the applicable procedure together with the reasoning and the decision factors considered. Those factors are dynamic rather than a static checklist, and include the nearest available crossover points, train bunching, station crowding and section characteristics. Real-time network state is sourced from ATS for train position & section status, and from AFC for station crowding. The crossover availability at stations to be derived from specific DMRC manuals forming part of the LLM knowledge base Solution details to be designed and finalized during the "Planning & Design" phase
106	7.2 (2); 7.2 (3) / Data sources & integrations	If optimisation is required, which sources provide the real-time network state (actual train positions, crossover availability, station crowding) versus SOP assumptions, and at what refresh latency?	
107	7.3 (4) / Optimization acceptance	Please define the optimization objective, mandatory constraints, expected output and acceptance criteria for Single Line Operation, Short Looping and other Traffic Management recommendations.	
108	7.1.3 (i) / Location identification	What is the authoritative data source for pinpointing the affected track / station / section? Where the conversation is ambiguous or lacks location detail, should the system return “no recommendation”, or is a clarification workflow with the TC preferred?	The affected track, station or section is to be derived from ATS, accessed via the RCW, together with the metadata generated by the CDRS recordings. The Traffic Controller and Train Operator may not state the location explicitly in conversation, and the solution is therefore expected to resolve it from the data available from these systems rather than relying on the spoken content.

109	7.1.2 / Decision authority	Please define the decision-authority model for each use case: informational assistance, AI recommendation requiring controller approval, or automated action.	The solution is a decision support system and does not take any autonomous action on operational systems. Voice capture, transcription and incident detection are informational, presented to the Traffic Controller for awareness. Incident resolution guidance and traffic management recommendations are advisory in both cases: the solution proposes the resolution steps, and the final decision rests with the Traffic Controller, who may accept or reject the recommendation. PRA write-back is initiated only once the Traffic Controller accepts the resolution guidance provided by the system. Where a recommendation is rejected, no PRA record is created.
110	7.4 (4) / Accuracy	What is the current exception-handling process when transcription accuracy is low, and what behaviour is expected of the new system at low transcription confidence?	At present, there is no automated exception process, since transcription is not currently performed. The expected behaviour of the new system at low transcription confidence is to surface the low-confidence segment to the Traffic Controller with the uncertainty explicitly flagged, and to withhold downstream incident detection and recommendation where confidence falls below an agreed threshold.
111	7.1.1 (i); 7.3 (7) / Timestamping for audit & playback	For “timestamped for later audit and playback”: is it sufficient for the audit log to reference CDRS recordings via timestamps, or must the solution store and serve the audio files for playback?	The audit trail must reference the source recording by identifier and timestamp so that the original audio can be reliably located in CDRS. For the PoC, the solution is not required to duplicate and serve the audio files for playback, provided the reference is reliable and the Traffic Controller can reach the source recording through existing means.
112	7.3 (7) / Traceability & retention	Must AI recommendations be traceable to specific SOP / manual clauses and live-system evidence, and what is the required audit-log retention period?	Yes. Every recommendation presented to the Traffic Controller must be traceable to the specific manual, SOP or circular clause relied upon, with the citation displayed alongside the recommendation. Citation of source is a functional requirement, and forms part of the guidance-quality acceptance criteria. The audit trail must retain the recommendation, its citations, the underlying transcript segment and the Traffic Controller's action. The audit-log retention period is addressed in the retention response above.
113	Sec 11 – Evaluation Criteria / AI Governance / Model Audit / Explainability Requirements	We would be grateful if the department could kindly confirm whether any mandatory AI governance, model audit, explainability, or human-approval control requirements apply beyond what is already mentioned in the RFP.	Please refer to RFP Section 7.3 for Technical requirements

114	7.3 - Technical Requirements / Prompt Engineering and Guardrails	Are there any mandatory AI governance requirements, model audit requirements, explainability requirements, or human-approval controls beyond those explicitly stated in the RFP?	
115	Sec 7.3.5(i) – LLM Model Selection / Model Origin / MeitY-Approved Cloud Restrictions	We would be thankful if the department could kindly share formal confirmation on whether AI models of Chinese origin, hosted on MeitY-approved cloud, would be acceptable for use in this solution.	Please refer to RFP Section 7.3.5 Technical Requirements - Large Language Model (LLM) Bidders are expected to mention the planned model/(s) usage in their solution in details. Given the operational criticality of the OCC and the sensitivity of the data involved, the proposed model stack is additionally subjected to a review during design stage and finalized.
116	7.3.5 - Large Language Model / Model, tool, and license constraints	Please confirm acceptance of open-source, commercially licensed, and managed models and specify any mandatory license, indemnity, vulnerability, export-control, or model-card requirements. Please also clarify whether delivery of model weights, fine-tuning adapters, prompts, evaluation sets, hyperparameters, and MLOps pipelines is required where third-party license terms permit.	
117	7.3 (5); 7.4 (5) / LLM / Localization	Are commercially hosted LLM APIs permissible provided India data residency is maintained, or must all models be self-hosted?	
118	Sec 7.3.5 (general model flexibility) / Commercial/Chinese-Origin Model Approval	Following up on the pre-bid discussion — has JICA/DMRC reached a decision on whether commercially available models hosted on MeitY-approved infrastructure (including non-US-origin models) are acceptable for this engagement?	
119	Sec 7.4.3 – Responsiveness / Sub-Second Latency Expectation	We would be thankful if the department could kindly confirm in writing the exact end-to-end latency expectation for the solution, given that a sub-second target was mentioned during the pre-bid meeting alongside the 10 second LLM query response time stated in this section. This would help us plan our infrastructure sizing appropriately.	Push to Talk (PTT) to CDRS file availability latency is sub-second. These files are expected to be transcribed (STT) and presented on the user application in real time. Once the conversations are transcribed and AI has classified the problem, 10 second is the max response time that the LLM should take to give a solution guidance in response to the identified incident.
120	7.4.3 - Responsiveness / Definition of 10-second response time	Please clarify whether the 10-second requirement applies only to LLM inference or to the full path from availability of the final problem-report audio segment through ingestion, STT, NLP, RAG, generation, guardrail validation, and display. Please specify percentile (average/P95/P99), measurement point, payload size, network assumptions, and whether the target applies at 15 concurrent queries.	
121	7.4 (3) / Responsiveness	For the 10-second LLM query response target, please clarify the measurement start point – completion of a PTT segment, detection of the incident / problem boundary, formation of the LLM query, or another event – given upstream CDRS latency.	
122	Sec 7.4.3, 7.3.3 / 10-Second LLM SLA Scope	For the LLM query response time target of 10 seconds (Section 7.4.3), please confirm what is included within this window — does the clock start from raw audio capture, transcript completion, or Part 1/Part 2 boundary detection? Does it include RAG retrieval and the full multi-agent reasoning chain, or only final response generation?	

123	7.4 - Non-functional Requirements	The RFP specifies a response time of 10 seconds. Please clarify whether this is measured from query submission or from the point of incident/problem detection. Kindly confirm whether the 10-second response time includes STT processing latency, or if STT is measured separately.	
124	Sec 7.4.1 – Performance & Scalability / Basis of the 15 Concurrent LLM Queries	We kindly request the department to clarify whether the 15 concurrent LLM queries mentioned refers to the total peak load across the solution, or whether each of the multiple agents (Subject Matter Expert, Ad Hoc Analysis, Explainer) is expected to separately handle up to 15 concurrent queries.	Each agent is expected to handle up to 15 concurrent LLM queries. The figure is therefore per agent rather than a total peak load across the solution.
125	7.4.1 and 7.4.2 - Performance and Availability / Availability and concurrent load	Please define the 99.5 percent availability measurement period, service boundary, planned-maintenance exclusions, dependencies excluded/included, incident severity, and service-credit/acceptance implications. Please also confirm whether 15 concurrent LLM queries represents the peak design load for the PoC and provide expected concurrent audio streams, dashboard sessions, report downloads, and batch workloads.	99.5% availability target applies to the solution components delivered by the Digital Partner, measured monthly across DMRC's OCC operating hours During the PoC, availability is a performance objective reported through monitoring rather than a service-credit-bearing obligation, and no liquidated damages attach to it Refer to answers on data volume on design load considerations
126	7.4 (2) / Availability	Please define the production hours and exclusions used to calculate 99.5% availability, including planned maintenance, source-system outages and DMRC network outages.	
127	7.4.4 - Accuracy / Definition of greater than 85 percent transcription accuracy	Please define the metric for the 85 percent threshold (for example, word accuracy/word error rate, sentence accuracy, or semantic/critical-entity accuracy), whether it is measured separately for Hindi, English, and Hinglish, the required performance under radio noise/interference, treatment of proper nouns and metro terminology, and the independent acceptance corpus and scoring method.	> 85% accuracy requirement is to be measured as word-level accuracy, equivalent to one minus Word Error Rate, computed on a held-out acceptance corpus drawn from Line 2 OCC recordings and representative of normal operating conditions including typical radio interference. The threshold applies to the corpus in aggregate across all languages. The acceptance corpus is to be constructed jointly during the design phase, held out from all training and tuning data, and signed-off by DMRC experts. If >85% threshold cannot be met because of inherent source audio quality like radio interference, code-switching complexity, rather than model or Digital Partner performance, this is not treated as a performance failure. Bidders are expected to demonstrate that the shortfall is attributable to source quality and not to their approach.
128	7.4 (4) / Transcription accuracy metric	How will the >85% transcription accuracy be measured – word accuracy / WER, sentence accuracy, critical-entity accuracy or another methodology – and against what test set, adjudicated by whom?	
129	7.4 (4) / Accuracy by language / noise	Must the STT threshold be achieved separately for Hindi, English, Hinglish and each defined noise category?	
130	Sec 7.4.4 / STT Accuracy	How will the 85% transcription accuracy target (Section 7.4.4) be measured — Word Error Rate, Character Error Rate, or task-level effectiveness — and against what reference/golden dataset? Please also clarify accountability if this target proves unachievable due to inherent source audio quality (radio interference, code-switching complexity) rather than model/vendor performance.	

131	7.4 - Non-functional Requirements	The RFP specifies transcription accuracy >85%. Please clarify the evaluation methodology, test dataset, and language mix (Hindi/English/Hinglish). Kindly confirm whether accuracy will be measured using WER, CER, or another operational performance metric.	
132	7.3 (2)(iii) / Speaker identification	Please provide the metric and threshold for correct Radio ID-to-role mapping and speaker diarization.	Radio ID-to-role mapping in DMRC's environment is based on a static, pre-assigned configuration (e.g., Radio xx = TC #, Radio yy = Station Controller #)
133	7.1.2 (i) / Incident extraction accuracy	Please specify acceptance thresholds for incident detection, incident classification and extraction of train ID, station, section, asset, fault and other PRA fields.	Bidders should propose thresholds for each, together with the measurement method, evaluated against the SME-adjudicated acceptance corpus. A distinction should be drawn between fields that populate a incident record and are therefore operationally consequential, and fields that are informational. Thresholds will be agreed and fixed with DMRC during the design phase, and recorded in the acceptance criteria. Where a proposed threshold depends on data quality that cannot be verified at bid stage, bidders should state the dependency/assumptions made
134	7.1.2 (ii); 7.3 (3) / LLM guidance quality	How will groundedness, citation correctness, ranking relevance, confidence scores, hallucination prevention and “top three” solution hit rate be measured and accepted?	Bidders should propose the metric definitions, the construction of the evaluation set and the target levels, which will be agreed with DMRC during design stage. Once the solution is operational, the Traffic Controller rating captured through the feedback loop provides the in-life measure and feeds model improvement, consistent with the self-learning requirement.
135	7.5 / Benchmark dataset	Who will create and approve the benchmark dataset, what is its approximate size, and will it be held separately from training and fine-tuning data?	The benchmark and acceptance dataset is to be constructed jointly during the design phase, drawn from Line 2 recordings and PRA records, and must be held out entirely from training, fine-tuning and prompt-development data. Selection, transcription and labelling of the candidate corpus sits with the Digital Partner within its data-preparation scope. Review & formal approval will be required from DMRC. The dataset must be large enough for statistically meaningful measurement across the language and noise categories in scope, the size to be agreed during the design phase and recorded in the test plan.
136	7.5 / Rare incident testing	May rare / high-severity incidents be tested through historical replay or simulation when they do not occur during the PoC?	The test approach & scenario set is to be agreed during the design phase and recorded in the test plan. It should cover comprehensive test scenarios in scope, including single-line operation and short-looping which could be rare. Incident types that are rare or high-severity may be evaluated through replay of historical recordings and PRA records, and through scripted simulation agreed with DMRC subject-matter experts
137	8 / Final report – KPIs	What are the predetermined KPIs for PoC evaluation referenced in the Final Report deliverable, and when will they be agreed?	Project success will be measured by the defined objectives laid out in the RFP Section 2 and achievement of functional & non-functional requirements.

138	8 - Deliverables, Milestones and Timelines	Please define the specific KPIs and success criteria that will determine successful completion and acceptance of the PoC across STT accuracy, LLM response quality, response latency, integration readiness, user adoption and operational impact. For AI-generated recommendations and traffic-management guidance, please also clarify the authoritative basis against which recommendation correctness will be evaluated (e.g., SOP/manual compliance, SME review, Chief Controller validation, historical PRA outcomes, or a combination of those).	The KPIs to be reporting in analytical dashboards are defined in the RFP section 7.1.6 Functional requirement - Front End Dashboards (Supervisory dashboard & analytics, System performance) The complete KPI set, with measurement definitions, measurement points and target levels, is to be agreed with DMRC during the design phase For Authoritative basis, please refer to an earlier answered question
139	7.4 (8) / Infrastructure sizing	For infrastructure sizing: please confirm the number of Line 2 talk groups / channels to be continuously monitored, the peak number of simultaneous conversations / PTT files to be processed, and typical daily conversation and incident volumes.	For Line 2: there are around 3 radio talk groups generating approximately 10,000 individual push-to-talk recordings per day. Between 20-30% of conversations relate to incidents requiring resolution guidance. Bidders should state the sizing assumptions adopted and the headroom provided in the solution
140	Sec 7.5 – Testing / Security & Penetration Testing Cost Ownership	We kindly request the department to confirm whether security and penetration testing would be conducted by a DMRC/CERT-In empanelled third party at the department's cost, or whether engaging such a party would be within our scope and cost.	Security and penetration testing are within the Digital Partner's scope and cost and must be included within the USD 300,000 ceiling. The bidder may engage any CERT-In empanelled auditor; no specific agency is designated. Findings are to be remediated by the Digital Partner and re-tested before Go-Live. This is separate from, and additional to, DMRC IT's own review and approval of the BRD, TRD, solution designs and deployment, which DMRC conducts through its stage-gate process at no cost to the Digital Partner
141	7.4 - Non-functional Requirements / 7.5 - Testing	Please confirm the specific version/edition of the CERT-In/MeitY Guidelines on Information Security Practices for Government Entities applicable to this project. Kindly clarify whether DMRC/JICA mandates the use of a designated security auditor/penetration testing agency, or if the Bidder may engage any CERT-In empanelled auditor for compliance assessment and security testing.	
142	Sec 5 – Location of Work / On-site Co-location Requirement	With reference to the co-location requirement mentioned in this section, we would be grateful if the department could kindly confirm the expected on-site versus remote working model (and the roles/duration involved) for the PoC duration. This will help us plan our resourcing thoughtfully.	Digital partner can function on a hybrid model. However the key team members (project manager, business analysts, etc) are required to be co-located for key project phases/ milestones. Bidders are expected to clealy communicate the phase wise staffing plan in the solution and highlight the roles planned on-site and the duration.

143	11.2 - Technical Evaluation / 100 percent availability of named experts	The RFP requires 100 percent availability of named experts from Day 1 but also asks for a phase-wise resource plan. Please clarify which roles must be full-time for the entire PoC, whether 100 percent means exclusive allocation or availability when required, permitted onsite/remote mix, minimum allocation by phase, and the approval timeline for replacement due to attrition, illness, or force majeure.	100% availability means that the named experts are committed to this engagement and available when required by the phase-wise plan, and that they will not be substituted or diverted to other engagements without prior approval. It does not require exclusive full-time allocation of every named role across the project duration. The bid should state, for each named expert, the allocation by phase as a percentage of full time, and the roles proposed as full-time throughout. The Project Lead is expected to be allocated substantially full-time for the duration. Onsite and remote mix follows the hybrid working model above. Replacement of a named expert on account of attrition, illness or force majeure requires prior written approval, with a substitute of equivalent or better qualification and experience proposed within an agreed period, ordinarily two weeks.
144	7.2 - System Integrations	If the required integration APIs are unavailable or access is restricted by OEM providers, will alternative integration mechanisms such as database views, file exports, message queues, or scheduled data extracts be accepted for PoC implementation? In such cases, please confirm the preferred fallback integration approach and any associated limitations or constraints.	The preferred approach is to build an integrated solution. DMRC is facilitating the discussions with OEM vendors for integration readiness. In-case of integration challenges, fall back options to be evaluated on a case-to-case basis, jointly agreed with the concerned parties including the digital partner
145	7.1 / 7.3 - Traffic Management Guidance / Optimization Models	Please clarify the authoritative rule hierarchy and availability of real-time operational data for traffic-management recommendations and confirm whether the system should withhold guidance and escalate to the Traffic Controller whenever the available information is missing, conflicting or unsafe. Example: If a train fails at Rajiv Chowk, historical records may suggest short-looping trains at a nearby crossover. However, if the crossover is unavailable or a newer safety circular prohibits that movement, should the system reject the historical solution, display the applicable restriction and request the TC's decision?	The planned solution should consider documented SOPs/manuals as authoritative rule hierarchy. If there are changes in operating conditions, circulars are updated and all concerned parties including TCs are kept updated on the same. The detailed solution blueprint should document all design decisions and get it signed-off.
146	Sections 2 and 7.1 - Objectives and Functional Requirements	Please define the minimum set of incident categories, fault types, asset classes, Line 2 sections and talk groups that must be supported for PoC acceptance. Is comprehensive coverage of every Line 2 incident expected?	A comprehensive coverage of every incident type occurring on Line 2 is expected for the PoC.

147	Sections 7.1 and 7.3 - Incident Identification and Problem Extraction	At what point should the application create an incident: the first relevant radio segment, completion of the problem report, AI classification, TC confirmation or another operational trigger? Which incident classes or recommendation types require approval from the Chief Controller or an asset-specific controller before the guidance is acted upon?	The system initiates incident resolution procedure once the incident has been identified based on live radio conversation between the TC and the TO. PRA writeback is triggered once the incident resolution guidance has been accepted by the Traffic Controller Traffic Controller acts on guidance for incidents as per defined SOPs/ manuals. Incident classes involving asset-specific decisions - rolling stock, signalling, traction require the relevant asset controller's input. Chief Controller involvement applies to critical incidents and to traffic-management decisions such as single-line operation and short-looping. All these functionaly/ design decisions should be finalized during the project design phase, documented and signed-off
148	Section 7.1 - Feedback Loop	What precisely does 'Accept' mean: the TC has reviewed the recommendation, selected it, communicated it, executed it, or confirmed successful incident resolution? What happens in a multi-incident case?	Accept' means the Traffic Controller has reviewed the recommendation and adopted it as the course of action to be communicated. It is not a confirmation that the incident has been resolved, and it is not an execution record. Acceptance of the recommendation provides the relevance signal for model feedback only. On concurrent incidents: the Traffic Controller may be handling several incidents at once. Each recommendation, acceptance and rejection must be bound to a specific incident identifier, and the front end must make clear which incident an action applies to. Acceptance of guidance for one incident carries no implication for any other.
149	Sections 7.1, 7.3 and 7.4 - Guidance, Guardrails and Safety	Are there incident types, emergency conditions or safety-critical topics for which the AI must not provide procedural guidance and must instead direct immediate human escalation?	The solution planned is expected to work as a decision support system to the Train Controller, where the incidents in metro operations are resolved as per the defined guidelines (SOPs, Manuals, Circulars) with required escalation to concerned authorities (internal/ external) The scope of guidance is limited to available/documented DMRC knowledge repository only.

150	Sections 7.1 and 7.3 - Radio Monitoring and Conversation Segmentation	Can a single incident involve multiple conversations, talk groups, or radio channels? Conversely, can a single conversation contain multiple unrelated incidents? Please clarify the expected correlation and incident-to-conversation mapping rules.	Both are possible and the design must accommodate a many-to-many relationship rather than assuming one conversation equals one incident. An incident will commonly have multiple to-and-fro conversations on radio between TC and TO . As radio channels/ talk groups are generally created per section, the same incident may span multiple sections/talk groups as well. For example- a door closing incident may be resolved partially and multiple conversations may happen in its operation across the radio talk groups. Conversely a conversation may reference more than one incident e.g- two different asset failures may happen and TO/TC conversations may be centered around resolving these incidents
151	Sections 7.1 and 7.3 - Confidence Score and Output Guardrails	Please clarify what the displayed confidence score represents: transcription accuracy, incident classification confidence, retrieval relevance, source consistency, recommendation completeness, or a composite metric derived from multiple factors.	Confidence score in the context of solution recommendation represents retreival relevance, to measure the accurate retreivel of solution from knowledge repository
152	8 - Deliverables, Milestones and Timelines / Program governance	Will DMRC nominate a single empowered Data point-of-contact (SPOC) for the PoC, authorised to approve data and system access, confirm the knowledge-base document set, and resolve data-related queries, and can DMRC indicate the expected turnaround for access requests and clarifications during the build?	The project team will include a DMRC POC for required support on data and system access requirements. Considering the strategic importance, DMRC support will be available for clarification, data requests. Digital Partner should actively work with the project governance/escalation channels to highlight any challenges/delays in the expected support
153	7.4 - Non-functional Requirements (Data Security) / 7.3 - KB Prep, Storage & Audit	OCC radio audio, the resulting transcripts, and PRA records contain personally identifiable information (for example staff names and radio IDs) and security-relevant operational detail. For the PoC, is the Digital Partner expected to implement de-identification / masking of PII as a processing step within the solution, or will DMRC provide data already de-identified where required? This determines whether a masking/redaction stage is built into the ingestion and transcription pipeline.	DMRC will provide data as held, without prior de-identification. Where masking /redaction is required, implementing it as a processing stage within the ingestion and transcription pipeline is within the Digital Partner's scope

154	Sec 7.4.5, 7.4.7 / Model Hosting vs. Data Localization	Given the data localization requirement (Section 7.4.5) mandates data storage within India, does this also require the LLM/STT inference endpoint itself to be hosted within an India region, or does the requirement apply only to data at rest?	The data localization requirement under Section 7.4.5 is intended to apply to data throughout its lifecycle including data in processing (e.g., inference) and not solely to data at rest. The expectation is that the LLM/STT inference endpoint itself be hosted within an India cloud region, such that voice/text data does not leave Indian jurisdiction at any stage of capture, processing, or storage
155	7.3 - Storage & Audit Trail / 7.4 - Localization, Data Security / 10 - IP and Other Considerations	Data Residency & Sovereignty Requirements. Please confirm whether all PoC data, including CDRS audio recordings, transcripts, PRA records, vector databases, audit logs, model artifacts, backups and monitoring data must remain within India. Are there any restrictions on cross-border data storage, processing, remote support access or transfer of operational data outside India during the PoC? Please clarify whether cloud-managed AI services hosted outside India are permitted for model inference, training or support activities.	