Due Diligence Questionnaire - AI/ML Capabilities

Stand 18. Juni 2026

BUNDESVERBAND MERGERS & ACQUISITIONS GEM. E.V.

AI/ML Components

If you have several products and there are differences between the products, please describe your AI/ML capabilities per product.

AI/ML Capabilities Strategy

Question	Response
Which AI tasks are addressed by the solution?	e.g., classification, NLP, document processing, image processing, reinforcement learning, anomaly detection, predictive analytics, etc.
Where are AI capabilities used in the product portfolio?	Please provide information about specific products or modules utilizing AI capabilities.
Which AI technologies are currently implemented and actively used?	List specific technologies, models, and approaches currently in production.
How do AI capabilities support business objectives and customer needs?	Describe alignment between AI features and value creation.

AI/ML Frameworks & Tools

Question	Response
What AI/ML frameworks are utilized?	e.g., LangChain, LangGraph, TensorFlow, PyTorch, Scikit-learn, Keras, Hugging Face Transformers, etc.
Are there any proprietary AI tools?	Please describe any internally developed AI tools or frameworks.
Which cloud providers are used for AI-related workloads?	e.g., Azure AI, AWS (SageMaker, Bedrock), Google Cloud (Vertex AI), others
How are AI components integrated with other enterprise software solutions?	Describe integration architecture, APIs, and middleware.

Legal & Data Handling

Question	Response
What types of data are used for training AI models?	structured, unstructured, proprietary, third-party, synthetic, etc.
How does the data ingestion process work?	e.g., collection, cleaning, transformation, integration, sampling, splitting, validation
Please provide information regarding data provenance and data lineage.	Documented trail accounting for the origin of data and its movement through systems.
Where is data stored and for how long?	e.g., database, cold storage, data lake, retention policies
Does the customer have the option of choosing one or more data storage locations?	Geographic and infrastructure options for data residency.
How is data ingested?	e.g., via REST API, 3rd party services, batch uploads, streaming
Is any solution for data synchronization in place?	Real-time vs. batch synchronization mechanisms.

Question	Response
How is the data isolated?	Multi-tenancy architecture and isolation mechanisms.
Is data encryption supported?	At rest and in transit encryption methods.
Is it possible to use a customer-provided key to encrypt the data?	BYOK (Bring Your Own Key) capabilities.
Is there any way to access customer data without the consent of the customer?	Emergency access procedures and legal requirements.
Is any special contract signed by the customer regarding data handling?	DPA, BAA, or other data processing agreements.
Is there any access control on customer data?	RBAC, ABAC, and access logging mechanisms.
Can the data be used for improving the models?	Opt-in/opt-out mechanisms for model improvement.
In case access to customer data is allowed, how is this controlled when access is revoked?	Data deletion, access revocation, and audit procedures.

AI/ML Model Development

Question	Response
Are foundation models integrated into products?	GPT, LLaMA, Claude, Gemini, Mistral, or others
Please share details of your AI models.	Are these models proprietary, third-party, or FOSS (Free and Open Source Software)?
How have the models improved?	What fine-tuning or adaptation strategies are used for LLMs in proprietary applications?
Do improvements rely on buying data from providers or customer data?	Data sourcing strategy for model enhancement.
Are AI models retrained periodically?	Retraining frequency and triggers for model updates.
How is model performance monitored?	Metrics, dashboards, alerting systems for model drift and degradation.

Question	Response
Which infrastructure is used by data scientists during experimentation?	e.g., cloud tooling, hyperscaler-specific infrastructure, local physical machines, GPUs, TPUs
What is the process of releasing a model from experimentation to production?	CI/CD pipelines, approval workflows, staging environments.
Which unit, integration, and performance tests are you running?	Testing frameworks and coverage for AI models.
How is the licensing of dependent libraries assessed?	License compliance scanning and management.
How are vulnerabilities of dependent libraries assessed?	Security scanning tools and vulnerability management processes.
Where is the final artifact stored?	docker image, python package, model registry location
Please provide a list of used Machine Learning libraries and tools.	Comprehensive inventory of ML dependencies.
Please provide information about ownership and licenses.	IP ownership and licensing structure for models and code.
Describe the pipelines used during training and model execution.	MLOps pipeline architecture and orchestration.

AI/ML Model Execution

Question	Response
Which ML infrastructure is used?	Is it hyperscaler agnostic (K8s, Kubeflow, MLflow, Databricks) or hyperscaler specific (AWS SageMaker, Azure ML, GCP Vertex AI)?
Are you using any specific ML services provided by hyperscalers?	OCR, Image Detection, Sentiment Analysis, Speech Recognition, others
What measures are in place for scaling AI workloads efficiently?	Auto-scaling, load balancing, resource optimization.
How do AI models interact with existing software architecture?	microservices, monolithic, distributed systems, event-driven

AI/ML Lifecycle Management

Question	Response
What is your pricing model?	Per-inference, subscription, usage-based, tiered pricing
How frequently can customers train models?	Training frequency limits and quotas.
Do customers pay to train a model?	Training cost structure and billing.
What happens if a customer creates a script that trains models consecutively?	Rate limiting, throttling, or abuse prevention measures.
What happens to models running but not used for inference for more than X hours?	Idle model management and resource cleanup policies.
Is scalability in place?	Horizontal and vertical scaling capabilities.
Is it possible to scale to zero (serverless)?	Serverless capabilities and cold start implications.
What is the loading time for CPU and GPU models?	Model loading latency and optimization strategies.
How do you handle model versioning and updates?	Version control, A/B testing, canary deployments.
How many models are supported in parallel?	Concurrent model execution limits per customer.
What happens if all customers try to train/infer simultaneously?	Resource allocation, queueing, and fair-use policies.

AI/ML Benchmark & Performance

Question	Response
Please share how AI performance is measured.	Include figures, metrics, benchmarks on internal or public datasets.
Which languages are supported by each service?	Natural language support across AI capabilities.
What is the effort to add a new language?	Localization and language extension processes.
Do you support specific business domains?	e.g., HR, CRM, Finance, Healthcare, Legal
What is the customer-specific effort required?	e.g., fine-tuning, training data preparation, customization
Please describe your sensitive data masking process.	PII detection, redaction, tokenization, anonymization.
Please provide documentation on scaling constraints.	Which scaling vectors are relevant, where are bottlenecks and inflection points, what mitigation strategies exist.
Do you require any special hardware to run your AI capabilities?	During training and in production (GPUs, TPUs, specialized accelerators).

Agentic AI & Autonomous Systems

Agentic AI Capabilities & Architecture

Question	Response
Are agentic AI capabilities integrated into your product portfolio?	Please specify which products and use cases.
What level of autonomy do your AI systems exhibit?	e.g., fully autonomous, semi-autonomous, human-in-the-loop, conditional autonomy
How do agentic AI models make decisions?	Describe decision-making frameworks, reasoning mechanisms, and planning approaches.
What oversight mechanisms exist for agentic AI decisions?	Human review, approval workflows, guardrails, circuit breakers.
What frameworks or tools are used to develop autonomous AI agents?	e.g., LangGraph, AutoGPT, CrewAI, AutoGen, OpenAI Agents, others

Multi-Agent Systems & Orchestration

Question	Response
Do you implement multi-agent systems?	Multiple specialized agents working together or competing.
How are multiple agents orchestrated and coordinated?	Orchestration frameworks, communication protocols, coordination patterns.
What specialized roles do different agents fulfill?	e.g., planner, executor, critic, researcher, validator, coordinator
How do agents communicate and share information?	Message passing, shared state, event-driven, API-based communication.
What conflict resolution mechanisms exist when agents disagree?	Voting, hierarchical decision-making, consensus algorithms.
How is task decomposition and allocation handled?	Breaking down complex tasks and assigning to specialized agents.
What mechanisms prevent agent deadlocks or infinite loops?	Timeout policies, maximum iteration limits, circuit breakers.

Agent Planning, Reasoning & Tool Use

Question	Response
What planning algorithms do your agents use?	e.g., ReAct, Chain-of-Thought, Tree-of-Thoughts, Plan-and-Execute
How do agents handle long-horizon tasks?	Tasks requiring multiple steps over extended time periods.
What reasoning techniques are employed?	Logical reasoning, causal reasoning, analogical reasoning, probabilistic reasoning.
How do agents decide which tools to use?	Tool selection mechanisms and decision logic.
What external tools and APIs can agents access?	Web search, databases, calculators, code execution, external services.
How are tool calls validated and secured?	Parameter validation, authentication, authorization, sandboxing.
What happens when a tool call fails or returns errors?	Error handling, retry logic, fallback mechanisms, graceful degradation.
How do agents learn from tool usage outcomes?	Feedback loops, reinforcement learning, experience replay.

Agent Memory & State Management

Question	Response
What types of memory do your agents utilize?	Short-term (working), long-term (episodic), semantic memory.
How is agent state persisted across sessions?	State management, checkpointing, session recovery.
What memory retrieval mechanisms are employed?	Vector search, semantic retrieval, temporal ordering, relevance ranking.
How do agents maintain context over long conversations?	Context window management, summarization, memory compression.
Are agent memories isolated per customer or shared?	Multi-tenancy and memory isolation architecture.
How is sensitive information in agent memory protected?	Encryption, access controls, retention policies, secure deletion.

Agent Autonomy Levels & Boundaries

Question	Response
What autonomy levels do you support?	Level 1 (basic responder), Level 2 (router), Level 3 (tool-calling), Level 4 (multi-agent), Level 5 (fully autonomous)
What boundaries and constraints limit agent actions?	Action whitelists, approval requirements, resource limits, domain restrictions.
How do agents escalate to human oversight?	Escalation triggers, confidence thresholds, complexity detection.
Can customers configure autonomy levels?	Customizable autonomy settings and guardrails.
What actions require explicit human approval?	Critical decisions, financial transactions, data modifications, external communications.
How are agent capabilities and permissions scoped?	Role-based permissions, capability-based security, least-privilege access.

Agent Learning & Adaptation

Question	Response
Do agents improve through continuous learning?	Online learning, reinforcement learning from human feedback (RLHF).
How do agents adapt to new tasks or domains?	Transfer learning, few-shot learning, in-context learning.
What feedback mechanisms enable agent improvement?	User ratings, explicit feedback, implicit signals, outcome tracking.
How is agent learning validated before deployment?	Testing, validation sets, human evaluation, A/B testing.
Are there mechanisms to prevent negative learning?	Safeguards against learning from adversarial inputs or edge cases.
How do you handle model drift in agentic systems?	Drift detection, monitoring, retraining triggers, version rollback.

Agent Observability & Monitoring

Question	Response
How are agent actions tracked and logged?	Comprehensive audit trails, action logging, decision traces.
What observability tools are used for agentic systems?	LangSmith, Weights & Biases, custom dashboards, tracing frameworks.
How do you monitor agent performance and reliability?	Success rates, task completion times, error rates, SLA compliance.
What metrics are tracked for agent evaluation?	Accuracy, relevance, consistency, efficiency, cost per task.
How are agent failures detected and diagnosed?	Anomaly detection, error pattern analysis, root cause analysis.
What alerting mechanisms exist for agent anomalies?	Real-time alerts, threshold-based notifications, incident management.
How is agent behavior made explainable?	Reasoning traces, decision explanations, chain-of-thought visibility.

Agent Safety, Alignment & Governance

Question	Response
How do you ensure agents remain aligned with human goals?	Value alignment techniques, objective specification, goal validation.
What safety mechanisms prevent harmful agent actions?	Content filters, action validators, ethical guardrails, safety classifiers.
How are agent risks assessed and managed?	Risk frameworks, threat modeling, scenario testing, red-teaming.
What governance frameworks are in place for agentic AI?	Policies, approval processes, oversight committees, compliance checks.
How do you prevent agents from executing unintended actions?	Confirmation requirements, sandboxing, action simulation, dry-run modes.
What mechanisms detect and prevent agent manipulation?	Prompt injection detection, adversarial input filtering, behavioral analysis.
How is agent behavior audited for compliance?	Audit logs, compliance reporting, regulatory alignment verification.
What are your agent incident response procedures?	Containment, investigation, remediation, post-mortem analysis.

Agent Performance & Evaluation

Question	Response
How do you evaluate autonomous agent performance?	Goal completion rates, efficiency metrics, quality assessments.
What benchmarks are used for agent evaluation?	Internal benchmarks, public datasets, domain-specific tasks.
How do you measure agent success across different autonomy levels?	Level-specific metrics (e.g., routing accuracy, tool selection precision).
What is the agent's Time-to-Value (TTV)?	How quickly agents deliver measurable business value.
How do you assess cross-domain transfer performance?	Generalization to unfamiliar tasks and environments.
What are common failure modes for your agents?	Hallucination, tool misuse, infinite loops, escalation failures.
How do you measure agent cost-efficiency?	Cost per task completion, resource utilization, ROI metrics.

Retrieval-Augmented Generation (RAG)

RAG Architecture & Implementation

Question	Response
Are RAG models utilized for AI-driven content generation or knowledge retrieval?	Specific use cases and implementations.
What sources of knowledge are leveraged in RAG pipelines?	internal databases, web-based retrieval, proprietary corpora, document stores
What architectures are used for RAG implementations?	e.g., LangChain, LlamaIndex, Haystack, custom frameworks
Do you use vector databases to store embeddings?	Which database (Pinecone, Weaviate, Chroma, Qdrant, Milvus, FAISS, pgvector)?
Which embedding models do you use?	OpenAI embeddings, sentence-transformers, domain-specific models.
How is document chunking and preprocessing handled?	Chunking strategies, overlap, metadata extraction, preprocessing steps.
What retrieval strategies are employed?	Dense retrieval, sparse retrieval, hybrid search, reranking.

RAG Quality & Reliability

Question	Response
How do you ensure accuracy, consistency, and trustworthiness of AI-generated responses?	Quality assurance processes and validation methods.
What safeguards exist to mitigate hallucination and misinformation?	Fact-checking, source citation, confidence scoring, grounding verification.
How is retrieval relevance measured and optimized?	Relevance metrics, retrieval evaluation, continuous improvement.
What mechanisms ensure retrieved information is current and accurate?	Data freshness policies, version control, update frequencies.
How do you handle conflicting information from multiple sources?	Conflict resolution, source prioritization, uncertainty communication.
Are RAG responses traceable to source documents?	Citation mechanisms, source attribution, provenance tracking.

AI-Orchestrated Workflows & Automation

Workflow Automation Capabilities

Question	Response
Do AI models assist in real-time decision-making or forecasting?	Use cases and decision-making frameworks.
Do you use AI to automate complex workflows?	How does AI automate complex workflows or decision-making processes?
Are AI agents used for dynamic orchestration?	e.g., multi-agent collaboration, adaptive workflow routing
How are AI-driven automation capabilities monitored and optimized?	Performance tracking, bottleneck identification, continuous improvement.
How do AI automation components integrate with other enterprise systems?	Integration patterns, APIs, middleware, event-driven architecture.
What workflow orchestration tools are used?	e.g., Apache Airflow, Prefect, Temporal, custom orchestration

Additional Considerations

Ethical AI & Responsible AI Practices

Question	Response
What ethical AI principles guide your development?	Fairness, transparency, accountability, privacy, safety.
How do you assess and mitigate AI bias?	Bias detection, fairness metrics, mitigation strategies, diverse training data.
What transparency mechanisms exist for AI decisions?	Explainability, interpretability, decision documentation.
How do you handle AI-related incidents?	Incident response, remediation, stakeholder communication.
What responsible AI governance is in place?	Ethics boards, review processes, impact assessments.
What workflow orchestration tools are used?	e.g., Apache Airflow, Prefect, Temporal, custom orchestration

Regulatory Compliance & Standards

Question	Response
How do you comply with AI regulations?	EU AI Act, GDPR, industry-specific regulations.
What standards and certifications do you follow?	ISO/IEC standards, NIST AI frameworks, industry best practices.
How do you ensure AI system accountability?	Audit trails, documentation, responsibility assignment.
What risk classification applies to your AI systems?	High-risk, limited-risk, minimal-risk categorization.

Vendor & Third-Party AI Management

Question	Response
What third-party AI services or models do you use?	OpenAI, Anthropic, Google, open-source models, others.
How do you conduct vendor due diligence for AI partners?	Data quality, bias testing, transparency, security assessments.
What are the licensing terms for third-party AI components?	Commercial licenses, usage restrictions, attribution requirements.
How do you manage dependencies on external AI providers?	Vendor lock-in mitigation, contingency plans, multi-provider strategies.
What happens if a third-party AI service becomes unavailable?	Fallback mechanisms, redundancy, business continuity.

End of Questionnaire


BUNDESVERBAND
MERGERS &
ACQUISITIONS GEM. E.V.