# Beijing Haizhi Technology Group Co., Ltd.
# 北京海致科技集團股份有限公司

*(A joint stock company incorporated in the People's Republic of China with limited liability)*

**(Stock Code: 2706)**

## ANNOUNCEMENT OF INTERIM RESULTS FOR THE SIX MONTHS ENDED JUNE 30, 2026

The Board of our Company is pleased to announce the unaudited consolidated interim results of our Group for the Reporting Period, together with the comparative figures for the six months ended June 30, 2025.

### FINANCIAL HIGHLIGHTS

| | For the six months ended June 30, | | |
|---|---|---|---|
| | **(Unaudited)** | (Unaudited) | |
| | **2026** | 2025 | Change |
| | ***RMB'000*** | *RMB'000* | |
| Revenue | **295,563** | 173,469 | 70.4% |
| Gross profit | **132,084** | 66,782 | 97.8% |
| Loss for the period | **(78,966)** | (127,653) | (38.1)% |
| **Non-IFRS measure:**[Note] | | | |
| Adjusted net loss | **(19,601)** | (28,596) | (31.5)% |

*Note: Please refer to section headed "Non-IFRS Measure" in this announcement for more details.*

## BUSINESS REVIEW AND OUTLOOK

### I. Overall Performance

We are pleased to announce the Company's financial results for the first half of 2026. The Company generated total revenue of RMB295.6 million during the Reporting Period, representing a year-on-year increase of 70.4% (compared to RMB173.5 million for the corresponding period of 2025, accounting for 27.9% of the full-year revenue). The Atlas AI Agent business, delivered particularly outstanding performance, achieving exceptional growth to RMB114.6 million, a year-on-year increase of 135.6%, with its revenue contribution reaching 38.8%, firmly establishing its position as the core driver of the Company's overall growth.

#### *(1) Steady Improvement in Operating Quality*

In the first half of 2026, the Company recorded a net loss of RMB79.0 million, narrowed by 38.1% year-on-year, with adjusted net loss narrowed by 31.5% year-on-year to RMB19.6 million. Meanwhile, benefiting from the continued increase in the revenue contribution of the Company's Atlas AI Agent business, as well as the accumulation of the Company's solutions across diverse industries and application scenarios, the Company's overall gross margin expanded by 6.2 percentage points year-on-year to 44.7% during the Reporting Period, among which the gross margin of Atlas AI Agent grew by 4.8 percentage points year-on-year to 52.9%.

#### *(2) Significant Optimization of Capital Structure*

As at June 30, 2026, the Company's total assets increased substantially to RMB1,626.6 million compared to the end of 2025, with equity attributable to equity shareholders of the Company reaching RMB1,199.1 million, representing an increase of RMB2,650.5 million from the end of 2025. The aggregate of cash and cash equivalents and financial assets measured at fair value through profit and loss exceeded RMB971.7 million, reflecting ample funding reserves; the overall debt-to-asset ratio is 26.3%. Overall, following the Listing, the Company's capital structure has achieved significant optimization, and the overall financial structure remains robust, providing a solid foundation for sustainable long-term development.

### *(3) Continuous Deepening of Strategic Cooperation and Technological Breakthroughs*

During the Reporting Period, the Company established comprehensive strategic partnerships with leading enterprises across the AI industry chain, including core computing power providers, such as Hygon Information; leading large language model (LLM) developers, such as Z.AI; and top-tier enterprises in scenario application, such as Bank of Ningbo, establishing an industrial-grade AI ecosystem spanning "computing power – foundation models – scenarios". On the technological front, the Company made significant progress in both its Harness and AI database infrastructure: the Atlas Harness is further evolving from an Agent orchestration framework for single tasks into an industrial-grade Agent Runtime and standardized operating platform, transforming industrial Agents into production-grade systems. Concurrently, the Company has been progressively advancing the platformization and standardization of Atlas Harness, evolving its project-specific implementations into a reusable, manageable, and continuously evolving industry AI infrastructure. To build the AI database foundation for the LLM era, the Company upgraded its AtlasGraph Graph Database into a multi-model database that natively supports unified analysis across various data types. By organizing heterogeneous data around unified business objects, the system reduces the need for object matching across disparate data spaces, result concatenation and repeated data migration – allowing long-chain reasoning to run end-to-end within a unified computing framework. With this upgrade, the graph database is transitioning from an "analytical infrastructure" to a "runtime memory and reasoning for agents", and is becoming increasingly integrated with the Ontology and Harness layers to form an end-to-end technology stack, further strengthening the Company's full-stack foundation for industrial-grade AI.

In addition, the Company was officially included in the Hong Kong Stock Connect and the Hang Seng Composite Index during the Reporting Period, further enhancing the accessibility and market liquidity of the Company's shares, marking a significant milestone in expanding institutional investor coverage and elevating market recognition.

## II. TECHNICAL AND PRODUCT PROGRESS

### *(1) Industry Trend: Ontology-Centric Harness Emerging as a Key Paradigm for Industrial AI Deployment*

In the first half of 2026, large language models and AI agent technologies further advanced from capability exploration toward industrial deployment. As domestic large language models and open-source models continue to improve their capabilities, enterprises gained access to a broader range of foundation models, which are increasingly becoming interchangeable underlying capabilities that can be selected flexibly based on performance, cost, security and deployment requirements. At the same time, the core challenge in Enterprise AI has also shifted: customers' focus is extending beyond model performance itself toward deeper business understanding by models, enabling seamless integration into business processes and reliable execution of assigned tasks.

This shift is driving AI applications beyond their early stages – once centered on Prompt Engineering, RAG, and knowledge Q&A – toward Context Engineering and Harness Engineering, and further into Loop Engineering and Graph Engineering, among which technologies such as Context Engineering, Loop Engineering, and Graph Engineering are progressively outlining a clear roadmap for enterprise AI Harness. An industrial-grade AI agent is no longer merely a large model capable of autonomously calling tools, but rather a comprehensive intelligent system that must continuously acquire business context, manage task states, invoke enterprise capabilities, execute multi-step workflows, and verify, trace, and govern execution results. This entire set of engineering capabilities supporting the runtime operation of such industrial-grade Agents is referred to as Harness.

Within this architecture, Ontology is re-emerging as a critical component of enterprise AI infrastructure. Currently, the global leading Enterprise AI solution provider have established an expressive paradigm that uses ontology to connect business semantics with the digital production environment. The interaction between AI and humans has leaped from mutual understanding of language to AI forming complex cognitive capabilities across data, computation, and operations in the digital world, opening a new window for AI to fully take over enterprise digital production environments in the future. A number of leading technology companies and data intelligence platforms have also recently strengthened their ontology and enterprise semantic layer capabilities, providing agents with unified business objects, relationships, rules, and context. Ontology is evolving from its traditional role as a "knowledge organization method" into a business semantic layer connecting the enterprise data world with AI agents.

At a time when model development has not yet converged while application demands have already stepped into deeper waters, the assets that enterprises truly need to build on the AI front are also undergoing a transformation. Foundation models can be upgraded and replaced as technology advances, but the knowledge network formed by enterprises' own data, computation, operations, terminology, processes, and more, all built upon business ontology, has become the core of enterprise AI infrastructure. Through agent service capabilities such as MCP, CLI, Skill and Loop, this knowledge network has unleashed irreplaceable value, and also possesses sustainable reuse value – these are precisely the AI assets that enterprises can accumulate in the AI era. These assets need to be decoupled from models, and therefore, accumulating assets at the Harness layer with continuous construction has gradually become the mainstream paradigm for industrial AI deployment. This represents the enterprise-level practice of sovereign AI: what needs to be autonomously controlled is not only models, computing power, and data, but also the enterprise's own knowledge systems, business semantics, and AI execution and governance systems.

Currently and in the future, the value of Haizhi in the enterprise AI landscape lies in helping enterprises comprehensively build business intelligence infrastructure beyond the model layer: a complete industrial-grade AI stack with a multi-model database as the data and computing foundation, ontology as the business semantic layer, and Harness as the intelligent execution engine. On top of this infrastructure, the architecture can continuously adapt to evolving large models, transforming model capabilities into enterprise-grade business productivity.

***(2) From Harness to Ontology: Enterprise Deployment of Industrial-grade Intelligent Agents***

*(i) Harness: From Model Capability to System Capability*

As large language models evolve from content generation to task execution, the bottleneck of AI systems is gradually shifting from "whether the model is smart enough" to "whether the model can consistently and reliably complete tasks". Prompt Engineering addresses how tasks should be articulated to the model, while Context Engineering addresses what information the model should receive. However, when agents begin to handle multi-step, long-horizon tasks, the model alone cannot reliably manage state, tool execution, failure recovery, permission control, and result verification. This has given rise to Harness Engineering.

At its core, the Harness builds system capabilities around the model–including **Context, Tools, State, Memory, Orchestration, Verification and Guardrails**–to transform probabilistic model intelligence into execution capabilities that are controllable, sustainable and verifiable. From this point forward, the focus of AI engineering has shifted from "optimizing model inputs" to "designing the model's operating environment".

*(ii) The Expansion of the Harness Concept: From Agent Reliability to Enterprise Operating Environment*

In terms of current industry terminology usage, the concept of Harness has developed a distinction between its narrow and broad senses. Harness in its narrow sense, i.e. Harness Engineering, primarily addresses the execution reliability of a single agent, such as context management, tool invocation, execution loops, retries, and verification. As Agent task's complexity increases, or as Agent begins to tackle problems in vertical industry scenarios, the scope of Harness gradually expands. It broadens to encompass workflows multi-agent collaboration, permissions, security, human-in-the-loop, and observability. At this point, Harness is no longer merely a narrow form of agent reliability engineering, but has begun to serve as the operating and control layer between the model and the external environment, which defines Harness in its broad sense. The core distinction lies in their respective objectives: the narrow sense enables a foundational model to become a general-purpose executable agent, whereas the broad sense ensures that AI operates reliably and securely within a specific, typical real-world enterprise environment.

When the problems agents address shift from open-world exploratory tasks to constrained-world problem-solving, and when the agent's environment shifts from general-purpose scenarios composed of files, web pages, and general-purpose tools to real-world enterprise business scenarios, the nature of the problems fundamentally changes, and consequently, the architecture and objectives of the Harness system evolve accordingly.

*(iii) Industrial-grade Agents: First and Foremost, a "Computable Enterprise"*

General-purpose agents can typically rely on linguistic knowledge and open-ended reasoning to perform tasks such as search, research, and programming. Industrial-grade agents, however, must form quantitative judgments based on the actual operating state of an enterprise. Decision-making in areas such as supply chain, manufacturing, energy, finance is not simply about understanding "what should theoretically be done", it requires knowledge of current inventory, production capacity, orders, costs, available resources, risks, and equipment status, and then feeding these variables into business formulas, predictive models, optimization algorithms, and constraint conditions for computation.

Therefore, the first prerequisite for deploying industrial-grade agents is not simply connecting to a few business systems, but rather establishing a unified enterprise data and computing foundation. This is precisely the role of the Company's DMC Data Intelligent Platform. Its value lies in uniformly ingesting, governing, transforming, and computing structured and unstructured data scattered across ERP, MES, CRM, databases, sensors systems and documents, while consolidating the metrics, models, and business computation relationships as reusable assets. While general-purpose models can comprehend decision logic, they can only generate executable, quantitative decisions when that logic is applied to real-world enterprise data.

*(iv) Ontology: From a "Computable Enterprise" to a "Understandable and Actionable Enterprise"*

While data and computation can answer "what is currently happening in the enterprise" and "what the underlying variables are," they cannot bridge the semantic gap. Industrial agents still need to comprehend the underlying business semantics of this data. The role of Haizhi Ontology is to further abstract the enterprise, on top of the underlying data and computation, into unified business objects, relationships, logic, and actions: suppliers, orders, factories, and equipment become Objects; dependencies between objects become Relations; business rules, models, and decision boundaries become Logic; and real-world operations such as procurement, scheduling, and order modifications become Actions.

In short, Haizhi Ontology provides AI with an understandable and actionable enterprise world model. The fundamental difference between industrial agents and general-purpose agents does not primarily stem from model capability, but rather from the environment in which they operate. Real-world enterprise environments contain well-defined business objects, processes, real-time status, role-based permissions, decision prerequisites, compliance requirements, and specific information systems. An agent must not only know "what can be done", but also understand "whether it should be done under the current enterprise state, who needs to approve it, and what impact it will have after execution". Haizhi Ontology contains all the enterprise semantics and operational modeling that agents require, providing rules, permissions, and action spaces, and this is precisely the core distinction between industrial AI Harness and general-purpose agent Harness. Agents operate on top of this business abstraction layer, invoking underlying data, models, and computational capabilities to formulate quantitative judgments. Through this architecture, natural language reasoning is truly transformed into actionable industrial decisions.

*(v) From Ontology to FDE: Continuously Encoding Industrial Knowledge into Ontology*

However, an enterprise's true operational logic cannot be automatically derived from databases alone. A substantial amount of critical knowledge resides in the experience of business practitioners, cross-departmental collaboration, historical processes, and implicit decision-making conventions – such as which variables truly impact decisions, under what circumstances exceptions can be made, and what risks warrant vigilance even when data appears normal. This tacit knowledge determines whether Ontology remains merely a data model or truly reflects how the enterprise actually operates.

The role of Forward Deployed Engineers ("**FDE**") is to go deep into real business sites, work collaboratively with customers to identify these critical objects, relationships, logic, workflows, actions, and decision boundaries, and continuously consolidate into Haizhi Ontology and platform capabilities. The fundamental difference between FDE and traditional B2B customization lies in their objectives: traditional customization aims to "build dedicated features for a specific customer", whereas FDE aims to "abstract reusable product capabilities and continuously enhancing Ontology in reverse by solving the most challenging problems of real customers". Thus, the data platform consolidates enterprise facts, Ontology consolidates the enterprise operating model, and FDE continuously encodes industry experience, which cannot be directly derived from data, into this model.

At this stage, the architectural framework for industrial-grade agents is fully established. Its core competency is no longer defined by a single model, but by a comprehensive, systemic capability: Models provide general intelligence; Atlas Harness orchestrates this intelligence into reliable execution; the data and computing foundation renders enterprise reality computable; Haizhi Ontology translates the way an enterprise operates into a dynamic world that AI can comprehend and act upon; and FDEs continuously update this world from the frontlines of industry practice.

### *(3) Expansion of Ontology Modeling Scope and Evolution into Standardized Platform Capabilities*

In the first half of 2026, the Company further upgraded its Haizhi Ontology capabilities, driving the evolution from a traditional "data semantic model" centered on entities and relationships toward an "enterprise operating model" encompassing business objects, states, rules, and actions. Through unified abstractions such as Objects, Properties, Relationships, and Actions, the Company incorporates objects – such as customers, contracts, accounts, equipment, orders, and events – along with their relationships, states, and executable actions into a single semantic system, and further associates Actions with Tools, Skills, and Workflows. As a result, Ontology not only describes "what the enterprise has", but also expresses "how the enterprise operates and what actions can be taken".

Building on this foundation, the Company is further productizing the capabilities for Ontology construction, configuration, management, and reuse, gradually developing standardized and automated toolchains to streamline Haizhi Ontology development and operations. These capabilities empower Agents to transcend the mere ability to execute tasks, evolving into production-grade systems that are fully observable, diagnosable, evaluable, and continuously optimizable. As a result, Haizhi Ontology has evolved from a semantic layer for data understanding into a unified business capability layer that connects enterprise data, business logic, and agent action capabilities.

It should be particularly noted that, given the long-standing preference of large Chinese enterprises for private deployment, graph-based ontology management technology demonstrates stronger robustness in the management of complex enterprise knowledge and the construction and operation of unified semantics. This represents the most significant difference between ontology construction practices in Chinese enterprises and standardized commercial software and SaaS-based environments of US enterprises. Looking ahead, three major trends are expected to become increasingly prominent: the development and operation of **enterprise-wide, end-to-end and scenario-spanning unified semantic Ontologies**; the automated construction and parsing of Ontologies; and the evolution of Ontology-based systems from **knowledge networks into complex decision networks**. In the future, an ontology-centric enterprise knowledge platform has the potential to further evolve beyond traditional data platforms, while the focus of enterprise-wide IT architecture may gradually shift from serving human employees toward serving enterprise AI systems.

### *(4) Harness Evolving Toward Industrial-grade Agent Runtime and Standardized Operating Platform*

In the first half of 2026, the Company's Atlas Harness further evolved from a single-task agent orchestration framework toward an industrial-grade Agent Runtime and standardized Agent operation platform. The Company has continuously enhanced its capabilities across Context construction, long-horizon task state management, Tool/Skill/Workflow orchestration, Trace, Replay, Metrics, Evaluation, and execution isolation. These capabilities enable Agents to advance from merely "**being able to execute tasks**" into production-grade systems that are **observable, diagnosable, evaluable and continuously optimizable**.

Building on this foundation, the Company's Atlas Harness is further advancing the platformization and standardization of these capabilities, gradually forming a unified engineering foundation covering task configuration, tool management, runtime observability, effectiveness evaluation, and continuous optimization. Going forward, depending on enterprise technology environments and business complexity, configuration, operation, maintenance and iteration can be jointly carried out by customer technical teams and the Company's FDE teams. Through this approach, Harness will continue to evolve from project-level engineering capabilities into **reusable, manageable and continuously evolving infrastructure for industrial-grade Agents**.

### *(5) From Graph Database to Multi-Model Database: Building a Unified Data Computing Foundation for Industry Agents*

In the first half of 2026, building on the continued iteration of its core AtlasGraph Graph Database engine, the Company further advanced the development of a next-generation multi-model database integrating graph, time-series, and vector data, driving the evolution of its underlying data capabilities from a focus on complex relationship computation toward a unified data computing foundation encompassing relationships, states, and semantics. Graph data describes the associations and dependencies between business objects; time-series data records changes in object states and events over time; and vector data carries semantic representations of unstructured knowledge such as documents and text. As AI applications increasingly demand joint analysis across multiple data types, the underlying data infrastructure is gradually trending from single-model toward multi-model convergence. The core of this shift is not about simply adding more database types, but rather organizing heterogeneous data around unified business objects, enabling the relationships, states, and knowledge associated with the same objects to be connected and understood within a unified computing framework.

Compared with traditional graph databases, the key advancement of this upgrade is the transition from single-model relationship computation toward cross-model integrated computing. Taking equipment fault analysis as an example: if a system uses separate graph, time-series, and vector databases, it must first identify equipment relationships, operational metrics, and maintenance knowledge in three independent data spaces, then rely on the application layer to match objects, splice results, and perform multiple rounds of data movement. In a unified multi-model system, however, equipment becomes the business object commonly referenced by all three data types, allowing the system to continuously complete fault knowledge retrieval, upstream/downstream impact analysis, and abnormal metric tracing around the same equipment. Thus, analysis is no longer a simple aggregation of results from three databases, but rather a joint computation across "relationships – states – knowledge" based on a unified object context, reducing cross-system integration complexity and providing a consistent data foundation for longer-chain machine reasoning.

More importantly, the multi-model database is further being integrated with Haizhi Ontology and Atlas Harness to form an end-to-end technology architecture. The underlying database is responsible for storing and computing the relationships, states, and knowledge that constitute the enterprise's real-world operations. Haizhi Ontology then maps these underlying data assets into unified business Objects, Relationships, States, and callable capabilities, allowing graph queries, time-series analysis, and vector retrieval to be encapsulated as business-semantic Tools or Actions rather than directly exposed to Agents. Atlas Harness subsequently selects and orchestrates these capabilities based on task objectives, context, and business constraints. Under this architecture, the multi-model database addresses "how enterprise reality can be computed"; Haizhi Ontology addresses "what these computations mean in business terms"; and Atlas Harness further addresses "how agents can reliably use these capabilities in specific tasks"; Together, they form the complete technology stack enabling industrial Agents to progress from underlying data computation to business understanding, and ultimately to decision execution.

***(6) Industry-Academia-Research and Technology Ecosystem Collaboration: Upgrade of Academician Workstation and LLM Ecosystem Synergy***

The Company continued to consolidate its technology foundation through closer collaboration across industry, academia and research and a more prominent role in shaping industry standards. During the Reporting Period, the Academician Expert Workstation completed a strategic upgrade, shifting its R&D focus from high-performance graph computing to the integration of "knowledge graphs and LLMs". Led by the Company's Chief Scientist, Academician Zheng Weimin, the workstation concentrates on large-scale distributed systems and the frontier of graph-LLM integration, marking a critical step for the Company in building a trustworthy application foundation for LLMs and providing stronger research support for the continued innovation of Graph-LLM integration technologies. In addition, the Company's wholly-owned subsidiary, Beijing Haizhi Xingtu Technology Co., Ltd., as a core drafting

unit, participated in the development of national standards including Information Technology – Neural Network Representation and Model Compression – Part 3: Graph Neural Networks and Intelligent Computing – Graph Computing Performance Testing Methods, which were officially released during the reporting period. The Company is progressively consolidating its years of technical practice experience into reusable and scalable industry standards, building a unified "technical language" for graph intelligence technologies in the era of LLMs and facilitating efficient collaboration across technology R&D, product interoperability, and engineering deployment within the industry.

In terms of technology ecosystem collaboration, the Company's Atlas AI Agent is currently compatible with over 100 LLMs. During the Reporting Period, the Company entered into a strategic cooperation framework agreement with Z.AI Co., Ltd. ("**Z. AI**"). Under the agreement, the Company and Z.AI will pursue strategic cooperation in areas such as model training and application scenario deployment, leveraging their respective technological advantages, industry expertise, and market resources to achieve complementary resources and synergistic development. In addition, in June 2026, the Company formally signed a strategic cooperation agreement with Hygon Information Technology Co., Ltd.. The two parties will focus on industries such as public services and security, jointly developing a full-stack solution of "domestic computing power + graph-LLM intelligence" to drive the large-scale deployment of domestically produced AI technologies in key industries. Going forward, the Company will continue to deepen its efforts in "Graph-LLM integration" and AI agent-related technologies, increase R&D investment in underlying technologies, and continuously iterate its full-stack self-developed product system, accelerating the large-scale deployment of trustworthy industry intelligence services in key sectors such as finance, energy, and public services. The Company will continue to advance the development of industrial AI infrastructure, creating sustainable long-term value for investors and partners.

## III. Business Progress

Leveraging its full-stack technological capabilities ranging from AI infrastructure to the development of industrial-grade AI agents, as well as its end-to-end industry-specific solutions and delivery expertises, the Company continues to strengthen its market position in key sectors such as public services, energy, and finance:

- **Public Services:** Continuously strengthening the leadership position in the security-related sector. Flagship customers achieved consecutive repurchases within the six-month period, covering core products including graph databases, graph ontology, and public services AI agents.

- **Finance:** While maintaining solid repurchase rates from existing customers, the Company is actively acquiring new customers. The Company officially signed a strategic cooperation agreement with Bank of Ningbo to deepen the integration of AI technologies with financial services. Meanwhile, breakthrough progress was made in new customer acquisition in the Hong Kong market.

- **Energy:** The energy sector maintained a strong growth momentum. As the development path of AI Autonomy is becoming increasingly clear, the Company has achieved deep coverage in the core application scenarios of power grid customers.

Furthermore, the Company has been ramping up its efforts in emerging sectors such as smart mining, telecommunications operators, and smart manufacturing, achieving milestone breakthroughs in emerging industry expansion:

- **Smart Mining:** The Company signed a multi-million RMB project contract with a large coal mining group to deploy its Atlas AI Agent, building an intelligent decision-making support system that covers production, safety, and operational management.

- **Telecommunications Operators:** Leveraging the proven delivery expertise and benchmark practices, the Company has continuously acquired new leading customers in the sector. This validates the Company's capabilities in rapid adaptation and horizontal replication across similar customer scenarios.

- **Oil and Gas:** The Company has successively won bids for multiple projects in core business areas such as production extraction and settlement operations for multiple large state-owned enterprise customers, ushering in a new phase of market expansion in this sector.

- **Other Emerging Industries:** The Company is actively exploring opportunities and has secured orders in areas such as China's aircraft manufacturing, smart vehicles, industrial robots, pharmaceutical production and distribution, and urban heating operations.

In terms of business outside Mainland China, the expansion into the Hong Kong market achieved phased progress. During the Reporting Period, the Company successfully established a local legal entity and dedicated team in Hong Kong, leading to a continuous expansion of the coverage in the financial service industry. Meanwhile, the Company is also focusing on the public services sector, further deepening local customers' understanding of its solutions through continuous engagement and scenario demonstrations with multiple public service departments and public organizations, laying a solid foundation for subsequent business deployment.

The Company has cumulatively served over 430 customers to date. In the first half of 2026, the average revenue per customer was RMB3.4 million, with the average revenue per Atlas AI Agent customer reaching RMB4.0 million. Among the Atlas AI Agent customers that contributed revenue, 69% had previously deployed the Company's Atlas Graph Solution products, demonstrating the strong customer relationships the Company has built.

### *(1) Atlas Graph Solution*

In the first half of 2026, the Company's Atlas Graph Solution generated revenue of RMB181.0 million, representing a year-on-year increase of 45.0%, serving 59 customers with an average revenue per customer of RMB3.1 million. The Company has continued to consolidate its customer coverage and engagement depth in areas such as public services, energy, and banking, while establishing business collaborations with typical customers across multiple sub-sectors.

- **A Major Public Services Customer:** This is a high-level public services customer with massive data volumes and stringent compliance requirements. The Company secured this project in 2025 and completed the main implementation in the first half of 2026. The Company delivered a comprehensive solution for the customer's systematic development of knowledge base, rule tools, and knowledge graph tools. Customer's data is characterised by diverse sources, inconsistent semantic standards and dense business rules, with demanding requirements for cross-level and cross-department collaboration. Through knowledge engineering, the Company established an integrated "knowledge base – rules – graph" knowledge service system spanning the full chain from data semantic alignment and knowledge processing to rule-based reasoning and intelligent applications, empowering scenarios such as intelligent Q&A, assisted review and correlation analysis.

- **A Leading Financial Leasing Enterprise:** As an industry leader with total assets exceeding RMB100 billion, this customer has selected the Company to build its AI data infrastructure. The Company provides an integrated solution built on its proprietary AtlasGraph Graph Database, Atlas Knowledge Graph Platform and GraphRAG. Financial leasing is characterized by lengthy approval processes and numerous regulatory constraints. With the implementation of the financing lease related clauses of Measures for the Administration of Financial Leasing Companies in 2026, the industry has entered an era of full-process supervision centred on leased assets, positioning AI as the new financial infrastructure. This marks the Company's first systematic deployment of a three-layer system of comprising metadata knowledge graphs, event logic knowledge graphs, and text knowledge graphs in the financial leasing industry. This system covers the entire chain from data comprehension to business reasoning to knowledge Q&A. By deeply linking the core elements such as counterparty qualifications, leased assets, guarantee measures and approval processes through the knowledge graph, this deterministic knowledge framework effectively minimizes hallucination risk, providing a trusted data foundation for AI applications in strictly regulated scenarios.

- **A Leading Domestic Power Grid Enterprise:** This customer, being a subsidiary of a leading domestic power grid enterprise, bearing the critical responsibility of building the grid's data center and core digital infrastructure. As the construction of new power systems and digital grids accelerates, the demand for data-driven intelligent grid operations continues to grow, rapidly driving the application of graph database in scenarios such as topology analysis, fault tracing, and dispatch analysis. By leveraging the proprietary integrated solution based on AtlasGraph Graph Database, the Company has built a foundational graph data platform for the customer. This platform enabled unified management of grid object relationships, topological connections, and real-time operational states. It supports high-concurrency, low-latency query and analysis of large-scale graph data, empowering critical scenarios across dispatching, power transmission, transformation and power distribution. Specifically, it facilitates fault localization, impact analysis, equipment status analysis, and real-time cross-section analysis, ultimately driving the objectification of grid relationship data, integration of time-series states, efficient topology analysis, and intelligent business applications.

### *(2) Atlas AI Agent*

In the first half of 2026, the Company's Atlas AI Agent generated revenue of RMB114.6 million, representing a year-on-year increase of 135.6%, serving 29 customers with an average revenue per customer of RMB4.0 million. Since the launch of its AI Agent products, the Company has achieved strong momentum in repeat purchases and new customer acquisition. In the first half of 2026, the customer conversion rate from Atlas Graph Solution to Atlas AI Agent reached 69%. Beyond continuously converting AI Agent products within established core industries, the Company has also been extensively exploring new customers and securing orders in numerous new industries and scenarios.

- **A Large Coal Mining Group:** The Company has signed a multi-million RMB project contract with a large coal mining group to deploy the Haizhi Atlas AI Agent at one of its national model mines. The main phase of the project has been completed. Driven by the national "AI+" initiative and policies such as the Guiding Opinions on Deepening Intelligent Mine Construction to Promote Mine Safety Development, the digital and intelligent transformation of the coal industry continues to accelerate. This transformation is a key driver of high-quality development in this industry. The Company leveraged the Atlas AI Agent to advance core operations including intelligent energy consumption analysis and dynamic optimization, predictive equipment maintenance, and intelligent coordination across the entire production process. At the same time, the Company is deploying multiple models for key businesses, covering intelligent safety risk warning, collaborative production scheduling, and assisted management decision-making, forming a end-to-end intelligent support capability from production to operation. It builds an intelligent decision support system for production, safety, and operational management for this model mine, driving the shift of mine operations from relying on manual experience to data-driven intelligent decision-making.

- **A Major Public Services Customer:** In the first half of 2026, the Company completed a multi-million RMB project for this major public services customer, delivering an integrated system based on its proprietary Atlas AI Agent. This intelligent support system, built upon a comprehensive Harness runtime architecture, is designed to meet the modern governance needs of mega-cities, seamlessly integrating full-domain perception, intelligent assessment, and closed-loop execution. Powered by the DMC Data Intelligence Platform and supported by a full-scale data warehouse and business scenario asset library, this system enables unified aggregation of multi-source data, full-lifecycle tracking and control, and the consolidation of knowledge assets. This establishes a standardized data governance closed loop and a solid data foundation for end-to-end business operations. Utilizing Ontology as the carrier, the system addresses three core scenarios: routine management, emergency response, and special operations. It enables panoramic situational awareness, traceable risk visualization, and closed-loop command execution that is closely aligned with frontline command and dispatch processes. With a large language model serving as the core intelligent engine, AI capabilities are deeply embedded throughout the entire business chain. This enables semantic data interpretation, automated risk assessment and classification, intelligent response plan generation, and precise command dispatch — upgrading traditional manual evaluation into intelligent, assisted decision-making. The deep integration of these three components propels the governance model from passive response to proactive anticipation, significantly enhancing the efficiency of risk prevention, decision support, and collaborative execution capabilities.

- **A Hong Kong-based Integrated Financial Services Institution:** The Company has signed a project contract with this customer to provide an intelligent agent solution for prospectus document review. The Company has built an intelligent review system that integrates semantic understanding, large language model logical reasoning, and a next-generation semantic rule engine for the customer. The system systematically models document drafting standards and the articulation relationships between financial and business data, establishing a configurable, reusable, and extendable skill library that enables a "human-AI collaborative, standards-aligned" intelligent review paradigm. This achievement marks a significant breakthrough for the Company's industrial-grade agent capabilities in the highly specialized investment banking segment within the financial sector. It also lays a scalable foundation for future expansion into more compliance-intensive document review scenarios.

- **A Subsidiary of a Leading Grid Power Enterprise:** The Company has initiated the construction of an AI innovation platform for this customer. With the accelerated deployment of large language models and AI agents, the power industry is experiencing a rapid surge in demand for AI knowledge services and secured, domestically-developed IT infrastructure. The Company has systematically built an ontology for power industry knowledge, enabling precise natural language-based Q&A and dynamic digital twin capabilities. Meanwhile, the Company has consolidated its power industry knowledge components and modules into standardized platform capabilities, significantly reducing delivery cycles and supporting the rapid deployment of digital and intelligent applications across multiple scenarios. This validates the long-term effectiveness of integrating ontology with large language models in large enterprise-level environments.

- **A Leading Pharmaceutical Enterprise:** The Company has successfully established a business partnership with a leading pharmaceutical enterprise. Pharmaceutical companies have long value chains and fragmented data sources scattered across R&D, production, quality control, and sales channels, making it difficult for management to obtain a unified, real-time, and reliable full view of business operations. As information systems in the pharmaceutical industry are progressively enhanced, pharmaceutical companies are continuously improving their digital infrastructure and accelerating data platform development. The Company provided the customer with the Atlas AI Agent, enabling real-time, drill-down data querying, in-depth data analysis, and decision support through natural language interaction. This marks a significant breakthrough for the Company's industrial-grade intelligent agents in the pharmaceutical industry, laying the foundation for subsequent scaled expansion in this sector.

## IV. Future Outlook

As foundation models continue to scale up in parameters, the improvement in model intelligence capabilities has raised higher expectations. However, in serious and complex enterprise and public service application scenarios, one thing has become clear: the role of an ontology-centric Harness system is just as critical as the model's inherent capabilities in the deployment of industrial-grade AI. The Company will continue to deepen the practical implementation of industrial-grade intelligent agents across various industries and advance capability building in the following areas:

### *(1) Continuously Build a Cross-Industry, Cross-Domain Core Technology Foundation*

The world is the sum of complex relationships, and graphs are used to accurately describe such complexities. Therefore, the graph forms the underlying architecture for large-scale, complex ontology systems. In industries with profound societal impact or highly intricate knowledge domains, such as manufacturing, finance and energy, AI applications must build a human-AI integrated system: Ontology. This ontology acts as the pivotal anchor connecting multiple trillion-dollar markets, including

AI infrastructure, large language model capability development, and industrial AI applications. Along this direction, the Company will continue to invest in the R&D of core underlying technologies, iterating its full-stack proprietary product ecosystem across two key layers: the industrial Harness and the multi-model database, ensuring that its core technologies possess cross-industry and cross-domain applicability and reusability.

In terms of industry expansion, building upon the steady and deep engagement in its established industries, including finance, public services, and energy, the Company will accelerate its expansion into multiple new industries and applications such as smart manufacturing, mining, telecommunications, and pharmaceuticals. The benchmark projects implemented by the Company in these areas during the Reporting Period have preliminarily validated the cross-industry replicability of the Company's technical capabilities.

***(2) Strengthening the Talent Pool of AI-Native and Industry Experts***

In terms of talents, the Company will continue to attract and cultivate AI-native talent, building a multidisciplinary R&D team that is proficient in ontology modeling and agent engineering methodologies, providing support for the continued evolution of the technological framework. In parallel, the Company will intensify its efforts to recruit multi-industry talent, bringing in professionals with deep domain expertise to continuously translate industry knowledge into ontology models, skill libraries, and solutions. This will support the expansion into new industries to transit from "single-point breakthroughs" to "systematic coverage".

***(3) Consolidating Standardized Industrial Products through Benchmark Cases with Flagship Customers, and Gradually Transitioning to Scalable Growth***

On the industrial and application fronts, the Company continues to consolidate standardized products and modules with industry-specific characteristics by serving benchmark customers, solidifying delivery experience into reusable ontology models, skill libraries, and scenario components. This drives the revenue growth model to gradually shift from being driven by "headcount and workforce productivity" to "scalable replication of products and services". The Company is actively exploring opportunities for rapid volume growth in established industries and scenarios. This approach supports the simultaneous enhancement of both revenue scale and revenue quality through an optimized marginal cost structure.

***(4) Continuously Enhancing Management's Learning Capabilities and Organizational Operational Efficiency***

As the business scale and organizational complexity expand, the Company's management team will continue to strengthen its self-learning capabilities, maintaining keen awareness and rapid absorption of emerging technologies and developments while simultaneously building systematic management capabilities commensurate with a larger enterprise scale. By optimizing management processes and organizational mechanisms, the Company will continuously improve operational efficiency, providing a solid organizational foundation for the Company's scaled development.

***Reaching for the Stars while Staying Grounded***

As the rapid iterations in AI technology repeatedly inspire us to glimpse the AGI horizon, the Company firmly believes that the path of industrial-grade AI lies in production lines, in the depths of mines, in the silent operation of oil wells and pipelines in the wilderness, and in the efficient and invisible protection of public safety. While technology advances in leaps, the intelligent development of enterprises is gradual and winding. It requires us to respect and revere the complexity of industry knowledge and the unknown yet to be reached by humanity. It demands that we adhere to the "First Principles" of starting from reality and the frontlines, with problem-solving as our ultimate driving force.

Under this conviction, the Company will continue to explore the frontiers of technology while maintaining a strong focus on operational quality, driving the simultaneous improvement in revenue scale, delivery efficiency, and profitability, and striving to achieve scalable profitability while consolidating its technological leadership. Through long-term, steady operations, the Company is dedicated to creating sustainable, long-term value for its shareholders.

## FINANCIAL REVIEW

### Revenue

The revenue of the Group amounted to RMB295.6 million during the Reporting Period, which was increased by RMB122.1 million (or approximately 70.4%) from RMB173.5 million for the six months ended June 30, 2025. This was mainly due to an increase in revenue generated from Atlas Graph Solutions and Atlas AI Agent.

***Atlas Graph Solutions.*** The revenue from Atlas Graph Solutions increased by RMB56.2 million (or approximately 45.0%) from RMB124.8 million for the six months ended June 30, 2025 to RMB181.0 million for the Reporting Period, serving 59 customers, mainly as a result of an increase in market demand for our Atlas Graph Solutions, driven by (i) our further consolidation of the scope and depth of customer coverage in public service, energy, banking and other sectors; and (ii) our secured business cooperation with benchmark customers across multiple niche segments.

***Atlas AI Agent.*** The revenue from Atlas AI Agent significantly increased by RMB66.0 million (or approximately 135.6%) from RMB48.6 million for the six months ended June 30, 2025 to RMB114.6 million for the Reporting Period, serving 29 customers, mainly as a result of an increase in market demand for our Atlas AI Agent, driven by (i) our continuous iteration of our AI agent products within advantageous industry verticals; (ii) our extensive customer acquisition across a broad range of new industries and scenarios, with tangible order conversion achieved; and (iii) an increase in demand for Atlas AI Agent from our existing Atlas Graph Solution customers, with a 69% customer conversion rate, driven by the synergies between our solution offerings that have established an integrated system capability that transforms model capabilities into enterprise-level business productivity.

### Cost of sales

Cost of sales of the Group increased by RMB56.8 million (or approximately 53.2%) from RMB106.7 million for the six months ended June 30, 2025 to RMB163.5 million for the Reporting Period. This was mainly due to an increase in customer demand for our Atlas Graph Solutions and Atlas AI Agent, which led to a corresponding increase in related costs of sales, and was in line with our increase in revenue.

### Gross profit and gross profit margin

Gross profit of the Group was RMB132.1 million for the Reporting Period which was increased by RMB65.3 million (or approximately 97.8%) from RMB66.8 million for the six months ended June 30, 2025. The overall gross profit margin of the Group increased from 38.50% for the six months ended June 30, 2025 to 44.7% for the Reporting Period, which was mainly due to (i) the improvement in our operational and fulfillment efficiency, driven by (a) the continuous consolidation of the Company's technological and business capabilities across different industries, and (b) our strategic optimization of cost structure; and (ii) our Atlas AI Agent business, which has higher gross profit margin, grew from 48.1% for the six months ended June 30, 2025 to 52.9% for the Reporting Period.

**Selling and marketing expenses**

Selling and marketing expenses of the Group were RMB43.9 million for the Reporting Period which was increased by RMB5.8 million (or approximately 15.1%) from RMB38.1 million for the six months ended June 30, 2025. This was mainly due to the increase in total sales staff costs arising from the expansion of the sales team to support the Company's business development into new industries.

**Administrative expenses**

Administrative expenses of the Group increased by RMB33.5 million (or approximately 70.4%) from RMB47.5 million for the six months ended June 30, 2025 to RMB81.0 million for the Reporting Period. This was mainly attributable to: (i) the recognition of RMB17.5 million in listing expenses relating to the Global Offering commenced in March 2025 for the Reporting Period under the accrual basis of accounting; (ii) the incremental amortization of share-based compensation expenses during the Reporting Period arising from the grant of post-IPO share options on June 15, 2025; and (iii) higher staff costs associated with the enhancement of the Company's organizational management.

**R&D expenses**

R&D expenses of the Group increased by RMB29.5 million (or approximately 81.5%) from RMB36.2 million for the six months ended June 30, 2025 to RMB65.6 million for the Reporting Period. This was mainly due to an increase in staff costs and outsourced R&D costs, and our enhanced efforts on research and development of AI agent, considering the increasing market demand and recognition of AI agents for the Reporting Period.

**Impairment losses on financial assets and contract assets**

The impairment losses on financial assets and contract assets slightly increased from RMB4.6 million for the six months ended June 30, 2025 to RMB5.0 million for the Reporting Period.

**Other income and loss, net**

We recorded a net other loss of RMB2.3 million during the Reporting Period, compared to a net other income of RMB0.4 million for the six months ended June 30, 2025, mainly due to a significant foreign exchange loss recognized during the Reporting Period.

**Finance costs**

Finance costs of the Group remained relatively stable at RMB0.2 million for the six months ended June 30, 2025 and 2026.

**Changes in the carrying amount of redemption liabilities**

The changes in the carrying amount of redemption liabilities decreased by 71.7% from RMB64.7 million for the six months ended June 30, 2025 to RMB18.3 million for the Reporting Period, primarily due to the termination and derecognition of the redemption liabilities upon the Listing of the Company on February 13, 2026.

**Changes in the fair value of financial liabilities measured at fair value through profit or loss ("FVPL")**

The changes in the fair value of financial liabilities measured at FVPL for the six months ended June 30, 2025 amounted to RMB5.7 million and no such loss was recorded for the Reporting Period, as the anti-dilution rights granted to pre-IPO investors were fully exercised in May 2025, and all remaining preferred rights were terminated upon Listing of the Company on February 13, 2026.

**Income tax credit**

We recorded an income tax credit of RMB1.3 million for the Reporting Period and RMB2.2 million for the six months ended June 30, 2025, primarily arising from deferred tax adjustments.

**Loss for the period**

As a result of the foregoing, our loss for the Reporting Period decreased by 38.1% from RMB127.7 million for the six months ended June 30, 2025 to RMB79.0 million for the Reporting Period.

**Non-IFRS measure**

To supplement our consolidated financial statements, which are presented in accordance with IFRS Accounting Standards, our Group also uses adjusted net profit/loss (non-IFRS measure) as an additional financial measure, which is not required by or presented in accordance with IFRS Accounting Standards. We believe that this non-IFRS measure facilitates comparisons of operating performance from period to period and company to company. We believe that this non-IFRS measure provides useful information to investors and others in understanding and evaluating our consolidated results of operations in the same manner as they help our management. However, our presentation of adjusted net profit/loss (non-IFRS measure) may not be comparable to similarly titled measures presented by other companies. The use of such non-IFRS measure has limitations as an analytical tool, and you should not consider it in isolation from, or as a substitute for an analysis of, our results of operations or financial condition as reported under IFRS Accounting Standards.

The following table reconciles our net loss for the period presented in accordance with IFRS Accounting Standards to our adjusted net loss (non-IFRS measure):

| | For the six months ended June 30, | |
|---|---|---|
| | **2026** | 2025 |
| | ***RMB'000*** | *RMB'000* |
| **Reconciliation of net loss to adjusted net loss *(non-IFRS measure)*** | | |
| Loss for the period | **(78,966)** | (127,653) |
| Add: | | |
| Equity settled share-based payments[(1)] | **23,560** | 15,538 |
| Changes in the fair value of financial liabilities measured at FVPL[(2)] | **–** | 5,744 |
| Changes in carrying amount of redemption liabilities[(3)] | **18,318** | 64,745 |
| Listing expenses[(4)] | **17,487** | 13,030 |
| **Adjusted net loss for the period *(non-IFRS measure)*** | **(19,601)** | (28,596) |

*Notes:*

(1) We adjust equity-settled share-based payment expenses because they were non-cash in nature.

(2) We adjust changes in the fair value of financial liabilities measured at FVPL because it was non-cash in nature. Our changes in the fair value of financial liabilities measured at FVPL were related to the derivative financial liabilities arising from anti-dilution rights granted to pre-IPO investors. These anti-dilution rights were terminated upon Listing.

(3) We adjust changes in the carrying amount of redemption liabilities because it was non-cash in nature. We recognized the redemption liabilities at present value of redemption amounts, with changes in such carrying amounts being booked in profit or loss, arising from redemption rights issued to pre-IPO investors. These redemption rights issued were terminated upon Listing and the carrying amount of the redemption liabilities has been re-designated from liabilities to equity as a result of the automatic conversion into H Shares upon Listing, such that the net liabilities position would turn into net assets.

(4) We adjust listing expense because they were related to the Global Offering.

## LIQUIDITY AND CAPITAL RESOURCES

The Group finances its operations primarily through a combination of (i) capital contributions from shareholders; and (ii) operating cash flows.

The cash and cash equivalents of the Group increased from RMB442.6 million as of December 31, 2025 to RMB831.4 million as of June 30, 2026. This was mainly attributable to net proceeds of approximately HK$655.4 million from the Global Offering (after deduction of the underwriting fees and commissions and the related costs and expenses) in the Reporting Period. These inflows were partially offset by cash used in operating activities and lease payments during the period.

As of June 30, 2026, the cash and cash equivalents of the Group were primarily denominated in Renminbi, Hong Kong dollars and U.S. dollars.

### Borrowings

As of June 30, 2026, we did not have any external borrowings or loans (as of December 31, 2025: Nil).

### Lease liabilities

The lease liabilities of the Group were relatively stable at RMB10.2 million as of June 30, 2026 and RMB11.8 million as of December 31, 2025.

### Pledge of assets

As of June 30, 2026, the Group did not have any pledged assets (as of December 31, 2025: Nil).

### Foreign exchange exposure

We generate substantially all of our revenue in RMB during the Reporting Period, while the net proceeds from the Global Offering were in HK$. Fluctuations in the exchange rate between the RMB and the HK$ will affect the relative purchasing power in RMB in terms of the proceeds from our Global Offering. We manage our foreign exchange risk by performing regular reviews of our Group's net foreign exchange exposures and trying to minimize these exposures through natural hedges, wherever possible.

### Contingent liabilities

As of June 30, 2026, the Group did not have any significant contingent liabilities (as of December 31, 2025: Nil).

## SIGNIFICANT INVESTMENTS HELD, MATERIAL ACQUISITIONS AND DISPOSALS OF SUBSIDIARIES, ASSOCIATES AND JOINT VENTURES

As of June 30, 2026, save as disclosed in this announcement, the Company had no significant investments.

During the Reporting Period, save as disclosed in this announcement, the Company had no material acquisitions and disposals of subsidiaries, associates and joint ventures.

## FUTURE PLANS FOR MATERIAL INVESTMENTS OR CAPITAL ASSETS

Save as disclosed in the section headed "Future Plans and Use of Proceeds" in the Prospectus, we did not have detailed future plans for material investments or capital assets.

## EMPLOYEES AND REMUNERATION POLICY

As of June 30, 2026, the Group had a total of 859 employees (as of December 31, 2025: 721). In the Reporting Period, the Group's staff cost (including share-based payments) amounted to approximately RMB145.4 million (for the six months ended June 30, 2025: approximately RMB105.3 million).

The Group's employee remuneration policy is determined by taking into account factors such as the remuneration in the local market, the overall remuneration level in the industry, operating efficiency, position and employees' performance. In addition to basic salaries, our employees are entitled to annual discretionary bonuses based on our employees' job performance. We have share incentive schemes to motivate our key management and skilled personnel.

## SIGNIFICANT EVENTS AFTER THE REPORTING PERIOD

The Company is not aware of any material subsequent events from June 30, 2026 to the date of this announcement.

## COMPLIANCE WITH THE CORPORATE GOVERNANCE CODE

As the H Shares of the Company were listed on the Main Board of the Stock Exchange on February 13, 2026, the Corporate Governance Code as set out in Appendix C1 to the Listing Rules is only applicable to the Company since the Listing Date.

The Company is committed to the establishment of stringent corporate governance practices and procedures. Since the Listing Date, the Company has adopted and applied the code provisions of the Corporate Governance Code, and will review its corporate governance practices and procedures from time to time to ensure that they comply with the statutory standards. During the period from the Listing Date and up to the date of this announcement, the Company has complied with all applicable code provisions set out in the Corporate Governance Code.

## COMPLIANCE WITH THE MODEL CODE

Since the Listing Date, the Company has adopted the Model Code as the code of conduct regulating dealings in securities of the Company by its Directors and employees who are in possession of inside information in relation to the Group or the Company's securities.

In response to specific enquiries made by the Board, all Directors confirmed that they have complied with the provisions of the Model Code since the Listing Date and up to the date of this announcement.

## PURCHASE, SALE OR REDEMPTION OF LISTED SECURITIES OF THE COMPANY

Neither the Company nor its subsidiaries had purchased, sold or redeemed any of the securities (including treasury shares) of the Company listed on the Stock Exchange during the period from the Listing Date and up to the date of this announcement. As of June 30, 2026, the Company did not hold any treasury shares.

## INTERIM DIVIDEND

The Board has resolved not to recommend the payment of an interim dividend for the six months ended June 30, 2026.

## AUDIT COMMITTEE

The Audit Committee comprises Mr. Zhang Yifan, Mr. Li Shu Pai and Mr. Ma Yeming, and Mr. Li Shu Pai is the chairman of the Audit Committee. The Audit Committee has reviewed the unaudited consolidated financial statements of the Group for the Reporting Period with the management of the Group. The Audit Committee has also discussed matters in relation to the accounting policies and practices adopted by the Company and internal control, risk management and financial reporting with management of the Group.

## SCOPE OF WORK OF THE AUDITOR

The Company's auditor, KPMG, has reviewed the unaudited condensed consolidated financial statements of the Group for the Reporting Period in accordance with Hong Kong Standard on Review Engagements 2410 "Review of Interim Financial Information Performed by the Independent Auditor of the Entity" issued by the Hong Kong Institute of Certified Public Accountants.

## CONSOLIDATED STATEMENT OF PROFIT OR LOSS

*For the six months ended 30 June 2026 - unaudited*
*(Expressed in Renminbi ("RMB"))*

| | | Six months ended 30 June | |
|---|---|---|---|
| | | **2026** | 2025 |
| | *Note* | ***RMB'000*** | *RMB'000* |
| **Revenue** | *3* | **295,563** | 173,469 |
| Cost of sales | | **(163,479)** | (106,687) |
| **Gross profit** | | **132,084** | 66,782 |
| Selling and marketing expenses | | **(43,895)** | (38,128) |
| Administrative expenses | | **(81,020)** | (47,543) |
| Research and development expenses | | **(65,643)** | (36,161) |
| Impairment losses on financial assets and contract assets | | **(4,966)** | (4,593) |
| Other income and loss, net | | **(2,280)** | 402 |
| **Loss from operations** | | **(65,720)** | (59,241) |
| Finance costs | | **(179)** | (197) |
| Changes in the fair value of financial assets measured at fair value through profit or loss ("FVPL") | | **3,968** | 100 |
| Changes in the carrying amount of redemption liabilities | | **(18,318)** | (64,745) |
| Changes in the fair value of financial liabilities measured at FVPL | | **–** | (5,744) |
| **Loss before taxation** | | **(80,249)** | (129,827) |
| Income tax credit | *4* | **1,283** | 2,174 |
| **Loss for the period** | | **(78,966)** | (127,653) |
| **Loss per share** | | | |
| Basic and diluted (RMB) | *5* | **(0.20)** | (0.38) |

## CONSOLIDATED STATEMENT OF PROFIT OR LOSS AND OTHER COMPREHENSIVE INCOME

*For the six months ended 30 June 2026 - unaudited*
*(Expressed in RMB)*

| | | Six months ended 30 June | |
|---|---|---|---|
| | | **2026** | 2025 |
| | *Note* | ***RMB'000*** | *RMB'000* |
| **Loss for the period** | | **(78,966)** | (127,653) |
| **Other comprehensive income for the period (after tax)** | | | |
| Item that may be reclassified subsequently to profit or loss: | | | |
| Exchange differences on translation of financial statements of overseas subsidiaries | | **(2,240)** | (51) |
| **Total comprehensive income for the period** | | **(81,206)** | (127,704) |

## CONSOLIDATED STATEMENT OF FINANCIAL POSITION

*At 30 June 2026 – unaudited*
*(Expressed in RMB)*

| | Note | At 30 June 2026 RMB'000 | At 31 December 2025 RMB'000 |
|---|---|---|---|
| **Non-current assets** | | | |
| Property and equipment | | **10,936** | 6,908 |
| Right-of-use assets | | **9,516** | 11,313 |
| Intangible assets | | **721** | 771 |
| Contract assets | | **16,177** | 15,151 |
| Restricted cash | | **805** | 1,909 |
| Deferred tax assets | | **14,674** | 13,391 |
| | | **52,829** | 49,443 |
| **Current assets** | | | |
| Inventories and contract costs | | **123,166** | 78,232 |
| Contract assets | | **29,470** | 25,266 |
| Trade and bills receivables | *6* | **395,002** | 292,045 |
| Prepayments, deposits and other receivables | | **52,140** | 26,248 |
| Financial assets measured at FVPL | | **140,277** | – |
| Restricted cash | | **2,268** | 22,223 |
| Cash and cash equivalents | | **831,426** | 442,641 |
| | | **1,573,749** | 886,655 |
| **Current liabilities** | | | |
| Trade payables | *7* | **192,912** | 142,827 |
| Other payables and accruals | | **71,640** | 49,953 |
| Contract liabilities | | **100,804** | 92,128 |
| Lease liabilities | | **3,351** | 3,255 |
| Redemption liabilities | | **–** | 2,063,421 |
| Provisions | | **25,380** | 16,207 |
| | | **394,087** | 2,367,791 |
| **Net current assets/(liabilities)** | | **1,179,662** | (1,481,136) |
| **Total assets less current liabilities** | | **1,232,491** | (1,431,693) |

| | Note | At 30 June 2026 RMB'000 | At 31 December 2025 RMB'000 |
|---|---|---|---|
| **Non-current liabilities** | | | |
| Lease liabilities | | **6,838** | 8,565 |
| Deferred income | | **26,584** | 11,209 |
| | | **33,422** | 19,774 |
| **NET ASSETS/(LIABILITIES)** | | **1,199,069** | (1,451,467) |
| **CAPITAL AND RESERVES** | | | |
| Share capital | | **40,043** | 37,240 |
| Reserves | | **1,159,026** | (1,488,707) |
| **TOTAL EQUITY/(DEFICIT) ATTRIBUTABLE TO EQUITY SHAREHOLDERS OF THE COMPANY** | | **1,199,069** | (1,451,467) |

## 1 BASIS OF PREPARATION

Beijing Haizhi Technology Group Co., Ltd. (the "Company", formerly known as Beijing Haizhi Technology Group Limited (北京海致科技集團有限公司)) was incorporated in Beijing, the People's Republic of China (the "PRC") on 23 August 2013 as a limited liability company. On 9 January 2025, the Company was converted from a limited liability company into a joint stock limited liability company. The Company's shares have been listed on the Main Board on the Stock Exchange of Hong Kong Limited ("the Stock Exchange") since 13 February 2026 by way of its initial public offering (the "Listing").

The Company and its subsidiaries (collectively, the "Group") are principally engaged in provision of Atlas Graph Solutions and Atlas AI Agent in the PRC.

This interim financial report has been prepared in accordance with the applicable disclosure provisions of the Rules Governing the Listing of Securities on The Stock Exchange of Hong Kong Limited, including compliance with International Accounting Standard ("IAS") 34, *Interim financial reporting*, issued by the International Accounting Standards Board ("IASB"). It was authorised for issue on 25 August 2026.

The interim financial report has been prepared in accordance with the same accounting policies adopted in the 2025 annual financial statements, except for the accounting policy changes that are expected to be reflected in the 2026 annual financial statements. Details of any changes in accounting policies are set out in Note 2.

The preparation of an interim financial report in conformity with IAS 34 requires management to make judgements, estimates, and assumptions that affect the application of policies and reported amounts of assets and liabilities, income and expenses on a year-to-date basis. Actual results may differ from these estimates.

This interim financial report contains condensed consolidated financial statements and selected explanatory notes. The notes include an explanation of events and transactions that are significant to an understanding of the changes in financial position and performance of the Group since the 2025 annual financial statements. The condensed consolidated interim financial statements and the accompanying notes do not include all of the information required for a full set of financial statements prepared in accordance with IFRS Accounting Standards as issued by the IASB.

The interim financial report is unaudited, but has been reviewed by the audit committee of the Company and approved for issue by the board of director on 25 August 2026. The interim financial report has also been reviewed by KPMG in accordance with Hong Kong Standard on Review Engagements 2410, *Review of interim financial information performed by the independent auditor of the entity*, issued by the Hong Kong Institute of Certified Public Accountants ("HKICPA").

The financial information relating to the financial year ended 31 December 2025 that is included in the interim financial report as comparative information does not constitute the Company's statutory annual consolidated financial statements for that financial year but is derived from those financial statements.

## 2 Changes in accounting policies

The IASB has issued a number of amendments to IFRS Accounting Standards that are first effective for the current accounting period. None of these developments have had a material effect on the Group's unaudited interim financial report. The Group has not applied any new standard or interpretation that is not yet effective for the current accounting period. The Group is assessing the impact of such standards and will adopt the relevant standards in the subsequent periods as required.

## 3 Revenue and segment reporting

### (a) Revenue

The principal activities of the Group are providing Atlas Graph Solutions and Atlas AI Agent. Further details regarding the Group's principal activities are disclosed in Note 3(b).

***Disaggregation of revenue***

Disaggregation of revenue from contracts with customers by major products or service lines is as follows:

| | Six months ended 30 June | |
|---|---|---|
| | **2026** | 2025 |
| | ***RMB'000*** | *RMB'000* |
| **Revenue from contracts with customers within the scope of IFRS 15** | | |
| Disaggregated by major products or service lines | | |
| – Atlas Graph Solutions | **180,974** | 124,822 |
| – Atlas AI Agent | **114,589** | 48,647 |
| | **295,563** | 173,469 |

Disaggregation of revenue from contracts with customers by the timing of revenue recognition and by geographic markets is disclosed in Notes 3(b)(i) and 3(b)(ii), respectively.

### (b) Segment reporting

The Group manages its businesses by business lines (products and services). In a manner consistent with the way in which information is reported internally to the Group's chief operating decision maker ("CODM") for the purposes of resource allocation and performance assessment, the Group has presented the following two reportable segments. No operating segments have been aggregated to form the following reportable segments.

- Atlas Graph Solutions: this segment includes providing data intelligence solutions, which are built on the Group's robust graph computing and data analysis capabilities and platforms, for customers to transform their diverse enterprise data into structured knowledge, enhance reasoning-based analytics and enable intelligent automation.

- Atlas AI Agent: this segment includes providing agentic solutions, which are built on the Group's graph-large language models integration technologies and platforms, for customers to apply in various application scenarios, such as risk identification, marketing analysis, regulatory auditing, equipment management and fault detection.

*(i) **Segment results***

For the purposes of assessing segment performance and allocating resources between segments, the Group's senior executive management monitors the results attributable to each reportable segment on the following bases:

Revenue and expenses are allocated to the reportable segments with reference to sales generated by those segments and the expenses incurred by those. The measure used for reporting segment result is gross profit.

The Group's senior executive management monitors the Group's assets and liabilities as a whole, accordingly, no segment assets and liabilities information is presented.

Disaggregation of the Group's revenue from contracts with customers by the timing of revenue recognition, as well as information regarding the Group's reportable segments as provided to the Group's senior executive management for the purposes of resource allocation and assessment of segment performance for the reporting period is set out below:

| | **Six months ended 30 June 2026** | | |
|---|---|---|---|
| | **Atlas Graph Solutions** | **Atlas AI Agent** | **Total** |
| | ***RMB'000*** | ***RMB'000*** | ***RMB'000*** |
| **Disaggregated by timing of revenue recognition** | | | |
| Point in time | **173,350** | **114,589** | **287,939** |
| Over time | **7,624** | **–** | **7,624** |
| Revenue from external customers and reportable segment revenue | **180,974** | **114,589** | **295,563** |
| Reportable segment gross profit | **71,509** | **60,575** | **132,084** |

| | Six months ended 30 June 2025 | | |
|---|---|---|---|
| | Atlas Graph Solutions | Atlas AI Agent | Total |
| | *RMB'000* | *RMB'000* | *RMB'000* |
| **Disaggregated by timing of revenue recognition** | | | |
| Point in time | 115,639 | 48,647 | 164,286 |
| Over time | 9,183 | – | 9,183 |
| Revenue from external customers and reportable segment revenue | 124,822 | 48,647 | 173,469 |
| Reportable segment gross profit | 43,376 | 23,406 | 66,782 |

*(ii) Geographic information*

The following table sets out information about the geographical location of the Group's revenue from external customers. The geographical location of customers is based on the location at which the services were rendered or the goods were delivered. All of the Group's non-current assets are located or allocated to operations located in the PRC.

| | Six months ended 30 June | |
|---|---|---|
| | **2026** | 2025 |
| | ***RMB'000*** | *RMB'000* |
| Chinese Mainland (place of domicile) | **295,563** | 173,376 |
| Hong Kong | **–** | 93 |
| | **295,563** | 173,469 |

## 4 INCOME TAX IN THE CONSOLIDATED STATEMENT OF PROFIT OR LOSS

| | Six months ended 30 June | |
|---|---|---|
| | **2026** | 2025 |
| | ***RMB'000*** | *RMB'000* |
| **Deferred tax** | | |
| Origination and reversal of temporary differences | **1,283** | 2,174 |

## 5 LOSS PER SHARE

### (a) Basic loss per share

The calculation of basic loss per share is based on the loss attributable to ordinary equity shareholders of the Company of RMB67,444,000 (six months ended 30 June 2025: RMB50,595,000) and the weighted average of 334,079,000 ordinary shares (six months ended 30 June 2025: 134,296,000 ordinary shares) in issue during the interim period.

The Group conducted the share subdivision immediately prior to the Listing, pursuant to which each of share with a par value of RMB1.00 subdivided into 10 Shares with a par value of RMB0.1 each. Accordingly, the weighted average number of shares throughout the periods presented has been adjusted retrospectively for the impact of such share subdivision.

Loss attributable to ordinary equity shareholders of the Company:

| | Six months ended 30 June | |
|---|---|---|
| | **2026** | 2025 |
| | ***RMB'000*** | *RMB'000* |
| Loss attributable to all equity shareholders of the Company | **(78,966)** | (127,653) |
| Allocation of loss attributable to ordinary shares with redemption rights | **11,522** | 77,058 |
| Loss attributable to ordinary equity shareholders of the Company | **(67,444)** | (50,595) |

Weighted average number of ordinary shares:

| | Six months ended 30 June | |
|---|---|---|
| | **2026** | 2025 |
| | ***'000*** | *'000* |
| Issued ordinary shares at 1 January | **37,240** | 33,051 |
| Effect of ordinary shares issued | **–** | 526 |
| Effect of issuance of ordinary shares by initial public offering | **21,042** | – |
| Effect of ordinary shares with redemption rights | **(5,741)** | (20,147) |
| Effect of shares subdivision | **281,538** | 120,866 |
| Weighted average number of ordinary shares at 30 June | **334,079** | 134,296 |

### (b) Diluted loss per share

Share options granted under the 2025 share incentive plan, ordinary shares with redemption rights and anti-dilution rights were not included in the calculation of diluted loss per share as their inclusion would have been anti-dilutive. Accordingly, the amounts of diluted loss per share were the same as basic loss per share.

## 6 TRADE AND BILLS RECEIVABLES

| | **At 30 June 2026** | At 31 December 2025 |
|---|---|---|
| | ***RMB'000*** | *RMB'000* |
| Bills receivables | **4,323** | 4,011 |
| Trade receivables | **434,448** | 327,161 |
| Gross amount of trade and bills receivables | **438,771** | 331,172 |
| Less: credit loss allowance | **(43,769)** | (39,127) |
| | **395,002** | 292,045 |

**Aging analysis**

As of the end of the reporting period, the aging analysis of trade and bills receivables of the Group, based on the date of revenue recognition and net of loss allowance, is as follows:

| | **At 30 June 2026** | At 31 December 2025 |
|---|---|---|
| | ***RMB'000*** | *RMB'000* |
| Within 1 year | **305,453** | 204,183 |
| 1 – 2 years | **68,935** | 68,207 |
| 2 – 3 years | **17,716** | 15,986 |
| 3 – 4 years | **1,434** | 2,434 |
| 4 – 5 years | **1,464** | 1,235 |
| | **395,002** | 292,045 |

All of the trade and bills receivables are expected to be recovered within one year.

## 7 TRADE PAYABLES

| | **At 30 June 2026** | At 31 December 2025 |
|---|---|---|
| | ***RMB'000*** | *RMB'000* |
| Trade payables | **192,912** | 142,827 |

As at the end of the reporting period, the aging analysis of trade payables of the Group, based on the invoice date, is as follows:

| | **At 30 June 2026** | At 31 December 2025 |
|---|---|---|
| | ***RMB'000*** | *RMB'000* |
| Within 1 year or on demand | **122,585** | 70,893 |
| Between 1 year and 2 years | **48,877** | 53,217 |
| Over 2 years | **21,450** | 18,717 |
| | **192,912** | 142,827 |

All of the trade payables are expected to be settled within one year or are repayable on demand.

## 8 DIVIDENDS

No dividends have been declared or paid by the Company during the six months ended 30 June 2025 and 2026.

## PUBLICATION OF INTERIM RESULTS AND INTERIM REPORT

This announcement is published on the website of the Stock Exchange at www.hkexnews.hk and the website of the Company at www.haizhi.com. The interim report of the Company for the six months ended June 30, 2026 will be made available for review on the same websites in due course.

## DEFINITIONS

In this announcement, the following expressions shall have the following meanings unless the context otherwise requires.

| | |
|---|---|
| "AI" | artificial intelligence, an area of computer science that focuses on machinery simulation of intelligence displayed by humans and other animals |
| "AI agent" or "intelligent agent" | an autonomous intelligent entity leveraging LLM technology capable of executing tasks with consistency and reliability, dynamically adjusting decisions according to real-time scenarios and adopting optimal strategies suitable for the current context |
| "AI solution" | a comprehensive bundle that may include software, such as an AI agent, hardware, services or processes which leverages AI technologies to solve specific problems, automate tasks or enhance decision-making across various domains |
| "Audit Committee" | the audit committee of the Company |
| "Board" | the board of Directors |
| "our Company" or "the Company" | Beijing Haizhi Technology Group Co., Ltd. (北京海致科技集團股份有限公司), a joint stock limited company established in the PRC, whose H Shares are listed on the Main Board of the Stock Exchange (stock code: 02706) |
| "Corporate Governance Code" | Corporate Governance Code, as set out in Appendix C1 to the Listing Rules |
| "Director(s)" | Director(s) of our Company |
| "Global Offering" | the Hong Kong public offering and the international offering of the Company, the details of which are described in the Prospectus |
| "graph" | a structural network composed of nodes and edges, intertwined by various types of entities and relationships |

| | |
|---|---|
| "graph database" | a database system that stores graph-structured data (nodes, edges, properties) in order to queries and process analysis complex relationships |
| "Group" or "our Group" or "the Group" or "we" or "us" | our Company and its subsidiaries from time to time |
| "H Share(s)" | ordinary share(s) in the share capital of the Company with nominal value of RMB0.1 each, which is/are listed on the Main Board of the Stock Exchange and subscribed for and traded in Hong Kong dollars |
| "Hong Kong" | the Hong Kong Special Administrative Region of the PRC |
| "HK$" | Hong Kong dollars, the lawful currency of Hong Kong |
| "IFRS Accounting Standards" | IFRS Accounting Standards as issued by the International Accounting Standards Board |
| "Listing" | the listing of the H Shares on the Main Board of the Stock Exchange on the Listing Date |
| "Listing Date" | February 13, 2026, the date on which our H Shares were listed on the Main Board of the Stock Exchange |
| "Listing Rules" | the Rules Governing the Listing of Securities on The Stock Exchange of Hong Kong Limited |
| "LLM(s)" | large language model, computerized language model that utilizes artificial neural network with a multitude of parameters, trained on substantial amounts of unlabeled text through self-supervised or semi-supervised learning techniques |
| "Main Board" | the stock market (excluding the option market) operated by the Stock Exchange which is independent from and operated in parallel with GEM of the Stock Exchange |
| "Model Code" | Model Code for Securities Transactions by Directors of Listed Issuers, as set out in Appendix C3 to the Listing Rules |
| "PRC" or "China" | the People's Republic of China, excluding, for the purposes of this announcement only, Hong Kong, Macau Special Administrative Region of the People's Republic of China and Taiwan |

| | |
|---|---|
| "Prospectus" | the prospectus of the Company dated February 5, 2026 in relation to the Global Offering and the Listing |
| "Reporting Period" | the six months ended June 30, 2026 |
| "RMB" | Renminbi, the lawful currency of the PRC |
| "R&D" | research and development |
| "Stock Exchange" | The Stock Exchange of Hong Kong Limited |
| "%" | Percent |

By order of the Board
**Beijing Haizhi Technology Group Co., Ltd.**
**Mr. Yang Zaifei**
*Chief Executive Officer and Executive Director*

Beijing, the PRC, August 25, 2026

*As at the date of this announcement, the Board comprises (i) Mr. Ren Xuyang, Mr. Yang Zaifei, Ms. Yang Juan and Mr. Wan Pengjiang as executive Directors; (ii) Mr. Li Jiaqing and Mr. Li Hongtao as non-executive Directors; and (iii) Mr. Zhang Yifan, Mr. Jiang Tian, Mr. Li Shu Pai and Mr. Ma Yeming as independent non-executive Directors.*