TECHNICAL SPECIFICATION
For the Open Tender

'Cloud-based software development, including database, analytics, and other functionality'

(Identification number LMT/BRAVO/2026-5**)**

The applicant shall submit a technical proposal in accordance with the procurement announced by AS LabMedTech in an open tender, "Cloud-based software development, including a database, analysis and other functionality" (identification number LMT/BRAVO/2026-5) the requirements of the Regulations and Technical Specifications, describing in detail in the Technical Proposal the fulfillment, provision and implementation of all requirements, including all activities related to development, delivery, and warranty technical support, which will result in the implementation of the Procurement Object. The description must be sufficiently detailed to convince the Customer of the Applicant's ability to fulfill all the requirements set out in the Regulations and Technical Specifications.

Technical specifications

#	Development/Delivery/Warranty technical support*	Quantity	Unit	Customer requirements**
1.	Provide a cloud-based platform (develop a new platform or adapt an existing platform) as a service for running medical trials on breath analysis	1	Platform	The platform must be hosted in the EU and meet other GDPR (and HIPAA) requirements. The platform must be able to receive data, perform real-time or near real-time data analysis, quality control with anomaly detection, and present user interface for researchers/doctors to interact with the data and results. The platform must implement appropriate audit logging and monitoring. The platform must provide drift compensation, normalization of signal data according to environment changes, as well as multi-instrument harmonization. The platform must be implemented as a cloud-native, multi-tenant web application on AWS or equivalent platform, using Amazon ECS on AWS Fargate or Amazon EKS (or equivalent services from different provider) for container orchestration, with all services containerized and stored in Amazon ECR or equivalent service. Core transactional and configuration data shall be stored in Amazon Aurora PostgreSQL (encrypted with AWS KMS), while high-volume sensor and log data is stored in Amazon S3 and, where time-series queries are required, in Amazon Timestream. All traffic shall be encrypted using TLS termination on

				Application Load Balancer (ALB) or Amazon (or similar) API Gateway, and data at rest shall use KMS-managed encryption in S3, Aurora (or equivalent solutions), and any attached volumes. Authentication and authorization shall rely on Amazon Cognito or AWS IAM Identity Center (or similar solutions), with role-based access control enforced on APIs exposed through API Gateway. Logging and monitoring shall use Amazon CloudWatch Logs, CloudWatch Metrics/Alarms, and AWS X-Ray or equivalent solutions, while compliance, security posture, and auditability are strengthened via AWS CloudTrail, AWS Config, AWS Security Hub, Amazon GuardDuty, and AWS Audit Manager for continuous controls and evidence collection (or similar solutions). The platform must be made available to the Customer after the project, also hosted on Customer's server (as a dockerized container or equal). The final platform should be implemented in a way that allows the Customer to use the data for model training and add their data processing and analysis steps/algorithms/methods implemented in Python and its common libraries within the time limit specified in the Regulations.
1.1.	Secure, automated data exchange	1	Module	Providing data transfer functionality with necessary data encryption during transfer (at least AES-256 encryption and TLS 1.3), loss-resilient transmission of data and storage at necessary security levels for measurement data (including patient and potentially sensitive healthcare data), instrument data and calibration data. Given the nature of the data stored in the system, the Supplier must provide incident response plans according to GDPR/HIPAA. The secure data exchange module shall expose RESTful endpoints and event-based integrations via Amazon API Gateway and Amazon EventBridge (or comparable solutions from other providers), with asynchronous delivery and buffering handled by Amazon SQS or Amazon SNS (or other solutions with similar functionality) where required for decoupling and reliability. For secure file-based transfers (e.g., from medical devices or partner systems), the platform shall use AWS Transfer Family or similar solutions (SFTP/FTPS/FTP) writing directly into S3 buckets encrypted with SSE-KMS keys or equivalent solutions. All endpoints shall enforce TLS in transit and KMS (or similar)-based encryption at rest, with credentials, API keys, and certificates stored in AWS Secrets Manager or AWS Systems Manager

				Parameter Store (SecureString) or equivalent solution. Access control for the exchange APIs shall be managed through Amazon Cognito or IAM authorizers or equivalent solutions integrated with API Gateway, and all access, data movement, and configuration changes shall be recorded via AWS CloudTrail and surfaced in AWS Security Hub, with threat detection from Amazon GuardDuty and optional investigation support through Amazon Detective or solutions with similar functionality to enable automated, compliant incident detection and response workflows aligned with healthcare and GDPR requirements.
1.2.	Data preprocessing, scaling, transformation, database	1	Module	Module implementing time-series and other data type preprocessing for electronic nose sensor data analysis, including: - measurement stage and breath phase detection and visualization of sensor baselines and responses, - appropriate normalization and calibration to compensate for drift due to temperature, humidity and pressure changes, sensor aging, including statistical validation for calibration, - harmonization among different instruments, - detection of erroneous signal (e.g. excessive peaks, missing signal and noise), - preprocessing of the signal, including feature extraction (delta normalized against absolute values, peaks and plateaus, breath fingerprints and other breath signal characteristics). The preprocessing and transformation layer shall be implemented using AWS Glue jobs or equivalent (Python/Spark) and/or AWS Lambda (or equivalent) functions orchestrated by AWS Step Functions or Amazon Managed Workflows for Apache Airflow (MWAA), operating over a data lake built on Amazon S3 with metadata managed by the AWS Glue Data Catalog or equivalent solutions from other providers. Sensor time-series and processed feature sets shall be stored either directly in Amazon Aurora PostgreSQL (for structured and relational queries) or in Amazon Timestream (for optimized time-series workloads) or similar solutions, while staged and historical data remains in S3 as encrypted Parquet or CSV. Data access, schema evolution, and harmonization logic shall be restricted and audited via AWS Lake Formation or equivalent, ensuring that different projects, sites, and roles see only permitted data. Batch and near-real-time transformations shall be monitored in Amazon CloudWatch or equivalent, with configuration drift and resource compliance tracked by AWS

				Config or equivalent, and sensitive data discovery in raw/preprocessed layers can be supported by Amazon Macie or equivalent solution to detect and protect personal data.
1.3.	Data analysis and model development	1	Module	A module implementing mathematical, statistical and machine learning based analysis of breath data (neural networks etc.), generating a result of prediction and group membership (disease and other characteristics), taking into account drift compensation and other breath data specifics. The module must also provide analysis of features and their effects across measurements and cohorts. The analysis and model development module shall be built primarily on Amazon SageMaker or an equivalent platform, using it as the central environment for data exploration, model development, and collaboration. Training jobs shall run as managed Training jobs using built-in algorithms or custom containers, with input data drawn from Amazon S3, Amazon Aurora, or Amazon Redshift or equivalent solutions as needed. Experiment tracking, model versioning, and lineage shall rely on SageMaker Experiments and SageMaker Model Registry, while data processing and feature engineering shall use SageMaker Processing jobs. Model explainability and bias analysis for medical use cases shall be supported by SageMaker Clarify, and pipelines for repeatable ML workflows shall be orchestrated using SageMaker Pipelines, or solutions with similar functionality can be used. All underlying infrastructure (S3, ECR, KMS, IAM and similar) remains fully managed by AWS or similar cloud service platform, with logs and metrics routed to Amazon CloudWatch and compliance/control evidence forwarded to AWS Security Hub and AWS Audit Manager or equivalent solutions where applicable.
1.4.	Model deployment for use by researchers and healthcare staff	1	Module	A module running the developed model(-s) on the platform, ensuring conformity with GDPR and HIPAA, scalability and deployment flexibility. The implementation of analysis models should be easy to use and provide a dashboard for researchers and healthcare specialists to analyze and visualize different cohorts. The model deployment module shall use Amazon SageMaker Endpoints (including Serverless Inference or Multi-Model Endpoints or equivalent solutions from other providers where appropriate) to host and scale approved models for clinical and research use, fronted optionally by Amazon API Gateway for consistent API management and throttling. Network

				access to endpoints shall be restricted using VPC endpoints, Security Groups, and optionally PrivateLink or equivalent solution to ensure that clinical trial systems and hospital networks connect over controlled, private paths. Authentication and fine-grained authorization to model inference APIs shall be enforced through Amazon Cognito, IAM-based or similar authorizers, and all invocations shall be logged with full traceability using Amazon CloudWatch Logs and AWS CloudTrail or equivalent solutions, with compliance posture monitored in AWS Security Hub or similar solution. Operational metrics (latency, error rates, throughput) shall be tracked via a Model Monitor and CloudWatch or similar solution, while model monitoring should also be configured to detect data and model drift, enabling controlled updates and safe model promotion workflows aligned with regulatory expectations.
1.5.	Interface for data browsing, research	1	Module	A module for doctors and researchers to interact with data (including aggregated data, metadata) and results. Role-based data access to aggregated data browsing, conforming to GDPR/HIPAA, including audit logs of activities involving data access. Generation of reports on patient data and statistics. Visualization of the results and data aggregation based on group membership (e.g., sex, disease etc.) and measurements. The data browsing and research interface shall be delivered as a secure web application hosted using AWS Amplify Hosting or as a static SPA on Amazon S3 fronted by Amazon CloudFront or equivalent solutions for global, low-latency access, with AWS WAF or similar services protecting against common web attacks. The frontend shall consume backend services exposed through API Gateway (REST or GraphQL via AWS AppSync or similar solution), which connect to data sources such as Amazon Aurora PostgreSQL, Amazon Timestream, Amazon Athena, Amazon S3 or similar. Cohort-level dashboards, exploratory analysis, and aggregated visualizations may be implemented directly in the web UI or complemented by Amazon QuickSight or similar dashboards embedded into the application for interactive analytics and reporting. Authentication for researchers, clinicians, and administrators shall rely on Amazon Cognito, IAM Identity Center or equivalent solutions, with role-based and attribute-based permissions enforced by the APIs and data-layer policies (Aurora, Lake Formation or similar). All UI-driven actions that access or export data shall be

				logged through API Gateway and CloudTrail into CloudWatch Logs and surfaced to AWS Security Hub and AWS Audit Manager or equivalent solutions for evidencing role-appropriate access and compliance.
1.6.	Platform module for patient meta-data collection and consent management	1	Module	Provides functionality to enter, store, edit and browse participant/patient meta data (including consent management), providing security and privacy level that conforms with GDPR and HIPAA. Patient metadata, enrollment, and consent records shall be stored in PostgreSQL (or optionally Amazon DynamoDB or similar data base for high-scale key–value access), with encryption at rest using AWS KMS or equivalent solution and strict network controls via VPC subnets and security groups. The application layer managing patient and consent logic shall be implemented as microservices on Amazon ECS Fargate or AWS Lambda functions behind Amazon API Gateway or equivalent solutions, with authentication through Amazon Cognito and fine-grained authorization using IAM or similar solutions and, where applicable, data-access controls via Lake Formation or similar. Audit trails for all create/update/delete operations, as well as consent changes and data-subject-rights actions (access, rectification, restriction, deletion), shall be captured via AWS CloudTrail and application logs in CloudWatch Logs, then aggregated and assessed in AWS Security Hub and AWS Audit Manager (using HIPAA/GDPR-aligned frameworks for control evidence) or equivalent solution. Discovery and monitoring of sensitive personal data in backups, S3 exports, and derived datasets shall be supported by Amazon Macie or similar solution, ensuring that patient-identifiable information remains properly protected, classified, and limited to authorized use only.
1.7.	Adapt the platform to work with the new data modalities (analysis etc.)	1	Updated platform	The initial platform must be updated to accommodate additional data modalities (MOX sensor readings supplemented by gas chromatography) that will be introduced after the breath analysis (electronic nose) device will be enhanced, including data preprocessing, analysis, aggregation and other functionality of the platform. The extended functionality should provide at least GC data ingestion, peak detection, VOC matching, and these data should be analyzed alone and cross-analyzed with sensor data, including reports and visualizations on

				patient or patient group levels. The updated platform must allow harmonization and pooling of respective data from both instruments (with and without GC). To extend the platform for combined eNose and gas chromatography (GC) data, the ingestion and processing pipelines shall be implemented as AWS Glue ETL jobs, AWS Lambda functions, or SageMaker Processing jobs (or equivalent solutions in other platforms) operating over GC output stored in Amazon S3, catalogued with the AWS Glue Data Catalog, and governed via AWS Lake Formation, or similar solutions. Time-aligned features and chromatographic peaks shall be persisted in PostgreSQL and/or Amazon Timestream or equivalent, while large-scale multi-modal feature sets are stored in S3 as Parquet (or equivalent) and queried using solutions like Amazon Athena or Amazon Redshift or similar where high-performance analytics is required. Multi-modal modeling (combining GC-derived VOC features with MOX sensor features) shall reuse the cloud-based ML platform (like Amazon SageMaker) stack (Training, Processing, Experiments, Pipelines, and Model Registry) so that single- and multi-modal models follow the same MLOps lifecycle. Integrated visual analytics for chromatograms, VOC profiles, and sensor fingerprints may be exposed via APIs (Athena/Redshift queries, Aurora SQL or similar) and rendered in the web UI or in embedded dashboards (e.g., Amazon QuickSight), with access controlled by Cognito/Identity Center or an equivalent solution and all usage monitored and evidenced via CloudTrail, Security Hub, and Audit Manager or tools with similar functionality, to maintain a consistent compliance posture across all data modalities.