OVERVIEW

Who We Are

We are the largest independent distributed AI storage solution provider in China, as measured by installed capacity in 2025. We focus on delivering enterprise-grade AI storage solutions that empower organizations to efficiently integrate data, decision-making, and operations at scale. We have developed our proprietary XSEA technology focusing on making data storage fast, resilient and efficient at scale, as well as our proprietary XScale technology focusing on managing large volumes of data in a flexible and scalable way. Leveraging these two proprietary technologies, we have developed our two main solutions, namely data lake storage solutions and storage solutions for Training and Inference ("**T&I**"), through which we enable deployment and implementation of AI storage into our enterprise customers' business operations, addressing critical storage needs during their AI-driven transformation. The top five distributed AI storage solution providers in China accounted for a combined market share of 47.9% by installed capacity in 2025, with us representing 10.6% of the market, and standing as China's second-largest distributed AI storage solution provider and the largest independent distributed AI storage solution provider.

We are highly focused on developing an independent, inclusive and compatible ecosystem for AI data infrastructure, with solutions deployed across multiple industry verticals. As of June 30, 2026, we had served more than 1,600 enterprise customers across key industry verticals, including intelligent manufacturers, telecommunications operators, internet and technology companies. Our solutions support diverse application scenarios such as AI LLMs, machine learning, high-performance computing, large-scale data analytics, scientific research, and retrieval-augmented generation (RAG). As of June 30, 2026, we had served 100 out of China's Top 500 largest companies by 2024 revenue. See "— Our Competitive Strengths Diversified and Retentive Customer Base Across the AI Value Chain." Our overall net dollar retention rate reached 160.0% in the six months ended June 30, 2026.

Our Journey

Since our founding in 2015, we have consistently placed technological innovation at the core of our strategy. Through independent research and development, we have built technological advantages. At the early stage of our founding in 2015, we strategically chose to enter the private cloud and virtualization market. During this phase, we developed early cloud storage products XEBS (block storage) and XEOS (object storage) systems, establishing the core roadmap of "software-defined" architecture and provided solutions to our first lighthouse customers. In 2017, we introduced our XEDP (Unified Data Platform) that broke down data silos by enabling interoperability across block, file, and object storage protocols, setting the standard of "one data center, one storage architecture" for our customers, reducing IT operation and maintenance complexity and improving data management efficiency. Since July 2022, we shifted towards building end-to-end AI data infrastructure. Our file system XGFS and object storage system XEOS have been integrated into the full lifecycle of AI training and inference, providing a solid data foundation for enterprises undergoing AI transformation. In January 2026, we launched AIMesh which further unified data access protocols and enables data flow from AI data lakes to T&I environments, providing a more efficient and unified data foundation for AI data infrastructure.

Our Market Opportunities

Against the backdrop of the rapid growth of China's AI infrastructure market, storage lies at the foundation of AI infrastructure. With AI workloads driving exponential growth in storage demand, traditional storage solutions face three major bottlenecks:

- The "I/O Wall": caused by high-concurrency read/write operations during AI training.
- The "Memory Wall": resulting from limited GPU memory capacity during AI inference.
- The "Gravity Wall": arising from high migration costs and inefficiencies when data is scattered across multiple systems, including different storage types and deployment environments.

Compared to legacy storage solutions, distributed storage solutions address these challenges through multi-node parallel read/write, data sharding with load balancing and redundancy for fault tolerance, delivering linear scalability in performance, capacity and reliability. This architecture transforms distributed storage from a general-purpose infrastructure into a critical AI enabler, unlocking significant market opportunities. China's distributed AI storage solution market reached RMB10.0 billion by 2024 revenue, and is expected to grow to RMB63.1 billion by 2030, representing a CAGR of 35.9%. Beyond the technical bottlenecks described above, the explosive growth of AI technologies has also triggered structural supply-demand imbalances and persistent price increases since the fourth quarter of 2025. To manage these rising hardware costs of traditional storage architectures, enterprises must adopt more cost-effective, scalable storage solutions.

Our Commercialization Path

We focus on serving leading enterprises in key industries by delivering customer-centric AI storage solutions. See "— Our Solutions." We have extended our focus on industry-leading customers across sectors including intelligent manufacturing, financial institutions, telecommunications operators and technology enterprises, which allows us to gain deeper insights into sector-specific challenges and attract more customers in a broad range of industry sectors. The number of enterprise end customers that entered into contracts with us increased modestly from 2023 to 2024, exceeded 520 in 2025, and reached over 260 in the six months ended June 30, 2026. In 2023, 2024, 2025 and the six months ended June 30, 2025 and 2026, we recorded revenue of RMB166.8 million, RMB172.5 million, RMB267.0 million, RMB99.5 million and RMB185.8 million, respectively, demonstrating a proven record of our commercialization strategy to date and the recognition of our solutions by enterprise customers.

OUR COMPETITIVE STRENGTHS

An Independent Distributed AI Storage Solution Provider in China with Competitive Barriers that Safeguard Sustainable Growth

We are the largest independent distributed AI storage solution provider in China as measured by installed capacity in 2025. As of June 30, 2026, we had been focusing on distributed storage for over a decade, with our technologies and products applied in cloud computing. We launched the industry's first distributed storage platform supporting all major protocols and the earliest to enable block, object, and file storage within a single cluster. We also introduced the industry's first distributed storage platform supporting heterogeneous chips within a single system.

Our capabilities and solutions have been recognized by various authoritative institutions and top industry associations. See "— Awards and Recognitions." We actively participate in national standard-setting and have contributed to several key standards, including *GB/T 37737-2019 Distributed Block Storage System Technical Requirements*, *GB/T 37732-2019 Cloud Storage System Service Interface Functions*, *GB/T 44271-2024 Edge Cloud General Technical Requirements*, and *GB/T 17966-2024 Floating-Point Operations for Microprocessor Systems*. In addition, we have led and participated in drafting multiple industry standards for object storage and distributed storage, continuously promoting standardization and technological advancement.

As the AI industry continues to develop, our capabilities, solutions and technical experience in AI data infrastructure support our competitive position and market participation.

AI-Driven Independent R&D Capabilities as the Foundation for Innovation

During the Track Record Period, we have focused on developing our storage solutions for AI applications, leveraging a decade of experience in distributed storage as well as a robust architectural foundation which enables us to build technologies tailored for AI scenarios. Specifically, as generative AI evolves toward inference-driven AI, we developed storage architectures that combine object storage as the core of the data lake featuring high throughput and scalability, layered with file systems for interface compatibility. Our tiered storage design places hot, warm, and cold data in separate pools, namely cache, and object storage, balancing performance and cost effectively.

We maintain an R&D strategy focused on distributed storage architecture and large-scale data management technologies. Our key storage solutions, which were all developed in-house, incorporate capabilities across storage engines, large-scale network transmission and scalable metadata management. See "— Our Proprietary Technologies." Our technologies benefit our customers by enabling faster AI model training, smoother operations for large-scale applications, and lower infrastructure complexity. By reducing delays and simplifying data management, our solutions help customers save time, cut costs, and unlock the full value of their data assets.

Building on this technological foundation, we introduced our data lake storage solutions and storage solutions for T&I, which support our customers' data access requirements for large-model training and real-time inference. In January 2026, we launched AIMesh. See "— Our Solutions — Commercialization — Strategic Role within AIMesh." AIMesh unifies data, memory, and compute resources in a single architecture to meet the unique demands of T&I at scale. It ensures that processes run efficiently without bottlenecks, reducing training time, improving inference speed, and lowering overall costs. It also simplifies operations by providing a single, consistent foundation for managing data across the entire AI lifecycle, enabling businesses to deploy AI solutions faster and more reliably.

Through more than a decade of continuous R&D and large-scale deployments, we have developed a unified R&D process that spans product design, development, testing, deployment, upgrades, and operations. This process, which was developed through the evolution of our products from XEBS, XEOS and XGFS to our all-flash XEBS-XINFINI products, increases our R&D efficiency by allowing cross-product sharing of foundational technical capabilities. Such R&D process also ensures consistent reliability, scalability, and operational experience across block, object, file, and all-flash products, forming a critical competitive advantage in the AI infrastructure market.

As of June 30, 2026, we operated R&D centers in Beijing, Shenzhen, Chengdu, Wuhan and Xiamen, with technical staff accounting for over 40% of our total employee headcount. We have applied for and obtained more than 570 intellectual property rights. Our core team members come from leading internet companies, cloud providers, and global storage leaders such as EMC, NetApp, and Huawei, bringing extensive experience in distributed storage and enterprise-grade product development.

Full-Protocol, Cross-Scenario Solutions Covering the Entire AI Data Lifecycle

Built on our self-developed capabilities on high-concurrency storage architecture, elastic metadata engine, and DataOS for heterogeneous hardware, our AI storage solutions offer full-protocol compatibility across block, file and object systems. This ensures compatibility across AI T&I, machine learning, high-performance computing, large-scale data analytics, scientific research, and retrieval-augmented generation (RAG). Our full-protocol, cross-scenario capability enables unified and consistent data foundations across diverse compute architectures, workloads, and deployment models, significantly reducing integration and migration costs.

Our unified data platform supports the entire AI data lifecycle from ingestion and cleansing to T&I, allowing data to flow freely across different stages of the AI data lifecycle such as data lakes, training, and inference after being written only once and maximizing data value throughout its lifecycle. Historically, these layers were often operated as separate silos, and customers frequently had to create duplicate copies of the same dataset. By supporting multiple access protocols within a unified platform and namespace, our solutions allow the same underlying data to be accessed directly for training without requiring such copying, enabling datasets to be used more quickly by GPU clusters and improving overall efficiency.

Independent and Highly Compatible Ecosystem

We have built an ecosystem that combines the performance of an integrated platform with the flexibility of a compatible standard and establishes an ecosystem barrier which we believe is hard to replicate by new industry entrants.

- **Multi-Compute Adaptability:** our AI storage solutions support compatibility with different types of computing chips, such as Nvidia, Hygon, and Ascend. This flexibility is highly valued by our customers who often use a mix of hardware from different vendors. Our technology provides a compatible, future-ready storage foundation that supports diverse computing environments without locking customers into a single vendor.

- **Cloud-Native & Sovereign:** Our solutions integrate with cloud-native environments while ensuring that data remains fully under the customers' sovereignty — an essential requirement for regulated sectors such as intelligent manufacturing and financial institutions. Customers gain the agility and scalability of a cloud-like experience on-premises, without compromising security or compliance.

- **Industry Ecosystem Collaboration:** We work with major computing ecosystem partners and our products integrate smoothly with mainstream architectures such as Intel and AMD. This compatibility and scalability enable our storage solutions to adapt naturally to diverse computing environments, supporting key workloads such as high-performance data analytics and improving customers overall data processing efficiency.

Together, these capabilities create an ecosystem that supports both domestic and global mainstream computing platforms. For customers, this means they can confidently build core AI data infrastructure that works across hybrid cloud environments and demanding scenarios like autonomous driving, large-model training, industrial AI, and financial technology. Our solutions help enterprises solve key challenges such as data ingestion, storage, management, retrieval, protection, and mobility, unlocking the full value of their data and making it dynamic and actionable.

Diversified and Retentive Customer Base Across the AI Value Chain

Through years of R&D and business operations, we have established industry recognition for our technologies and product capabilities through landmark projects for industry leaders, enabling us to expand our customer base across the entire AI value chain, from upstream model and computing power providers to downstream software developers and industry customers. We strategically focus on serving leading large-model companies, top research labs, and major software developers. Our customer portfolio includes multiple China Top 500 enterprises across intelligent manufacturing, telecommunications operators and technology enterprises. Besides, given the infrastructure nature of our AI storage solutions, our products are deeply embedded in customers' core data workflows, resulting in high switching costs and high deployment-level customer stickiness. Such stickiness is reflected in our overall net dollar retention rate of 160.0% in the six months ended June 30, 2026.

As of June 30, 2026, our customer base covered three of the top five LCD panel manufacturers, four of the top ten power battery manufacturers, three of the top five photovoltaic manufacturers in China by their 2025 revenue; in the financial sector, our customer base covered five of 12 joint-stock banks, 13 of 33 banks with assets over RMB1.0 trillion as of June 30, 2026, and four of the top ten life insurers in China by their 2025 revenue.

We have achieved large-scale commercial deployment in core AI applications such as autonomous driving, large-model training and industrial AI. For example, in autonomous driving, we support the full workflow from production-line storage to R&D and road test data processing. In large-model training, we provide file systems and hybrid cloud data mobility for clusters ranging from thousands to tens of thousands of GPUs. In industrial AI, we are deeply embedded in core production processes of leading panel and semiconductor customers.

Through continuous software updates, extended services, and ecosystem integration, we address enterprise needs for system stability and long-term operational support, reinforcing our position as a trusted AI infrastructure partner.

Experienced Management Team and Diverse Investor Base

Our management team brings deep expertise in AI storage and AI infrastructure. Our Chairman and Founder, Mr. Xu Xin, has been involved in China's open-source cloud computing sector since its early development and is a co-founder of China's first OpenStack vendor, with nearly a decade of technical and managerial experience. He is recognized as a Zhongguancun High-End Talent, a Beijing Science and Technology Commission expert, and a Haidian Youth Federation member. Our Director and CTO, Mr. Wang Haomai, previously led storage platform development at UnitedStack, establishing China's first unified framework for cloud block storage and driving industry standards for distributed storage.

As of the Latest Practicable Date, we had completed eight rounds of financing, attracting a diverse group of institutional investors, including Boyu Capital, Northern Light Venture Capital, China Reform Venture Capital, Qiming Venture Partners, Redpoint, Tencent, NEA, CICC Alpha, Legend Capital, Broad Vision Funds, Source Code Capital, Kunlun, Yishang, Shanghai Guoxin, V Fund and Hundsun Technologies. Multiple investors, including Northern Light, Qiming, and Redpoint, have participated in successive rounds, underscoring their confidence in our long-term growth potential. Backed by these investors' deep industry insights and resources, we are well-positioned to capture growth opportunities, strengthen our market leadership, and deliver sustainable success.

OUR STRATEGIES

We have formulated the following development strategies:

Strengthen R&D Capabilities and Attract Top Talent

Our strategy on R&D and talent attraction is designed to address the evolving enterprise storage needs in the age of AI, breaking through traditional bottlenecks in scale, performance, and latency, and advancing toward AI-native data platforms. Specifically, our R&D efforts will focus on the following areas:

- ***R&D for next-generation AI-native storage platform.*** We will continue to develop next-generation AI-native storage platforms for large-model training, AI inference and online services. This initiative represents an upgrade of our existing offerings, including MeshFS, MeshFusion, which form part of AIMesh, an upgraded iteration of our existing software products, and is intended to help customers manage and use massive AI datasets more efficiently, improve GPU utilization, shorten model training cycles and reduce infrastructure costs. See "— Our Solutions — Commercialization — Strategic Role within AIMesh." Specifically, we plan to achieve sustained throughput at the TB/s level within a single cluster, reduce training data cache access latency to as low as approximately 20 microseconds, and progressively support the storage needs of ultra-large clusters with up to 100,000 GPUs. We also plan to build a unified object and file storage foundation, enabling customers to support AI applications such as online inference, vector search and content generation through a simpler storage architecture, with less data movement and higher deployment efficiency.

- ***R&D for expansion and vertical adaptation of AI storage portfolio.*** We will expand our AI storage solution portfolio and adapt our solutions to key industries and typical use cases such as financial services, intelligent manufacturing, media video cloud and scientific research. We also plan to enhance our data management and migration capabilities in multi-cloud and hybrid-cloud environments and develop edge AI storage products and lightweight solutions, enabling customers to manage and use AI data more flexibly across on-premises, cloud and edge environments.

- ***R&D for ecosystem compatibility development.*** We will improve the compatibility of our products with mainstream hardware, cloud platforms and AI software frameworks through ecosystem adaptation and validation of our existing solution offerings, lowering the barriers for customers to deploy, integrate and operate AI storage solutions. We also plan to build new platform capabilities, including an API/SDK-based developer integration platform to enable third-party developers and ISVs to integrate with and extend our storage platform more easily.

To that end, we will continue investing in next-generation AI-native storage platforms, distributed high-performance storage engines, and data management capabilities tailored for emerging AI scenarios such as large-model T&I, AIGC, and multimodal applications. Our goal is to enhance our product portfolio across object storage, distributed file storage, and block storage, and further develop our AI storage infrastructure capabilities.

We also plan to recruit more talents across various technical fields, including large models, distributed systems, cloud-native storage, data lakes, and high-performance computing. By attracting experienced architects from leading internet companies, cloud vendors, and research institutions, and improving our expert and engineer career development pathways, we aim to build an experienced R&D team which we believe is key to maintaining our long-term competitive barriers and achieving sustainable growth.

Expand and Optimize Our Solution Portfolio

We will continuously enrich our solution portfolio and broaden our coverage of AI storage application scenarios to meet the real-world technical needs and business scenarios of customers in different industries and at various stages of development with a focus on next-generation AI-native storage platform development and related vertical expansion and adaptation. By tailoring solutions to industry-specific requirements and standardizing them into ready-to-use products, we expect to significantly enhance delivery speed, reduce implementation complexity, and improve replicability. We expect this approach to accelerate the adoption of our AI storage solutions in key industries, enabling enterprises to deploy them more quickly, cost-effectively, and easily.

We will also systematically upgrade our XSEA and XScale technology foundations to adapt to large-model training, inference, and multimodal workloads. Specifically, under our AIMesh strategy, we will focus on unifying the "AI data and memory plane" that integrates data, memory, and compute. AIMesh is designed to unify data and memory into a more integrated layer that sits closer to the computing environment. By adopting a converged compute-and-storage approach, AIMesh shortens the data access path and simplifies the overall system stack, improving performance efficiency and reliability for mission-critical AI workloads.

Expand Sales Channels and Build Industry-Focused Teams

We will strengthen investments in sales and delivery teams, shifting from a traditional geography-based model to an industry-focused approach. Through a "key account list" strategy, we will prioritize sectors with AI infrastructure demand, such as intelligent manufacturing, financial institutions, telecommunications operators and technology enterprises. In support of this strategy, we intend to expand sales coverage by building vertical-focused sales, pre-sales and delivery teams with expertise (i) for the finance sector, in financial workloads and AI infrastructure performance requirements; (ii) for the internet sector, in large-model training and inference across the AI lifecycle; and (iii) for the manufacturing sector, in industrial AI, machine vision and data-intensive workloads. In parallel, we plan to extend our compatibility and strengthen our performance validation across diverse chipset platforms by procuring representative systems, testing system specifications and performance benchmarks, and conducting performance tuning, with the objective of forming repeatable best practices that can be applied across mainstream hardware environments.

We also plan to deepen ecosystem cooperation around the AI T&I framework by engaging broadly with framework providers, being developers and maintainers of mainstream AI model training and inference software frameworks, and industry partners. We intend to co-develop and standardize interfaces (including through SDKs where appropriate) to facilitate integration, encourage partners and downstream developers to build add-on features on top of our platform, and introduce incentive programs to support third-party development and accelerate ecosystem adoption.

We will cultivate AI-focused sales capabilities by co-creating storage solutions with leading large-model vendors and independent software vendors (ISVs), enabling our teams to gain deep insights into AI training, inference workflows, and data lake management. This will help us build a sales force with hybrid expertise in both industry and technology. Additionally, we will expand localized sales and delivery teams to strengthen our full-cycle services capabilities for key customers and ecosystem partners, accelerating solution deployment and customer stickiness.

We will also deepen collaboration with leading AI technology companies and cloud service providers by engaging in joint application scenarios development, testing, and ecosystem building, which we believe will accelerate our own product refinement and technology iteration, as well as foster industry-wide collaboration.

We also plan to establish an AI infrastructure customer demonstration and ecosystem validation center, which will be used to (i) host live demonstrations of large-scale AI infrastructure in real operating environments, showcasing the high-throughput and low-latency performance of our storage systems under full AI training workloads; (ii) support regional sales and marketing teams by providing localized services covering POC testing, solution presentations and on-site delivery of integrated systems; and (iii) organize developer salons and AI technology forums, bringing customers' algorithm engineers together with our storage experts to facilitate direct technical engagement and strengthen our market presence.

Build an Inclusive and Compatible Ecosystem

We will continue enhancing product compatibility and technical adaptability across heterogeneous compute environments, creating a highly interoperable AI data infrastructure platform. Our ecosystem strategy will cover diverse hardware systems, broad industry scenarios, and key partners, enabling deployment in demanding environments such as intelligent manufacturing, financial institutions, telecommunications operators and technology enterprises.

We will maintain collaboration with upstream and downstream partners to apply distributed storage, data lakes, and AI T&I technologies in more application scenarios. By building a compatible platform for ecosystem partners and developers, we will support third-party applications and AI platforms to integrate and innovate on our storage foundation, creating network effects.

Our inclusive ecosystem approach compatibility in both technology and partnerships. Through sustained investment in core software capabilities and cross-ecosystem compatibility, we aim to establish a future-ready, inclusive and compatible AI data infrastructure platform.

OUR SOLUTIONS

We primarily offer (i) AI storage solutions, comprising our data lake storage solutions and storage solutions for T&I, which may be offered either on a standalone basis or as a bundled offering, depending on the customer's specific needs, (ii) AI storage services, including tail-end warranty, maintenance or renewals services that accompany our AI storage solutions. The use of "AI" in our solution names is because our storage solutions are specifically optimized for AI workloads, and not because they are storage products embedding AI algorithms or performing autonomous learning or decision-making functions.

AI Storage Solutions

A simple way to explain the difference between and the primary roles of our two AI storage solutions, namely Data Lake Storage Solutions and Storage Solutions for T&I, is through a comparison of these solutions to the process through which a personal computer stores data. A personal computer typically has a (i) large-capacity drive (commonly known as "hard drive") that keeps all files (including photos, videos and documents) for the long term at relatively low cost, and (ii) a smaller-capacity but much faster drive, such as a random access memory (RAM) or cache, that holds the files being actively worked on when a software is actively running on the computer. Our AI data lake storage solutions play a role broadly similar to the former at enterprise scale: they store, organize and govern an enterprise's large volumes of data for the long term, keeping such data ready for AI and other business use. Our storage solutions for T&I play a role broadly similar to the latter: they deliver data at very high speed to GPU computing clusters while AI models are being trained or are generating inference outputs. The two segments are often deployed together, with data flowing from the data lake to T&I scenarios as and when needed. This analogy is a simplified illustration only: it does not seek to imply that we develop or sell these personal computer related products or solutions. Rather, our solutions are enterprise-grade distributed systems deployed across large numbers of servers, with capacity, performance and reliability requirements far beyond those of personal computing devices.

Within each of our Data Lake Storage Solutions and Storage Solutions for T&I categories, we further offer various types of sub-solutions depending on customers' data usage scenarios and the storage formats required. Enterprise data is stored and accessed in three basic formats, and the choice among them depends on how the data is used.

- **Block storage** divides storage capacity into fixed-size "blocks" made directly available to an application, in the way a computer uses its internal disk; it is the fastest format with the lowest latency and is typically used by databases, virtual machines and other performance-critical systems.

- **File storage** organizes data as files within folders, in the way documents are kept on a shared office drive; multiple users and systems access the same files through familiar directory paths, which suits collaboration scenarios and AI training datasets that are read file by file.

- **Object storage** keeps each piece of data as an independent "object" with a unique identifier and descriptive tags in a flat structure without folder hierarchies, similar to the way a cloud photo album manages a massive number of items; it scales almost without limit and is best suited for very large volumes of unstructured data such as images, videos, backups and archives.

The following table summarizes the key functions and features of each of our business segments and their respective sub-segments:

Business segment	Subsegment	Key functions, features and differentiation	Specific value added along the data lifecycle/type of processing applied to data	Relevant quantitative supporting metrics
1. AI storage solutions				
(1) Data lake storage solutions	**Overall segment**	Helps enterprise customers collect, store, organize and manage large volumes of raw and unstructured data before such data is used for AI training, analytics or other business applications. It is differentiated from AI T&I storage solutions because it focuses on long-term data retention, organization, governance and lifecycle management, rather than real-time high-speed access during model training or inference.	Adds value at the data ingestion, storage, preparation and governance stages. It helps customers consolidate scattered data, manage different data formats, reduce data silos, and prepare data for future AI use.	See below
	XEOS	An enterprise-grade object storage solution for long-term storage of large volumes of unstructured data, such as images, videos and documents. It is differentiated from XEDP and EasyData because it mainly focuses on object storage and long-term retention, which means data is stored as individual units in a flat structure, making it easier to manage large volumes of unstructured data, and is particularly suited for use cases such as content repositories, backups and archival storage, as compared to file storage and block storage.	Adds value at the data storage, retention, governance and distribution stages. It stores raw and unstructured data, supports backup and archiving, prevents data loss through object lock and version control, optimizes small-file storage, and automatically moves data among hot, warm and cold storage tiers (with hot, warm and cold determined based on access frequency of data following rules set by customers), thus lowering the cost of long-term data retention for customers while keeping data accessible.	Validated in hundred-PB single-cluster production deployments; supports a single bucket scaling to hundreds of billions of objects; supports stable throughput and latency; supports policy-driven hot/warm/cold tiering.

Business segment	Subsegment	Key functions, features and differentiation	Specific value added along the data lifecycle/type of processing applied to data	Relevant quantitative supporting metrics
	XEDP	A unified data storage platform integrating file, object and block storage into one system. It is differentiated from XEOS because it is not limited to object storage, and is differentiated from EasyData because it focuses more on unified storage infrastructure rather than data collection and preparation.	Adds value at the data consolidation and management stage. It allows customers to manage file, object and block data under one platform, so data used by different departments or systems can be stored and accessed under consistent policies.	Supports file, object and block storage in one system; works together with XEOS, XGFS and XEBS to provide unified data infrastructure. Its performance metrics may vary depending on the specific deployment architecture and customer workload profile.
	EasyData	A data management platform for handling large volumes of unstructured data from diverse sources. It is differentiated from XEOS and XEDP because it focuses on data ingestion, metadata management and data preparation, rather than underlying storage capacity or unified storage infrastructure.	Adds value at the data ingestion, preparation and utilization stages. It automatically collects data from different devices and systems, unifies metadata management, supports search and access control, and applies processing such as tagging, image analysis, data labeling and format conversion when instructed by customers.	EasyData has been tested on a single bucket with two billion unstructured objects. It allows users to manage all this data in a unified way, organize it with tags, and quickly search across it using flexible search methods, with results returned in seconds. It also allows up to 300 different user groups, so access to data can be controlled and managed securely.
(2) Storage solutions for T&I	**Overall segment**	Provides fast and reliable data access for AI model training and inference. It is differentiated from AI data lake storage solutions because it focuses on using and delivering data at high speed during GPU-based training and inference tasks, rather than storing and governing data for long-term use.	Adds value at the model training, inference and high-performance data access stages. It helps feed processed and performance-critical datasets to GPUs and computing clusters, reducing bottlenecks such as insufficient I/O performance and memory limitations.	See below
	XEBS-XINFINI	A fully SSD-based block storage solution for mission-critical and performance-intensive workloads, such as large-scale AI training, real-time inference, vector databases, distributed databases and enterprise cloud platforms. It is differentiated from XEBS because it is designed for more demanding and latency-sensitive workloads, while XEBS is designed for general-purpose applications with more balanced requirements on performance, scalability and cost efficiency.	Adds value at the real-time data access, AI training, AI inference and vector search stages. It enables many systems to access storage in parallel with low latency, reduces data-access delays, and supports high-speed workloads that require predictable performance, such as large-scale AI training, real-time inference, vector databases, distributed databases and enterprise cloud platforms.	Provides high IOPS and microsecond-level latency; read and write latency can be as low as approximately 100 microseconds; built on a Shared-Everything architecture where every server node can directly access all storage devices in the cluster; uses all-flash SSD pools and supports high throughput.

Business segment	Subsegment	Key functions, features and differentiation	Specific value added along the data lifecycle/type of processing applied to data	Relevant quantitative supporting metrics
	XEBS	A distributed block storage solution supporting virtualization, databases, containers and hybrid cloud environments. It is differentiated from XEBS-XINFINI because XEBS is positioned as a general-purpose block storage solution for core enterprise applications, while XEBS-XINFINI is an all-flash, ultra-low-latency solution for mission-critical AI and high-speed workloads.	Adds value at the enterprise application and database data-access stage. It provides stable block storage for systems that require reliable low-latency operation, such as databases, virtualized environments and hybrid cloud applications.	In a performance test using three servers and a high-speed network, XEBS was able to handle a large number of data operations at once. It achieved over 110,000 small write operations per second and nearly 500,000 read operations per second, with 12 users generating the workload simultaneously.
	XGFS	A distributed file storage system for AI training, big data processing and other data-intensive workloads. It is differentiated from XEBS-XINFINI and XEBS because XGFS is a file storage solution, while XEBS-XINFINI and XEBS are block storage solutions, with file storage organizing data in hierarchical file systems for shared access and content management, typically used for AI training datasets, big data processing and collaborative workloads, whereas block storage stores data in fixed-size blocks for high-performance, low-latency transactional workloads such as databases and enterprise applications.	Adds value at the large-dataset access and AI training stage. It supports unified file access, broad protocol compatibility and high-concurrency access to large datasets, helping GPU clusters and data-intensive applications read and write many files efficiently.	Supports high throughput and low latency when handling billions of files; scalable to thousands of nodes; supports multiple data access protocols and domestic hardware ecosystems.
2. AI storage services	**Tail-end warranty, maintenance and renewals services**	Provides post-deployment support for customers that have purchased AI storage solutions. It is differentiated from AI storage solutions because it does not primarily involve initial system deployment, but supports system continuity after delivery.	Adds value at the post-deployment operation, maintenance and expansion stage. It helps customers keep deployed systems stable, renew warranties or maintenance arrangements, and expand storage capacity when business data increases.	Customers who purchase our AI storage solutions are also provided with AI storage services.

AI infrastructure today faces three structural challenges known as the "Three Walls", including the "I/O Wall", the "Memory Wall" and the "Gravity Wall". In addition, the global AI storage industry is also facing increasingly severe supply-chain challenges since the fourth quarter of 2025. The explosive growth of AI technologies has triggered intense demand for high-performance storage media, such as enterprise-grade NAND Flash and DRAM, leading to structural supply-demand imbalances and persistent price increases. This places heavy cost pressure on traditional storage architectures that rely on low efficiency and high redundancy. To manage these rising hardware costs, enterprises must improve storage-density efficiency and adopt more cost-effective, scalable storage solutions.

Against this backdrop, our AI storage solutions fall into two complementary categories, each addressing a different part of the AI pipeline:

Aspect	Data Lake Storage Solutions	Storage Solutions for T&I
Main Goal	Long-term data retention, organization, and governance	Fast data delivery for AI model T&I
Data Type	Raw, unstructured and diverse datasets	Processed, performance-critical and small working datasets
Key Optimization Focus	Scalability, cost-efficiency and lifecycle management	Throughput, latency, cache expansion and real-time access
Typical Use Case	Storing and preparing data for future AI use	Providing access to data for GPU-based T&I tasks

The following chart illustrates the relationship among our technology layer, product layer and solution offerings:

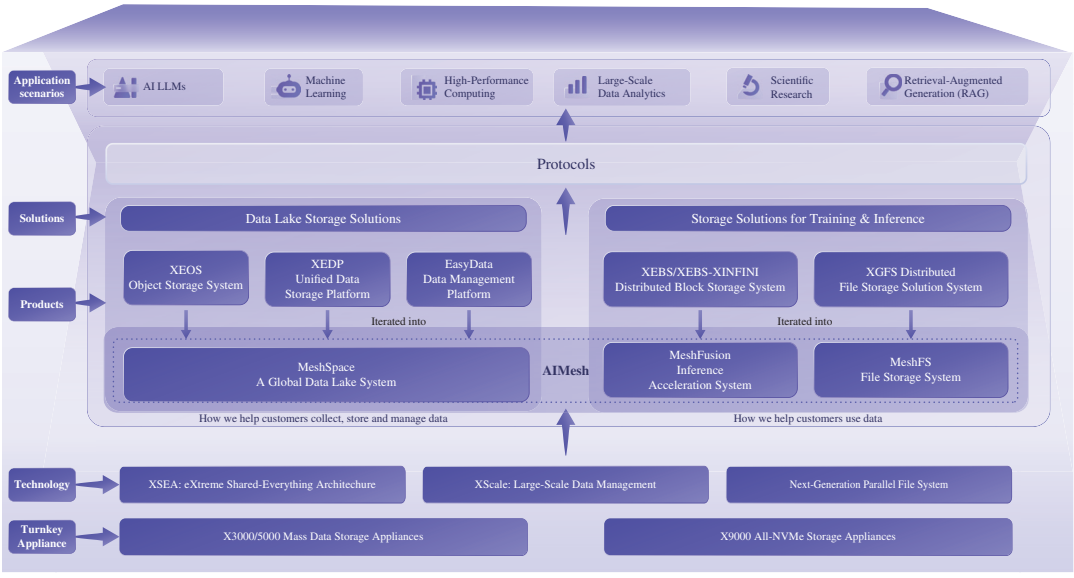

Our products can be configured and combined into solution offerings, including data lake storage solutions and storage solutions for T&I, which are the commercialized solutions delivered to customers. Our AI storage solutions may consist of software sold independently or as part of integrated hardware appliances, and customers may elect to procure either software-only solutions or turnkey appliances under each business subsegment. When delivering software alone, customers

are required to procure hardware themselves that meets our specific technical requirements. For turnkey appliance deliveries, we source hardware from third-party suppliers and perform development and customization work to embed our proprietary software.

Our Solutions and AI Data Lifecycle

The following chart illustrates a standard AI data lifecycle, from data ingestion and preparation to model T&I, with a feedback loop whereby inference results may be fed back into earlier stages for continuous optimization. In practice, customers are not required to adopt our solutions for all stages of the AI data lifecycle, nor are they required to follow the entire workflow depicted. Customers may deploy our products selectively, in combination with their existing infrastructure or third-party solutions, depending on their technical architecture, business needs and operational preferences. The AI data lifecycle illustrated in the chart is typically implemented and operated within customers' own environments.

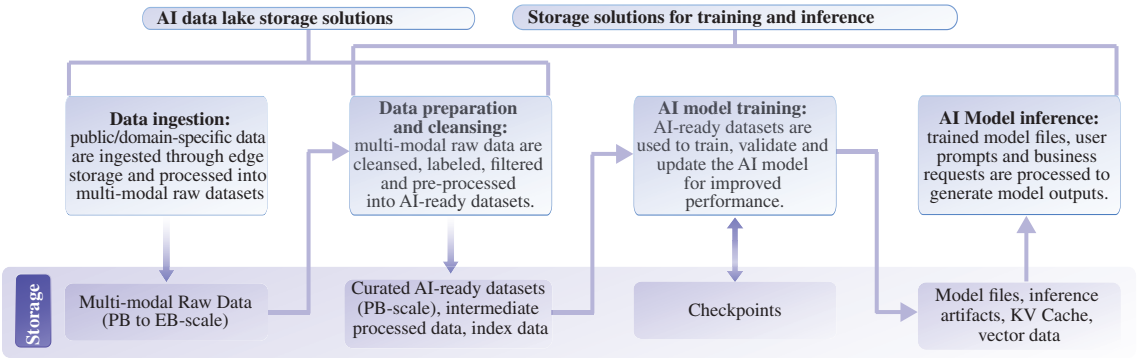

Our AI storage solutions streamline the entire data lifecycle, from data ingestion and storage to management and utilization, by providing an integrated platform designed to meet the requirements of AI applications, such as high throughput, low latency and scalability for massive datasets.

(i) **Data ingestion:** Raw unstructured data from multiple sources is automatically aggregated into the AI storage system without the need for manual scripting. The system continuously monitors source-side file changes and moves new data into centralized storage, which is commonly called "AI data lake", according to user-defined schedules. This process ensures timely, efficient and consistent ingestion of large and diverse datasets into AI data lake required for AI workloads;

(ii) **Storage:** Our AI storage solutions provide a unified and scalable repository capable of handling massive volumes of data across diverse AI-driven applications. Unlike traditional storage solutions that operate in silos, our solutions enable better resource sharing. They also include built-in features like lossless compression and image optimization, which help reduce storage space and overall cost without compromising data quality;

(iii) **Management:** Our centralized management method allows customers to monitor and allocate resources across both file and object storage clusters. Our customers can also use the platform to schedule data collection tasks, automate data tiering based on usage patterns and manage user access permissions, all from a single interface, which is essential to maintaining data availability and governance in AI environments; and

(iv) **Utilization:** Our customers benefit from an interface that supports easy upload, download, preview and search for data through a user-friendly portal. Our AI storage solutions support access to different types of storage solutions and provide real-time monitoring of CPU, memory and disk usage to maintain performance for AI T&I. It also enables data classification using automatic or rule-based tagging. Additionally, built-in tools for image processing, decompression, and customizable workflows help speed up data usage and improve operational efficiency for AI applications.

Beyond operational efficiency, our AI storage solutions also address the emerging strategic need for enterprises to retain full control over proprietary data. As general-purpose models become commoditized, enterprise-specific data, such as manufacturing parameters, proprietary codebases or confidential customer information, has become the key source of differentiation. Private, on-premises distributed storage ensures that such core assets remain fully controlled and isolated from external environments, reducing risks of data leakage, model contamination or loss of intellectual property. It also enables enterprises to build their own long-term "data flywheel," where business feedback data is continuously accumulated and used to refine internal models. This closed-loop, self-reinforcing system strengthens data assets over time and forms a sustainable competitive advantage.

Data Lake Storage Solutions

Our data lake storage solutions form the foundational data layer for enterprises building AI capabilities. They are designed to aggregate, cleanse and retain massive volumes of raw and unstructured data, serving as the starting point for the full AI data lifecycle. From the perspective of the "three walls", our data lake storage solutions primarily address the Gravity Wall. By aggregating multi-source, unstructured data into a single governed namespace and automating hot-warm-cold tiering across heterogeneous media, our data lake storage solutions keep data logically close to compute, reduce fragmented copies and manual data transfers, and provide a single system of record that downstream T&I layers can access on demand. Our customers may elect to procure either software-only solutions or turnkey appliances under each business subsegment.

The chart below demonstrates our data lake storage solutions (XEOS, XEDP and EasyData) and their distinctive features:

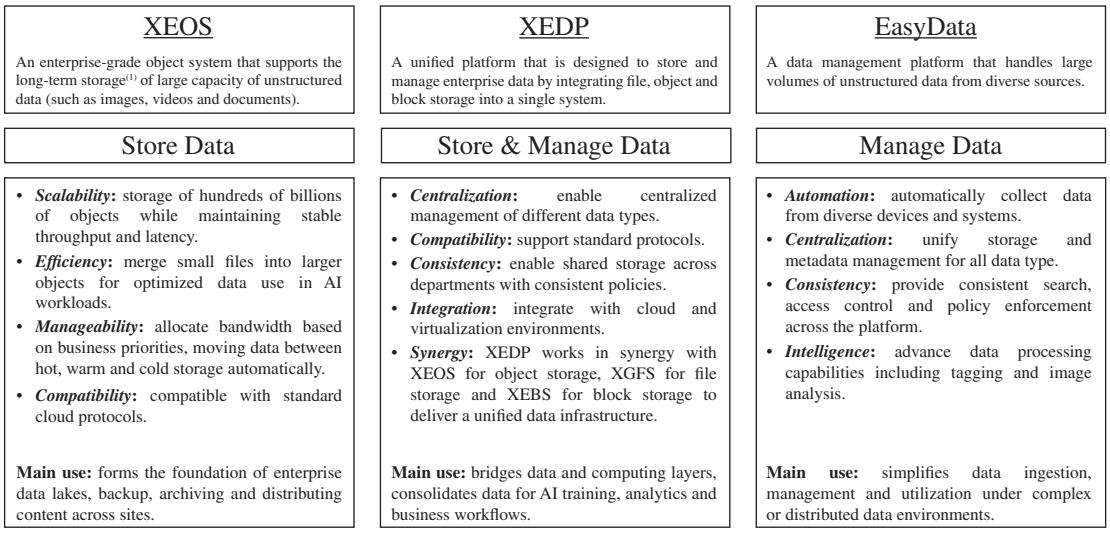

Note:

(1) Long-term storage refers to storage designed for the long-term retention and archival stage within the data lifecycle. In industry practice, such storage is typically associated with data retention periods of five years or more, depending on customers' specific requirements.

XEOS: Scalable Object Storage System for Enterprise Data

XEOS (XSKY Enterprise Object Storage) is an enterprise-grade object storage for long-term retention, governance and distribution of unstructured data. It serves as the foundational system for enterprise data lakes, backup and archiving, and cross-site content distribution, providing a scalable, highly available and easy-to-manage object storage system. XEOS supports both x86 and ARM processors, including domestic processors, and has been validated in hundred-PB single-cluster production deployments. XEOS helps customers continuously track how frequently each piece of data is accessed and, following rules pre-set by the customer's administrator (for example, moving data untouched for 90 days to a lower tier), automatically relocate data among storage media of different performance and cost levels — frequently accessed "hot" data stays on high-performance media for fast retrieval, less-visited "warm" data resides on standard media, and rarely accessed "cold" data is moved to the lowest-cost, high-capacity media. No manual effort is required and data remains retrievable at all times. This significantly lowers the cost of long-term data retention while keeping data accessible, in a manner similar to a smartphone photo application that keeps recent photos on the device while automatically archiving older photos to the cloud.

Built on standard protocols, XEOS integrates with most upstream applications without system changes. It provides a comprehensive S3-compatible implementation in China and supports advanced features such as object lock and version control to prevent data loss and retain historical data versions. XEOS has the following key capabilities.

- **Massive Single-Bucket Scalability:** XEOS uses a flat, non-directory object architecture with key-range distribution and independently managed metadata services. This enables a single bucket to scale to hundreds of billions of objects while maintaining stable throughput and latency, supporting cloud-native expansion for large-scale AI training datasets.

- **Small-File Optimization for Higher Performance and Lower Cost:** For AI workloads dominated by small files, XEOS aggregates small files into larger objects and applies multi-level caching with adaptive data-merging mechanisms. These measures reduce I/O overhead, improve performance and lower storage costs.

- **Priority-Based Resource Allocation:** XEOS provides built-in quality-of-service (QoS) controls that allocate resources based on business priority. High-priority applications receive preferential performance, while background workloads are governed through user-level and bucket-level controls to maintain system stability.

- **Automated Cross-Site Lifecycle Management and Multi-Site Deployment:** XEOS supports policy-driven tiering of hot, warm and cold data across private, public and hybrid environments. Data can be automatically migrated for cost-efficient retention and quickly recalled when needed, with support for active-active access and cross-region disaster recovery.

The XEOS workflow enables end-to-end data movement through a unified platform combining cache pools and heterogeneous data pools. Through its scalability, compatibility with mainstream protocol, resource governance and lifecycle management, XEOS provides a durable and cost-efficient foundation for enterprise data lakes and long-term AI development.

Use Case: Large Language Model (LLM) Infrastructure

In large language model infrastructure, XEOS can serve as the scalable object storage foundation for internet platforms, AI companies, research institutes and state-owned enterprises operating AI compute centers. LLM training and inference workloads generate and reuse large volumes of text, image, video, log, model-version and checkpoint data. When such data is distributed across isolated systems, data preparation becomes slow, duplicated copies increase storage cost and cross-cloud data mobility becomes difficult. XEOS addresses these issues through an enterprise-grade object storage architecture built on standard S3 protocols, enabling customers to retain, retrieve and manage large-scale unstructured datasets through a common object storage layer. Its large-scale single-bucket scalability supports datasets that may reach petabyte scale and contain a large number of files, while small-file optimization, QoS controls and automated lifecycle management help customers manage hot, warm and cold data according to workload priority and retention requirements. With multi-site deployment and cross-site lifecycle management capabilities, XEOS provides a durable and cost-efficient foundation for long-term LLM data retention, model iteration, backtesting and retraining.

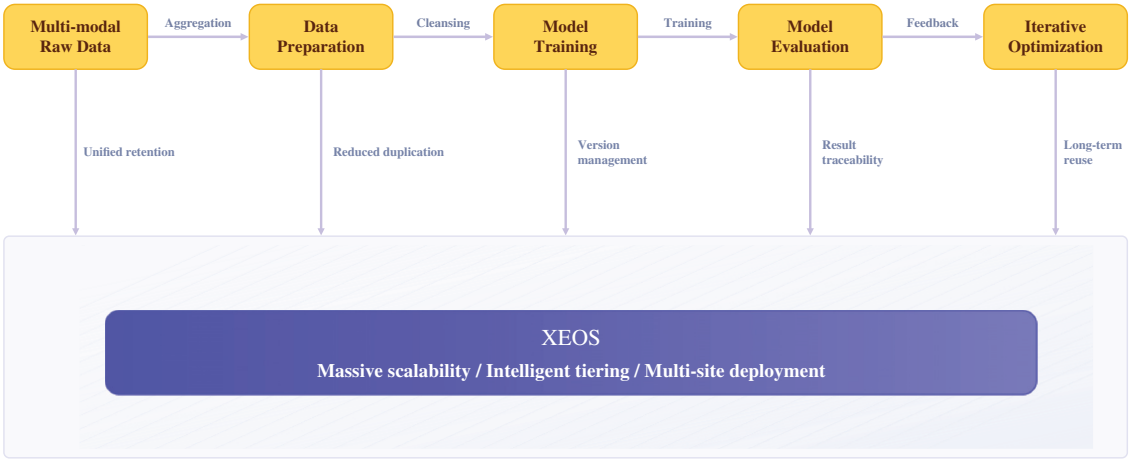

XEDP: Unified Enterprise Data Storage Platform

XEDP (XSKY Enterprise Data Platform) is a unified data storage platform that integrates file, object and block storage into a single system. Customers typically deploy XEDP together with other AI storage products, including XEOS, XGFS and XEBS. In simple terms, file storage organizes data in folders for easy human access, object storage manages large volumes of data as independent units with flexible metadata, and block storage splits data into smaller pieces to enable high-performance processing. The platform supports three core storage models as described below.

(i) **File Storage:** File storage uses a hierarchical structure of files and directories and presents data through familiar drive letters and folder paths. It aligns with common enterprise usage patterns, supports collaborative workflows and remains compatible with traditional IT systems, making data browsing, management and access intuitive and efficient.

(ii) **Object Storage:** Object storage organizes data as independent objects with metadata and unique identifiers in a flat structure, typically grouped within logical containers known as buckets. This model supports efficient management of massive unstructured datasets at petabyte scale and beyond, is optimized for distributed access and cloud-native environments, and is commonly used for cold data storage, long-term archiving and AI model version management.

(iii) **Block Storage:** Block storage divides data into fixed-size blocks managed through logical volumes, delivering high IOPS and low latency. It is used for mission-critical workloads such as financial transactions and high-performance databases, where stability and predictability are essential. In AI workloads, block storage supports high-performance training and short-term caching for online inference.

Without a centralized platform, organizations often deploy separate storage systems across teams or applications, leading to inefficient resource usage, inconsistent data protection and higher operational complexity. XEDP addresses these issues by enabling shared storage, applying consistent policies, quotas, performance controls and security settings across all storage types while preserving existing applications.

Use Case: FinTech AI Solutions

XEDP provides a unified enterprise data storage platform for licensed consumer-finance institutions, banks, insurers and securities firms building private knowledge bases, risk models, compliance analytics and AI-enabled customer-service applications. These applications depend on large training datasets and substantial volumes of historical unstructured data, such as tickets, call recordings, screenshots and customer interaction records, which must be retained, governed, retrieved and reused for model training, inference and continuous model refinement. Financial institutions also need to move data efficiently across collection, preprocessing, model training and inference while maintaining business continuity, regulatory compliance and consistent performance. Traditional systems often rely on separate file, object and block storage environments, making data movement slow and increasing management complexity. XEDP addresses these challenges by integrating file, object and block storage into a single platform and applying consistent policies, quotas, performance controls and security settings across different storage types. It supports a wide range of standard access protocols, enabling financial customers to consolidate historical data into an enterprise-grade data lake and make such data available for AI workflows without repeated migration. XEDP also provides data protection capabilities such as snapshots, backups, replication and cross-site disaster recovery, helping financial customers shorten model iteration cycles, reduce data silos, transform archival cold data into AI training assets and maintain stable operations for AI workloads.

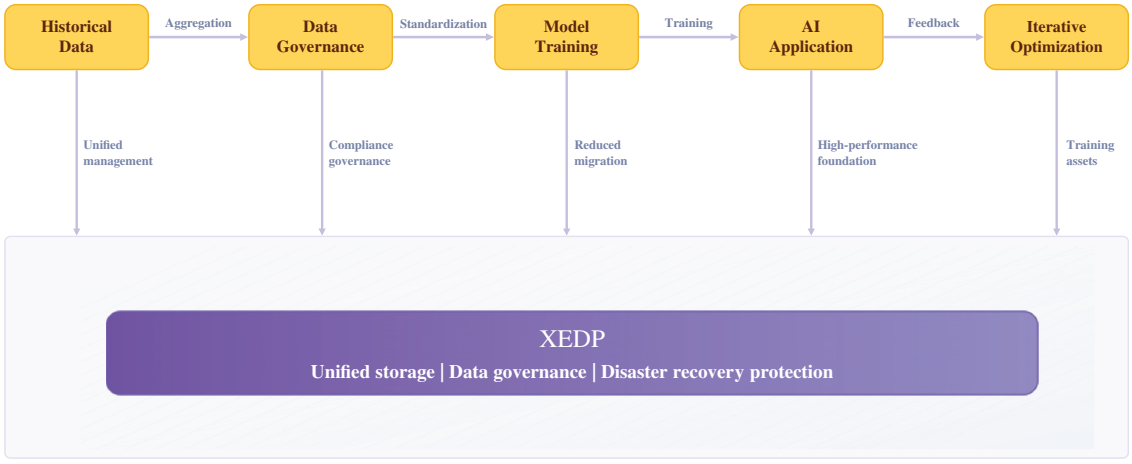

In deployment, XEDP connects data and computing layers to support end-to-end workflows. Data from block and object storage is consolidated into a centralized data lake, supporting access for AI training analytics.

Below is an illustrative example of the user interface of XEDP:

EasyData: Unified Data Management Platform for Large-Scale Unstructured Data

EasyData is an enterprise-level data management platform designed to simplify large and complex data environments. As organizations generate increasing volumes of data across diverse systems, EasyData provides a unified platform for data lifecycle management, covering data ingestion, storage, management and utilization. Positioned above the storage layer (XEOS, XEDP, XEBS and XGFS), EasyData collects data from multiple sources, writes it to underlying storage systems, and enables users to efficiently locate, process and use data.

- **Ingestion:** EasyData automatically collects data from local devices and third-party storage systems through software agents. It supports integration with legacy storage without requiring data migration or changes to business processes, and enables automated data expiration at the source to improve efficiency.

- **Storage:** Collected data is centrally written to object storage or file storage systems, with real-time metadata synchronization to ensure consistent visibility and data integrity across the platform.

- **Management:** EasyData provides unified data search, access control, policy management and operational monitoring. Multi-tenant permission controls and IP-based restrictions enhance security, while automated tagging and policy-based filtering support structured data catalogs.

- **Utilization:** Users access data through a web console for search, preview, download and processing. Processed results can be returned to the platform or consumed directly by AI and analytics systems, eliminating manual copying and format conversion.

At its core, EasyData offers a centralized, web-based management console supporting a multi-tenant architecture. Different departments operate in logically isolated environments with role-based access control for administrators and users. Integration with enterprise authentication systems such as AD and LDAP provides unified identity management and security.

EasyData supports automated ingestion from distributed sources, including local devices, third-party systems using standard protocols, and XSKY's distributed storage infrastructure, without modifying existing applications. Data retrieval is enabled through a metadata-driven query

engine capable of querying large-scale datasets using attributes such as filenames, timestamps, sizes and tags. Built-in processing services include decompression, image processing and automated tagging, with extensibility through custom plugins.

Security and compliance are supported through fine-grained access controls, audit logging, cold data archiving to tape and unified search across hot and cold tiers. Real-time resource monitoring, automated alerts, scheduled tasks and reporting enhance operational efficiency. Overall, EasyData provides a robust and extensible platform for secure, end-to-end enterprise data management in complex and high-growth data environments.

Use Case: Advanced Manufacturing and Industrial AI

EasyData helps manufacturers such as MFlex Electronics convert scattered production-line data into searchable, governed and AI-ready data assets. Manufacturers of LCD panels, printed circuit boards and photovoltaic panels typically operate many production devices and business systems, each generating data in different formats. This makes it difficult to create a unified data asset, and production data is often difficult to clean, label and retrieve for AI model development. At the same time, manufacturers must retain large-scale production data for compliance and traceability while maintaining efficient access. EasyData addresses these challenges through software-agent-based ingestion, access to distributed sources such as local devices and third-party storage systems, and real-time metadata synchronization. Through EasyData governance, the platform can extract metadata such as equipment identifiers, serial information and yield attributes, and build multi-dimensional indexes that make raw production files structured, searchable and usable for AI. By combining unified search, permission control, policy management and lifecycle governance, EasyData supports a closed loop from production-line inference to model training, helping improve quality inspection and customer-issue handling.

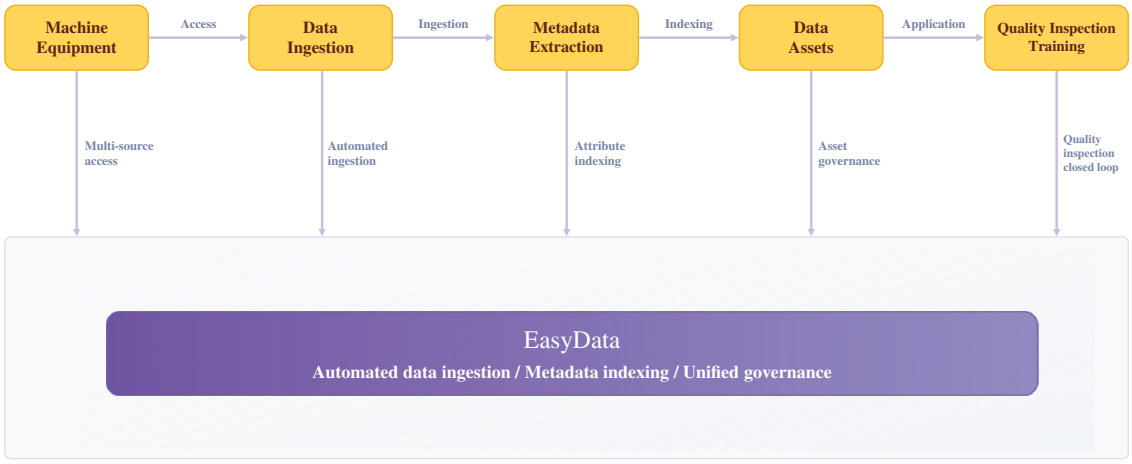

Storage Solutions for T&I

Storage solutions for T&I serve as the engine of AI data infrastructure, responsible for feeding data to GPUs at high speed and overcoming the industry's two major bottlenecks—the IO Wall in large-model training and the Memory Wall in long-context inference. Our core offerings in this business line include XEBS, XEBS-INFINI (block storage system) and XGFS (file storage system), which support all-flash deployments and deliver the throughput, IOPS and ultra-low latency required to keep large GPU clusters fully utilized. Our customers may elect to procure either software-only solutions or turnkey appliances under each business subsegment.

In AI training, data must be delivered to GPU clusters efficiently and consistently. Training workloads typically involve repeated reading of large datasets and frequent writing of intermediate outputs as GPU clusters repeatedly access training data across multiple iterations and continuously write checkpoints, which record a model's state so that training can resume from the latest saved point if interrupted. Our storage solutions for training are designed to support high-bandwidth data loading and stable concurrent access patterns that are common in large-scale training tasks.

In AI inference, the system must respond to end users or downstream applications in real time. Inference workloads therefore emphasize low latency, high concurrency and stable performance. In addition to model files and feature data, inference may require rapid access to auxiliary data structures such as cache data used to accelerate multi-turn interactions and vector data used to support retrieval-based AI applications. Our storage solutions for inference are designed to provide efficient access to such performance-sensitive data and to support production-grade AI applications.

From the perspective of the "three walls", our storage solutions for T&I are designed to break through both the I/O Wall in large-scale training and the Memory Wall in long-context inference. For training workloads, our storage solutions are designed to deliver fast and stable data access, allowing large volumes of training data to be read efficiently and model progress to be saved regularly. This helps ensure that computing resources are used effectively, rather than sitting idle while waiting for data. For inference workloads, our solutions support efficient data caching and memory expansion, enabling AI systems to handle longer conversations and a higher number of simultaneous requests without requiring significant additional hardware investment.

The chart below demonstrates our storage solutions for T&I (XEBS-XINFINI, XEBS and XGFS) and their distinctive features:

XEBS-XINFINI	XEBS	XGFS
An All-flash block storage system that provides high-performance, mission-critical workloads such as large-scale AI training and real-time inference.	A distributed block storage system that supports virtualization, databases, containers and hybrid cloud setups.	A distributed file storage system that facilitates AI training, big data process and enterprise workloads through unified namespace and broad protocol.
Access and Use Data		
• ***Synergy:*** leverage a "Shared-Everything" architecture to deliver maximum speed and consistent reliability across clusters. • ***Accessibility:*** enable direct access to all storage devices to support data availability. • ***Efficiency:*** enable accelerated data transfer and sustained performance. • ***Responsiveness:*** build for ultra-low latency and high throughput to ensure rapid response. **Main use:** powers demanding AI workloads, vector databases and enterprise cloud platforms.	• ***Compatibility:*** work with multiple virtualization platforms and database environments. • ***Responsiveness:*** deliver consistently stable operation with low latency. • ***Flexibility:*** support hybrid cloud operations and enable smooth data movement between on-premises infrastructure and public clouds. **Main use:** provides high-performance storage for core business applications and databases.	• ***Compatibility:*** support multiple data access protocols and provide broad application compatibility. • ***Responsiveness:*** provide high throughput and low latency even when handling billions of files. • ***Scalability:*** expand to thousands of nodes. • ***Localization:*** optimized for domestic hardware ecosystems to maximize compatibility and efficiency. **Main use:** enables fast, stable access to large datasets for AI model training and other data-intensive tasks.

XEBS-XINFINI: All-Flash Block Storage System

XEBS-XINFINI (XSKY Enterprise Block Storage-XINFINI) is our all-flash block storage system designed for mission-critical and performance-intensive workloads, including large-scale AI training, real-time inference, vector databases, distributed databases and enterprise cloud platforms. As a block storage system, it is engineered to maximize IOPS and deliver highly predictable performance. In high-speed analytics and AI training scenarios, it provides high IOPS and microsecond-level latency under both single-stream and highly concurrent workloads, while maintaining stable data access. XEBS-XINFINI integrates with mainstream private cloud and virtualization platforms, including VMware, Kubernetes and OpenStack.

XEBS-XINFINI is underpinned by three core architectural features:

- **Shared-Everything Architecture:** XEBS-XINFINI is built on our eXtreme Shared-Everything Architecture (XSEA), under which every server node can directly access all storage devices in the cluster. This flat architecture eliminates unnecessary data hops, reduces latency variation and delivers uniform performance for distributed workloads. By sharing storage, memory and network resources across the system, XSEA improves hardware utilization and avoids common performance bottlenecks even under heavy concurrent access.

- **End-to-End High Performance Storage Architecture:** The system adopts a fully end-to-end high performance storage architecture, supporting multiple protocols for customer access within the internal storage network. All SSDs are accessed directly without intermediate protocol layers, reducing overhead and latency. This design enables read and write latency as low as approximately 100 microseconds, meeting the requirements of online inference, financial transactions and high-speed search workloads.

- **Single-Tier All-Flash Architecture:** XEBS-XINFINI uses a single tier of SSDs to form a unified all-flash storage pool. Combined with technologies designed to improve storage efficiency, the system achieves high capacity efficiency while maintaining low latency and high throughput, reducing total cost of ownership without sacrificing performance.

XEBS: Distributed Block-Storage System

XEBS (XSKY Enterprise Block Storage) is our distributed block-storage system supporting mainstream virtualization platforms, transactional and analytical databases, container environments and hybrid cloud deployments. Compared with XEBS-XINFINI, which targets ultra-high-performance AI workloads, XEBS is designed as a general-purpose system for a broad range of enterprise applications.

Key features of XEBS include:

- **Multi-Virtualization:** XEBS integrates with leading virtualization platforms and presents storage as a unified shared pool, simplifying capacity management, performance policy enforcement and data protection.

- **Multi-Database:** XEBS supports both transactional and analytical databases with stable low-latency performance. It provides snapshots and clones for backup, recovery and workload isolation, together with performance controls to prevent resource contention.

- **Multi-Container:** XEBS integrates with mainstream container platforms and includes built-in data services, accelerating deployment and reducing operational overhead.

- **Hybrid Cloud:** XEBS enables data movement and backup between on-premises infrastructure and public clouds without application changes, supporting flexible hybrid cloud operations.

XGFS: Distributed File Storage System for Enterprise Workloads

XGFS (XSKY Global File System) is a distributed file storage system designed to support AI training, big data processing and enterprise workloads through a unified global namespace and broad protocol support. It enables different teams and applications to access the same data without duplication or format conversion and delivers high throughput, low latency and stable performance within a single cluster. In all-flash deployments, XGFS can deliver tens of gigabytes per second of bandwidth for AI training environments.

Key features of XGFS include:

- **Full Protocol Interoperability:** XGFS supports multiple protocols within the same global namespace, allowing data written through one protocol to be accessed through another without duplication.

- **High Performance:** XGFS minimizes overhead through a self-developed metadata service, all-flash or hybrid storage pools and the XCache caching layer, delivering low latency even for billions of small files.

- **Scalability and Elastic Expansion:** A small number of metadata nodes can support tens of billions of files, while the data layer scales to thousands of nodes, enabling cost-efficient expansion.

- **Deep Optimization for Domestic Ecosystems:** XGFS is optimized for domestic CPU ecosystems through software-level enhancements and close vendor collaboration, improving resource efficiency.

- **Lower Operational Costs and Dual-Site Deployments:** Erasure coding reduces hardware overhead while maintaining reliability, and active-active dual-site support ensures continuous data availability across locations.

Use Case: Autonomous Driving Data Infrastructure

XGFS can be deployed as the file storage foundation for new-energy vehicle OEMs, Tier-1 suppliers and autonomous-driving technology companies, including intelligent driving and cabin solution providers such as Minieye Technology. Autonomous-driving road tests generate large volumes of camera, LiDAR and radar data every day. These datasets need to be collected, cleaned, labeled and then supplied to model-training pipelines with minimal delay. In a traditional environment, data often passes through separate systems and access protocols across collection, labeling and training, creating duplicated copies and long waiting periods. XGFS addresses these issues through a unified global namespace and multi-protocol access, so that the same dataset can be used across different workflow stages. When deployed on all-flash storage, XGFS provides high-throughput parallel file access to GPU clusters and supports faster perception-model iteration. Historical road-test datasets can be retained on hybrid-flash capacity tiers with erasure coding and compression, helping customers shorten the cycle from road test to model release while controlling storage cost.

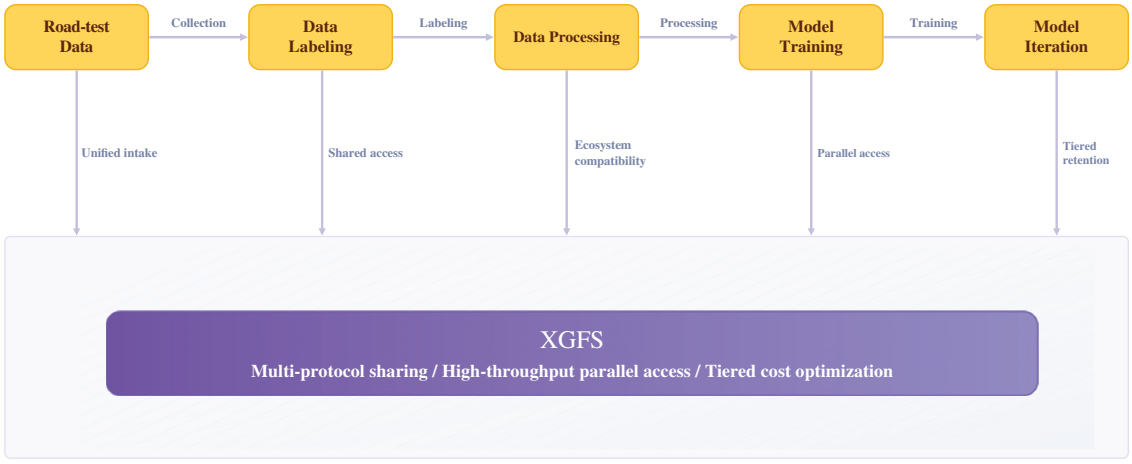

From an AI perspective, managing and governing large-scale unstructured data for downstream training and inference requires more than generic storage capabilities. It demands support for large-scale metadata management and searchability, policy-based lifecycle and tiering across hot, warm and cold layers, and secure multi-tenant controls and auditability for enterprise-grade AI data governance, which our data lake storage solutions are designed to meet. Also, by providing a unified namespace and flexible multi-protocol access, our solutions facilitate data organization, tiering and movement across storage layers and environments, reducing data silos and

operational complexity and enabling enterprises to accumulate proprietary data assets that support continuous AI optimization. The following table illustrates the AI-related features and functions of our products in comparison with traditional storage solutions that are not specifically designed or optimized for AI workloads:

Product	AI-Related Features	How This Is Different from Traditional Data Storage
XGFS/MeshFS	Uses very fast SSD caching, a design that avoids system lock-ups, and direct high-speed data transfer between servers, enabling high concurrency of thousands of GPU read-write actions which are common and necessary in AI training and inferences	Ordinary file systems cannot handle thousands of GPUs reading and writing data at the same time.
MeshFusion	Stores frequently-used AI inference data so servers can access it quickly, reducing the need for expensive GPU memory. This is specifically built for AI LLM inference.	Traditional storage does not have such needs.
XEOS/MeshSpace	Can store extremely large amounts of data in one place, across regions, and automatically move data between hot (frequently-used) and cold (non-frequently used) storage. This is designed to support massive data collection and cleaning before AI training.	Traditional storage is not built for this scale.
EasyData	Helps users label data, change data formats, and prepare data when instructed by the customer.	Supports AI data preparation, but not specifically for AI purposes.
XEBS-XINFINI	Very fast shared block storage that allows many systems to access data in parallel with low latency and delay, providing fast access for AI vector search and high-speed workloads.	General block storage is slower and allows much lower level of parallel access.

Our Turnkey Appliances

Leveraging our extensive experience in distributed storage, we have deeply integrated compute, networking and multi-protocol storage capabilities to develop a portfolio of turnkey appliances. Our offerings include the X3000/X5000 turnkey appliances as well as the Y3000/Y5000 turnkey appliances based on domestic processors, which provide pre-configured, out-of-the box deployment options for customers seeking scalable data infrastructure. Each of these appliances supports all of our data lake storage solutions as well as our storage solutions for T&I.

The 3000-series appliances are optimized for high performance workloads. The 5000-series appliances are designed for capacity-centric scenarios. The X series is our standard non-domestic line intended for mainstream enterprise environments, while the Y series is our domestic-ecosystem version, designed for full compatibility with domestic CPUs, operating systems and trust-stack components.

The following table sets out details of our distributed storage appliance models:

Model	Supported Storage Types	Capacity	Reliability	Application Scenarios
X3000/X5000 Turnkey Appliances **(X3000)**	Block, File, Object	X3000: HDD 20~240TB; SSD 9~153TB X5000: HDD 20~720TB Dual CPU; up to 512 GB memory Up to EB-level capacity, 4,096 nodes	Comprehensive data protection strategies implemented through our software architecture, complemented by disaster recovery capabilities	• Cloud computing • Virtualization • Virtual desktop • Database • Intelligent manufacturing • Content and media services • Cloud-native applications • Financial imaging systems • Medical picture archiving and communication system (PACS)
Y3000/Y5000 Turnkey Appliances based on Domestic Processors **(Y5000)**	Block, File, Object	Y3000: HDD 20~240TB; SSD 9~92TB Y5000: HDD 20~720TB Dual CPU; up to 512 GB memory; Up to EB-level capacity, 4,096 nodes	Comprehensive data protection and reliability design ensure data safety, backed by five-nines-plus data durability	• Content and media services • Cloud-native applications • Financial imaging systems • PACS

AI Storage Services

AI storage services comprise tail-end warranty, maintenance or renewal services that accompany our AI storage solutions. Services typically include software updates and patches, proactive health checks, performance tuning for AI T&I workloads, capacity and lifecycle planning, security enhancement and component replacement coordination under tiered Service Level Agreements (SLAs). We also offer enablement services that are designed to maximize the value of our AI storage solutions, aiming to help customers embed our storage solutions into their existing IT and AI workflows. AI storage services are the necessary continuation of AI storage solutions following deployment. They cover ongoing operational and technical support essential to AI environments, including performance tuning for AI training and inference workloads, capacity planning, and security enhancement. These services ensure that AI-oriented storage systems continue to operate in line with the demanding performance and reliability requirements of AI workloads. They are an indispensable component of AI-enabling infrastructure, ensuring that underlying storage platforms supporting AI workloads remain stable, functional, and continuously available throughout the AI lifecycle.

For the purpose of classifying our revenue under Chapter 18C, we consider sales of our AI storage solution software and turnkey appliances to be revenue from our Specialist Technology Products. See "— Business Model" for a detailed discussion of our AI storage solutions delivered as software only and turnkey appliances. We also consider sales of our AI storage services to be revenue from our Specialist Technology Products. The following table sets forth a revenue breakdown by product mix in absolute amount and as percentage of our total revenue for the years/periods indicated:

	Year ended December 31,						Six months ended June 30,			
	2023		2024		2025		2025		2026	
	RMB'000	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*
							(Unaudited)			
Software-based solutions	56,793	34.1	70,022	40.6	114,806	43.0	46,871	47.1	62,701	33.7
Turnkey appliance	80,961	48.5	68,028	39.4	114,125	42.7	34,525	34.7	102,666	55.3
AI storage services	29,087	17.4	34,434	20.0	38,082	14.3	18,133	18.2	20,393	11.0
Total	**166,841**	**100.0**	**172,484**	**100.0**	**267,013**	**100.0**	**99,529**	**100.0**	**185,760**	**100.0**

Our Directors are of the view that, based on the technology characteristics and use cases described above, each of the foregoing solutions and services falls within an acceptable sector of a Specialist Technology Industry under Chapter 18C of the Listing Rules. Based on the above analysis and the view of the Directors and CIC, the Sole Sponsor is of the view that each of our solutions and services qualify within an acceptable sector of Specialist Technology Industry as defined under Chapter 18C of the Listing Rules.

Commercialization

We adopt a transaction-based model for our offerings. The following chart illustrates the commercialization timeline of our offerings, reflecting our continuous commercial application of technologies:

Specialist Technology Products	Launch	Start of Revenue Generation
Data lake storage solutions	July 2022	September 2022
Storage solutions for T&I	July 2021	September 2021
AI storage services	July 2021	September 2021

Our commercialization centers around increasing engagement with customers across key industries, expanding our reach among enterprise customers. By serving customers across different verticals, we gain deeper insights into sector-specific challenges, making our solutions more attractive to new customers. This enables efficient scaling and drives customer acquisition while boosting our brand influence and market penetration.

The number of customers remained stable at 251 in 2023, 257 in 2024, 269 in 2025, 223 in the six months ended June 30, 2025 and 248 in the six months ended June 30, 2026. Cumulatively, we served 490, 535, 603, 559 and 635 customers as of December 31, 2023, 2024, 2025 and June 30, 2025 and 2026, respectively. Our revenue increased from RMB166.8 million in 2023 to RMB172.5 million in 2024. Our revenue further increased by 54.8% from RMB172.5 million in 2024 to RMB267.0 million in 2025. Our revenue increased from RMB99.5 million in the six months ended June 30, 2025 to RMB185.8 million in the six months ended June 30, 2026.

The following table sets forth our revenue breakdown in absolute amounts and as percentages of our total revenue for the years/periods indicated:

	Year ended December 31,						Six months ended June 30,			
	2023		2024		2025		2025		2026	
	RMB'000	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*
							(Unaudited)			
AI Storage Solutions										
Data lake storage solutions	62,060	37.2	70,141	40.7	129,696	48.6	39,733	39.9	102,017	54.9
Storage solutions for T&I	75,694	45.4	67,909	39.3	99,235	37.2	41,663	41.9	63,350	34.1
AI Storage Services	29,087	17.4	34,434	20.0	38,082	14.2	18,133	18.2	20,393	11.0
Total	**166,841**	**100.0**	**172,484**	**100.0**	**267,013**	**100.0**	**99,529**	**100.0**	**185,760**	**100.0**

Strategic Role within AIMesh

AIMesh, which we launched in January 2026, is a significant addition to our solution portfolio and a cornerstone of our long-term strategy for AI infrastructure. AIMesh comprises MeshFS, MeshFusion and MeshSpace, which represent upgraded iterations of our existing software products. These three products can be offered as a package together with our existing products. AIMesh is part of our Specialist Technology Products and is designed as a unifying software-defined layer that spans both our data lake storage solutions and our storage solutions for T&I. In particular, AIMesh is developed based on the same proprietary distributed storage technologies as our existing solution offerings and is designed to address similar AI-related data storage and processing requirements. It represents an extension and upgrade within our existing technological framework rather than an entry into a new business line. AIMesh is a unified AI data and memory fabric — a software-defined layer that connects different types of storage and memory resources — designed to overcome three critical bottlenecks in AI systems:

- **MeshFS** — a high-speed data layer for AI training, engineered to overcome the I/O wall, which occurs when data cannot be delivered fast enough to keep powerful processors fully utilized;
- **MeshFusion** — an inference memory layer that expands memory capacity for large-context AI applications, addressing the memory wall, which limits the ability to handle long conversations or complex reasoning; and
- **MeshSpace** — a global data lake for managing and moving massive datasets, solving the gravity wall, which makes it difficult and costly to transfer and govern large volumes of data.

Together, these components form AIMesh, a single, integrated platform that connects diverse resources — including local solid-state drives, all-flash arrays, high-capacity hard drives, tape libraries, and public-cloud storage — into a consistent, software-defined system that connects the compute system, storage system and business and operations management systems, enabling efficient interaction and unified management across the entire AI infrastructure. Through a global namespace (a unified addressing system) and a central control plane, data written once can move across training, inference, and archival layers without manual migration or application changes. AIMesh represents our next-generation architectural standard for AI infrastructure, building on our proven XSEA shared-everything architecture and XScale metadata engine, both widely deployed in enterprise environments.

To further illustrate the differences between AIMesh and our existing solution offerings, the following table sets forth a comparison based on key functional roles and quantitative performance metrics:

AIMesh Component	Functional Role	Related Existing Product	Key Quantitative Differences
MeshFS	High-performance data layer for AI training (addressing I/O bottlenecks)	XGFS	MeshFS enhances the performance of XGFS by reducing access latency from 100 microseconds to as short as approximately 20 microseconds and enabling throughput at multi-GB/s per node, aggregating to tens of GB/s at the cluster level.
MeshFusion	Memory extension layer for inference (expanding KV cache for large-context AI workloads)	XEBS-XINFINI	MeshFusion improves our current portfolio by enhancing system throughput and efficiency in AI workloads. Test results indicate performance improvements of approximately 10 to 30 times under certain configurations, with throughput generally reaching tens of thousands of tokens per second. It also improves data transfer efficiency and reduces response latency. In addition, it supports terabyte-scale KV cache pools on a per-cluster basis, with cache access latency at the sub-millisecond level.
MeshSpace	Unified global AI data lake (cross-region and cross-cloud data orchestration)	XEOS	MeshSpace improves XEOS by enabling a unified namespace across multiple clusters and cloud environments. In a representative deployment, it supported a single namespace of 100PB usable capacity, achieving read throughput of up to 3 Tbps and write throughput of up to 2.5 Tbps, as well as approximately 4 million 4KB read requests per second. It is designed to support exabyte-scale storage capacity with multi-site deployment capability.

As illustrated by the quantitative metrics set out in the table above, AIMesh not only enhances performance levels but also introduces additional functionalities and architectural capabilities that extend our existing solution offerings, including:

- **Lower training-data latency:** MeshFS improves XGFS by reducing data access latency from around 100 microseconds to about 20 microseconds, achieved by using local SSDs as a cache layer so GPUs receive data faster and spend less time waiting;

- **Higher training throughput:** MeshFS delivers higher data throughput, reaching multi-GB/s per node and tens of GB/s at cluster level, achieved by reading and writing across storage nodes in parallel and transferring data directly between servers without involving the CPU;

- **Larger inference memory at lower cost:** MeshFusion improves our inference capability by extending usable memory beyond local GPU and server memory, achieved by pooling DRAM and SSD resources across servers as an external KV Cache layer that is cheaper than adding more high-bandwidth memory;

- **Faster and more efficient inference:** MeshFusion improves system throughput and efficiency for AI workloads, with test results showing roughly 10 to 30 times performance improvement under certain configurations, achieved through shared-everything parallel I/O, low-latency retrieval design and a zero-copy data path that avoids unnecessary software overhead; and

- **Faster, lower-cost access to large datasets across locations:** MeshSpace improves XEOS by giving AI teams one place to access data spread across clusters, cloud environments and third-party storage, achieved by using a unified global namespace and automatically moving or tiering data between high-speed storage, hybrid tiers and lower-cost storage based on policy and usage.

MeshFS: High-Speed Training Data Layer

MeshFS is the next-generation evolution of XGFS, a storage solution for T&I designed specifically for large-scale AI training workloads. It retains compatibility with standard file system protocols while introducing optimizations for AI workloads:

- **Ultra-Low Latency Cache**: MeshFS uses local SSDs as a Tier-0 cache — the fastest layer of storage — reducing data access time to about 20 microseconds. This prevents expensive graphics processors (GPUs) from sitting idle and accelerates training.

- **Parallel Metadata Access**: AI training often involves millions of small files accessed simultaneously. MeshFS distributes file information (metadata) across all nodes in a lock-free, parallel design, ensuring stable performance even under extreme concurrency.

- **High-Throughput Architecture**: MeshFS enables simultaneous reading and writing across all storage nodes using a technology that transfers data directly between servers without CPU involvement. This eliminates bandwidth bottlenecks and supports multi-gigabyte-per-second data delivery.

- **Flexible Tiering and Data Sharding Algorithms**: MeshFS employs distributed data and metadata sharding algorithms to enable parallel file access and load balancing across multiple nodes, while supporting fault awareness and data rebalancing. In addition, through tiering policies and hot- and cold-data migration, MeshFS improves storage resource efficiency while sustaining stable T&I performance.

MeshFusion: Expanding Memory for AI Inference

We believe that as AI evolves, industry investment will gradually shift from model training to large-scale inference. Inference workloads increasingly involve multi-round conversation, long-context reasoning and the simultaneous execution of many AI agents. These patterns generate rapidly expanding KV Cache data (intermediate memory states used to accelerate inference), which must be retrieved and updated in real time. As models grow and context windows lengthen, the high-bandwidth memory (HBM) within a single inference server will no longer be sufficient to hold all inference-time state.

To address this challenge, we are developing MeshFusion as a cost-effective alternative to HBM. Built on our XSEA architecture, MeshFusion acts as an L3 persistent KV Cache layer (an external memory layer that extends usable memory beyond local GPU and server memory) by aggregating Dynamic Random Access Memory (DRAM) and SSD resources across multiple inference servers into a unified, high-performance pool. This centralized resource enables inference clusters to maintain much larger contexts, reduce memory pressure on individual GPUs, and sustain stable performance even when workloads exceed single-machine capacity.

MeshFusion is powered by three core innovations:

- **Shared-everything parallel I/O:** MeshFusion allows all inference servers to access all storage devices directly, increasing concurrency and ensuring consistent performance across large clusters;

- **End-to-end erasure-coded writes and one-hop reads:** MeshFusion minimize retrieval latency for KV Cache and other inference-critical data; and

- **Zero-copy data path:** MeshFusion minimizes additional processing requirements and delivers near-hardware-level performance.

Key capabilities of MeshFusion include:

- **L3 external memory layer for inference:** MeshFusion creates a unified L3 layer using local SSDs to store persistent KV Cache when the first two layers — GPU high-bandwidth memory (L1) and system DRAM (L2) run out of space. This greatly increases the amount of context that can be handled during inference, at a much lower cost than adding more HBM.

- **Dual data channels:** MeshFusion supports two types of data paths: standard POSIX paths for storing and accessing model files at high speed, and specialized channels for inference frameworks like vLLM and SGLang. These channels allow zero-copy access, meaning data can be used directly without extra duplication, improving efficiency.

- **Integrated deployment model:** Unlike traditional external storage, MeshFusion runs directly on inference servers as containerized services, bringing memory expansion closer to the compute plane, reducing network hops and improving cluster efficiency.

Together, these capabilities form a scalable external memory pool for inference that supports long-context reasoning, improves cluster-wide efficiency and reduces dependence on costly high-density memory.

MeshSpace: Global AI Data Lake

MeshSpace is the next-generation evolution of XEOS, designed for managing exabyte-scale data across regions and platforms. It introduces a global namespace and object-virtualization framework that unifies multiple XEOS clusters — and even third-party storage systems — into a single data fabric. This eliminates data silos and gives organizations a consistent view of data across regions and platforms.

- **Unified Global View Across Regions and Platforms:** MeshSpace creates a virtualized layer over multiple storage systems, such as XEOS clusters in different locations, third-party storage and public-cloud object storage. No matter where data physically resides, AI developers access it through a single, unified address space using the standard protocol. Behind the scenes, MeshSpace automatically decides where data should be stored and moves it between high-speed flash storage, hybrid tiers, and cloud storage based on predefined policies and usage patterns.

- **Global Data Mobility and Distribution:** MeshSpace enables automated, policy-driven movement of datasets across different sites, ensuring that data stays synchronized between central data centers and branch locations. It can also automatically move "cold" data — information that is rarely accessed — from expensive high-performance storage to more cost-efficient options, such as hard disk drive (HDD) pools or cloud storage. This tiering approach helps reduce overall storage costs while maintaining accessibility throughout the AI data lifecycle.

Since its launch in January 2026, AIMesh has secured a customer order and commenced generating revenue as of the Latest Practicable Date.

Key Operating Data

The following table sets forth key metrics of our AI storage business:

	Year ended December 31,						Six months ended June 30,			
	2023		2024		2025		2025		2026	
	Data lake solutions	Storage solutions for T&I	Data lake solutions	Storage solutions for T&I	Data lake solutions	Storage solutions for T&I	Data lake solutions	Storage solutions for T&I	Data lake solutions	Storage solutions for T&I
Number of customers during the year/period[(1)]	79	53	60	56	92	72	46	37	59	48
Number of new customers[(2)]	49	22	22	21	49	27	16	7	19	19
Cumulative customers served	103	98	125	119	174	146	141	126	193	165
Average customer value[(3)] (RMB in million)	0.8	1.4	1.2	1.2	1.4	1.4	0.9	1.1	1.7	1.3
No. of early termination/ cancelation of contracts[(4)]	23	14	16	18	7	9	2	2	4	1
No. of orders during the year/period[(5)]	257	270	364	293	508	384	185	145	205	128
Range of amount per order (RMB in thousand)	2.0~ 16,225.8	1.9~ 14,983.5	1.8~ 4,860.6	4.5~ 4,681.9	2.4~ 9,300.0	3.2~ 4,420.0	5.4~ 4,164.0	8.0~ 4,420.0	4.8~ 15,083.4	9.2~ 9,205.0
Average selling price[(6)] (RMB in thousand)	241.5	280.3	192.7	231.8	255.3	258.4	214.7	287.3	497.6	494.9
Average customer acquisition cost (for all AI storage business)[(7)] (RMB in million)	0.4		0.3		0.3		0.2		0.1	
Overall customer retention rate (for all AI storage business)[(8)]	84.8%		82.9%		78.2%		75.4%		74.1%	
Overall net dollar retention rate (for all AI storage business)[(9)]	46.5%		91.9%		132.5%		117.2%		160.0%	

Notes:

(1) Number of customers during the year/period is defined as those in the specified year/period. As part of our sales are conducted through distributors, such number reflects contractual counterparties instead of the number of end users.

(2) Number of new customer equals to the number of customers in the current year/period who did not contribute revenue in the previous year/period. The number of new customers of our AI data lake solutions decreased from 49 in 2023 to 22 in 2024, primarily because we only commenced generating revenue from AI data lake solutions since September 2022. As a result, most customers procuring the data lake solutions in 2023 were classified as new customers, leading to a relatively higher number of new customers compared to 2024. The number of new customers of our AI data lake solutions subsequently rebounded from 22 in 2024 to 49 in 2025, primarily due to the enhanced market recognition of our AI data lake solutions and the effective execution of customer expansion initiatives, which is consistent with the growth in revenue from our AI data lake solutions in 2025 as compared to 2024.

(3) Average customer value for a given year/period is calculated by dividing revenue in that year/period by the number of customers for the same year/period.

(4) Early termination or cancellation of contracts during the Track Record Period was rare and was primarily attributable to normal commercial adjustments, such as end customers' project timetable delays, budget cycle changes or minor scope adjustments, rather than customer loss.

(5) The number of orders during the year/period is defined as the number of contracts that recognized revenue during the year/period. The number of orders for our data lake solutions and storage solutions for T&I increased from 364 and 293 in 2024 to 508 and 384 in 2025, respectively, primarily attributable to the continued acceleration of the commercialization of these solutions.

(6) Average selling price is calculated by dividing revenue for the year/period by the total number of orders in the same year/period. The average selling price of our data lake solutions and storage solutions for T&I decreased from 2023 to 2024 primarily due to a higher proportion of software-only projects, which generally carry lower contract values. The average selling price subsequently increased in 2025, primarily attributable to a higher proportion of higher-value projects and improved profitability of our AI storage solutions. The average selling price further increased in the six months ended June 30, 2026, primarily attributable to larger project sizes as customers accelerated their investments in AI-related data infrastructure and expanded their data storage clusters, as well as a higher proportion of hardware components delivered as part of our turnkey solutions.

(7) Average customer acquisition cost is calculated by dividing the selling and marketing expenses for the year/period by the number of customers acquired in the same year/period. Our average customer acquisition cost decreased from RMB0.4 million in 2023 to RMB0.3 million in 2024, and remained relatively stable at RMB0.3 million in 2025. Our average customer acquisition cost decreased from RMB0.2 million in the six months ended June 30, 2025 to RMB0.1 million in the six months ended June 30, 2026.

(8) Overall customer retention rate is calculated by subtracting the number of new customers acquired during the year/period from the total number of customers at the end of the year/period, and dividing the result by the total number of customers at the beginning of the year/period (i.e., the total number of customers at the end of the previous year/period) and multiplied by 100%. The decline in our overall customer retention rate was mainly due to the project-based nature of customer procurement, under which customers may not place orders in every year/period, as well as the expansion of our customer base, and did not reflect any material deterioration in our customer relationships.

(9) Overall net dollar retention rate equals the revenue of a current year/period from customers that contributed to our revenue for both the current and previous years/periods divided by the revenue of the previous year/period and multiplied by 100%. Our overall net dollar retention rate increased from 46.5% in 2023 to 91.9% in 2024 and further to 132.5% in 2025, and from 117.2% in the six months ended June 30, 2025 to 160.0% in the same period in 2026, primarily attributable to the increased order volume and expanded deployment scale from existing customers, supported by high customer stickiness arising from the infrastructure nature of our AI storage solutions, under which customers' core data is deeply embedded on our platform.

The following table sets forth the number of orders and number of end customers, including those under both direct sales and the distributorship model, by nature of business of the customers during the Track Record Period:

	Year ended December 31,						Six months ended June 30,			
	2023		2024		2025		2025		2026	
	Number of orders	Number of end customers	Number of orders	Number of end customers	Number of orders	Number of end customers	Number of orders	Number of end customers	Number of orders	Number of end customers
Technology Enterprises	652	226	831	243	906	269	714	219	753	248
Financial Institutions	391	149	458	170	534	194	394	168	453	184
Intelligent Manufacturing	183	119	291	161	450	216	310	152	457	196
Healthcare	159	128	183	139	208	142	164	123	163	117
Education	112	76	122	77	139	92	100	69	128	85
Government	149	115	143	109	131	97	103	83	100	78
Telecommunications Operators	29	14	34	14	35	19	24	13	29	19
Others[(1)]	31	20	41	25	45	27	34	10	41	10
Total	**1,706**	**847**	**2,103**	**938**	**2,448**	**1,056**	**1,843**	**837**	**2,124**	**937**

Note:

(1) Primarily including power grid, petrochemical and oil industries.

OUR BUSINESS MODEL

Customers of our AI storage solutions may choose to purchase software only or turnkey appliances with our software integrated. In substance, customers procure a complete, pre-integrated system designed to solve storage problems, rather than separable hardware and software components. Software and hardware are synergistic and complementary within the overall solution, and turnkey appliances are priced and contracted as a single deliverable, without itemized allocation between hardware and software. Customers generally derive no meaningful benefit from acquiring hardware or software alone.

Where customers request software-only deployment, we require hardware procurement strictly in accordance with our prescribed bill of materials and may decline delivery if specifications are not met. Our staff deploy the software on customer-procured hardware, resulting in a functional outcome equivalent to purchasing turnkey appliances. Although hardware costs can be quantified separately for accounting purposes, these commoditized components are not the primary source of differentiation or value. Our core value lies in our proprietary software stack, system architecture, and integrated design, as well as our assumption of end-to-end responsibility for system design, hardware selection, procurement, deployment, risk management and spare-parts support. Consistent with the substance of these arrangements and applicable accounting standards, revenue from turnkey appliances is recognized on a gross basis.

Pricing for our AI storage solutions is determined on a case-by-case basis and reflects the overall value delivered, market conditions and customer-specific negotiations, rather than a sum of hardware and software prices. Through hardware-software co-design, our appliances can deliver required performance and reliability more efficiently, in some cases achieving comparable or superior outcomes with fewer hardware resources. Because components are jointly optimized, performance and functionality depend on their combined operation.

The duration of our projects, excluding after-sales maintenance, typically ranges from approximately three to seven months, with an average duration of approximately four to five months. Our typical sales cycle generally consists of the following stages: (i) project requirement solicitation (one to three months); (ii) solution design and proof-of-concept, where applicable (two to four weeks); (iii) project bidding and contract signing (two to four weeks); (iv) production and assembly, where applicable for integrated appliance deployments (one to three weeks); (v) deployment and implementation (one week); and (vi) training and go-live support (one week). Our standard contracts generally include a three-year warranty period and customers may purchase extended maintenance services after the warranty period. The duration for providing AI storage services is typically less than seven days per service request, depending on the specific scope and customer requirements.

Our standard services follow a four-stage process. First, we conduct requirements discovery and, where appropriate, pre-sale proof-of-concept testing, followed by solution design and implementation planning. Second, we perform software installation, configuration, set up automated tiering policies under which data is allocated across appropriate storage tiers based on usage patterns, performance requirements and cost-efficiency considerations, (i.e. setting rules that automatically place frequently used data on higher-performance storage and less frequently used data on lower-cost storage) and carry out data migration to integrate with existing IT environments. Third, we provide role-based training and post-deployment support. Fourth, we deliver tiered technical support, operations and maintenance services, including periodic health checks and optimization recommendations.

Our AI storage solutions support block, file and object storage, work with industry-standard protocols and diverse hardware platforms, and are compatible with major public and private cloud environments. As a data-lake foundation, they support multi-source data ingestion and migration across on-premises and cloud environments with integrity assurance and without business disruption.

While our solutions are designed for on-premises deployment, they support customer-initiated connectivity with public and private clouds to enable hybrid-cloud data management, such as (i) automated data lifecycle tiering (e.g. archiving cold data to lower-cost cloud object storage), (ii)

off-site backup and disaster recovery, (iii) unified access and management across multiple private cloud environments, and (iv) temporary access by cloud compute instances during peak computing demand. Such interoperability does not change the on-premises deployment nature of our solutions, and core data and AI T&I workloads remain on customers' on-premises infrastructure.

Our customer base demonstrates high retention, supported by stable and secure lifecycle data access. In addition to solution sales, we provide AI storage services, primarily warranty, maintenance and renewal services associated with deployed solutions. These services include software updates, health checks, performance tuning, capacity planning, security hardening and component replacement coordination under tiered service-level agreements, ensuring long-term stability and performance continuity.

OUR PROPRIETARY TECHNOLOGIES

Our core technology stack is purpose-built for AI-era data infrastructure and forms the foundation of our differentiated solutions capabilities. It primarily comprises:

XSEA: Distributed Storage Architecture for All-Flash Data Centers

XSEA is a proprietary distributed storage architecture independently designed and developed by our Group for all-flash data centers. It is designed to be fast, resilient and efficient at scale, especially well suited for AI applications. Unlike traditional architectures where each node owns a fixed subset of disks, XSEA adopts a "shared-everything" distributed metadata design, under which every server node in the cluster can directly access all storage devices, rather than being limited to its own disks as in traditional systems. This core architectural concept was developed in-house to address performance bottlenecks, scalability constraints and capacity inefficiencies inherent in conventional storage systems. This design is capable of achieving end-to-end read/write latency of approximately 100 microseconds, depending on network speed, workload level and deployment configuration. The architecture is also designed to reduce data access delays and support service continuity in the event of individual disk or server failures. As of the Latest Practicable Date, a single logical volume supported by XSEA can reach up to approximately 4.1 million IOPS, enabling fast, parallel access essential for AI workloads, vector retrieval and high-speed transactions.

Specifically, XSEA achieves three core advantages: (i) approximately 100 microseconds latency through optimized data paths and minimal processing overhead, (ii) approximately 100 milliseconds recovery through rapid detection and recovery from hardware and system failures, and (iii) approximately 100% effective capacity utilization, enabled by technologies designed to improve storage efficiency, which reduce storage costs without compromising performance.

XSEA is implemented through a layered architecture developed by our Group, comprising (i) a persistent layer that supports ultra-low latency writes and centralized metadata management, (ii) a service layer that provides block and file services with enterprise-grade data protection features, and (iii) a protocol access layer supporting multiple access protocols. These layers work together to form a flat, fully shared resource pool across storage, memory and network resources.

As part of the development of XSEA, we have accumulated significant proprietary know-how and intellectual property. As of June 30, 2026, we had obtained 103 registered patents, several of which directly support XSEA, including patents relating to data processing methods and metadata indexing.

XScale: Core Storage Engine for AI and Multi-cloud Era

XScale is designed to manage datasets of up to approximately 100 PB within a unified system in a flexible and scalable way. It is our core storage engine for the AI and multi cloud era and was also independently designed and developed by us. The principal technical breakthrough of XScale is the decoupling of metadata services from physical data storage, enabling independent and

scalable expansion: namely, XScale separates the system that tracks and manages data (metadata) from the physical storage of the data itself, so each part can grow independently as needs increase. By supporting multiple standard access protocols and using efficient data protection techniques, XScale allows large scale storage systems to expand easily, operate efficiently, and handle complex workloads without redesigning the underlying infrastructure.

XScale adopts a three-layer architecture consisting of (i) access layer, which serves as the entry point for users and applications and supports a wide range of protocols, while integrating load balancing, high availability and small-file handling mechanisms; (ii) catalog layer, which implements distributed metadata services to provide a unified namespace, cross-protocol and cross-site consistency, indexing and lifecycle management; and (iii) data foundation layer, which is responsible for secure data storage through erasure coding and replication, together with tiered storage, automated hot- and cold-data migration and online scalability.

By supporting multiple standard access protocols and applying efficient data protection and data management techniques, XScale enables storage systems to operate reliably under complex, high-concurrency workloads while maintaining scalability and cost efficiency. As the core engine of our unstructured data management capabilities, XScale supports use cases including AI-driven data lakes, multi-cloud collaborative storage and large-scale data lifecycle management.

As part of the development of XScale, we have accumulated proprietary technologies and intellectual property. Supporting patents include those relating to object bucket snapshots, object processing, distributed QoS control and data storage methods.

The following table sets forth how we apply our proprietary technologies across different products to address distinct storage scenarios.

Proprietary Technology	Products/Services Where Applied	Key Features and Outcome
XSEA	XEBS, XEBS-XINFINI, MeshFS, MeshFusion	• *Fast*: Fast data storage by allowing very low latency data access; • *Resilient*: Instantly switching to backup systems in case of system failure; • *Efficient*: Maximize efficiency of storage space utilization, and improved ability to concurrently handle large number of data requests.
XScale	XEOS, XEDP, MeshSpace	Allows the system to: • store and manage extremely large amounts of data; • scale up easily as more data is added; • handle many small files efficiently, and automatically move data between high-cost and low-cost storage to reduce overall costs.

Each product therefore differs in its technical focus and functional outcomes depending on which proprietary technology is applied, enabling us to customize solutions for different storage use cases.

XSEA and XScale technologies deliver faster, more resilient and more scalable storage by significantly improving performance and efficiency across the entire system. Data can be accessed with much lower latency, allowing applications to respond quickly without waiting, while system

failures are addressed through near-instant recovery, minimizing disruption. Storage resources are used far more efficiently, reducing wasted capacity and enabling customers to obtain more usable storage from the same infrastructure. Capacity and performance can be expanded independently, avoiding unnecessary over-investment as needs grow. The distributed design eliminates single control points that could slow down the system, ensures small files are processed efficiently rather than becoming a bottleneck, and allows the overall platform to scale to much larger data volumes without requiring fundamental redesign or replacement. The following table sets forth the comparison analysis between our technologies and the traditional methods.

Comparison Dimension	Traditional Storage Methods	Our Proprietary Technologies (XSEA/XScale)	Difference & AI Relevance
Architecture Design	Data is tied to specific servers, so when many users access the same data, some servers become overloaded.	All storage is pooled and shared, allowing any server to access data directly.	AI training may involve thousands of GPUs reading the same data, and a shared design avoids traffic jams and scales better.
Access Latency	Typical high-end systems show around 1–2 milliseconds worst-case read delay.	Worst-case read delay is reduced to around 100 microseconds, about 10 times faster.	AI training is limited by worst-case delays, and lower tail latency directly improves GPU utilization.
Recovery Time	When a failure happens, training jobs are often interrupted and must restart.	Recovery completes in about 100 milliseconds without interrupting training.	In AI, avoiding job restarts matters more than recovery speed, saving minutes or hours of GPU time.
Storage Capacity Efficiency	Data is stored three times for safety, leaving only about 33% of the hardware usable.	Erasure coding and compression deliver about 2–2.5 times higher effective storage efficiency.	AI datasets are measured in tens of PBs, so higher efficiency materially reduces hardware cost.
Scalability	Computing and storage scale together, and expansion can slow the system.	Compute and storage scale independently, and aggregate bandwidth grows roughly linearly to thousands of nodes.	In AI infrastructure, GPU capacity often grows faster than dataset size, so independent scaling avoids wasted capital spending.
Metadata Processing	Centralized metadata becomes a bottleneck when directories reach tens of millions of files.	Metadata is fully distributed, and performance increases as more nodes are added.	Large AI datasets commonly contain 10^7–10^9 files, making metadata speed critical at job startup.
Small-File Processing	Small files are handled individually, creating high overhead.	Small files are automatically combined and processed more efficiently.	AI datasets are dominated by small files, so aggregation significantly improves training efficiency.
Maximum Data Scale	Datasets above the PB level must be split across multiple systems.	Datasets of 100 PB can be managed within one unified system.	AI users benefit most from keeping all training data in one system rather than splitting across systems.

RESEARCH AND DEVELOPMENT

Our ability to originate core storage technologies, engineer new products and continuously enhance existing systems is critical to sustaining our market position in AI data infrastructure. We focus R&D on software-defined, protocol-agnostic data platforms that scale from hundreds of billions of objects to EB-class pools and deliver microsecond-class I/O for T&I. All of our key Specialist Technology Products and their underlying core technologies have been independently developed in-house by our Group. We own all intellectual property rights arising from our R&D activities and products. We do not rely on any material third-party licensed technologies, nor did we outsource or engage in any material R&D collaborations during the Track Record Period.

R&D Team

Our R&D organization consists of professionals with experience in technology companies and research institutions. They specialize in building high-performance systems that help businesses store, manage and use data more effectively, especially in areas like software engineering, storage development and enterprise system architecture.

Each of our core R&D team members has their specialized areas and the following table sets out their profile:

Core R&D team member	Profile
WANG Haomai (王豪邁)	Mr. Wang obtained a bachelor's degree in information management and information system from Southwest Jiaotong university and has over ten years of experience in storage development. He is primarily responsible for technology R&D, product architecture decisions and core technological innovation.
LU Jiafu (魯加福)	Mr. Lu obtained a bachelor's degree in mechanical design and manufacturing from Beihang University and has over 25 years of experience in software engineering and enterprise system architecture. He is currently the general manager of our storage business unit.
QU Tianshan (瞿天善)	Mr. Qu obtained a master's degree in computer technology from Nanjing University and has around nine years of experience in storage development. He is currently serving as an architect at our Company.
CHI Xinze (池信澤)	Mr. Chi obtained a master's degree in computer science and technology from the University of Science and Technology of China and has over ten years of experience in software engineering. He is currently an architect and manager of our platform R&D department.
LI Erlang (李二郎)	Mr. Li obtained a master's degree in pattern recognition and intelligent systems from Huazhong University of Science and Technology and has over ten years of experience in storage development. He is currently an architect and manager of our platform R&D department.
LIU Mingxin (劉名欣)	Mr. Liu obtained a bachelor's degree in computer science and technology from Central South University and has over ten years of experience in storage development. He is currently an architect and manager of our platform R&D department.
YANG Ruifeng (楊瑞峰)	Mr. Yang obtained a bachelor's degree in computer science and technology from Inner Mongolia University of Science and Technology and has over ten years of experience in enterprise system architecture. He is currently serving as an architect at our Company.

Our R&D focus has evolved through different stages alongside the maturation of our technology platform. During the initial phase, R&D resources were primarily devoted to building core distributed storage architectures and foundational engines at an early stage of development. As these core technologies matured, we entered a second phase focused on productization and expansion into AI-related storage solutions through the reuse of established engines and modules. In the current phase, R&D is centered on platform-level integration and AI-native capabilities, with development efforts concentrated on priority initiatives built on the existing technology base rather than on new foundational architectures.

Our R&D team consisted of 156, 133, 129 and 132 members as of December 31, 2023, 2024, 2025 and June 30, 2026, respectively, representing 43.5%, 42.1%, 41.1% and 40.2% of our total employees during the same years/periods. All of our R&D activities were conducted in-house during the Track Record Period. We incurred research and development expenses of RMB110.0 million, RMB88.0 million, RMB77.2 million, RMB38.0 million and RMB41.4 million in 2023, 2024, 2025 and the six months ended June 30, 2025 and 2026, respectively, representing 65.9%,

51.0%, 28.9%, 38.2% and 22.3% of our total revenue for the respective years/periods. The continuous decline in R&D expenses as a percentage of revenue during the Track Record Period was primarily attributable to (i) the maturation of our core underlying technologies after over a decade of continuous development, resulting in a shift in R&D emphasis toward platform iteration and productization; (ii) the implementation of a unified and modularized R&D process and architecture across business lines, enabling extensive reuse of existing core components and reducing incremental R&D resources required for new solution development; and (iii) the adoption of AI-enabled quality engineering systems and productivity tools that automate parts of coding, testing and documentation workflows, thereby improving output per engineer. During the same period, R&D remained a key priority within our operating expenditure, with R&D expenses accounting for 37.9%, 42.6%, 43.9%, 44.6% and 45.2% of total operating expenses in 2023, 2024, 2025 and the six months ended June 30, 2025 and 2026, respectively. Specifically, we intend to build a talent pool to support our R&D initiatives for next-generation AI-native storage platforms and related vertical expansion and adaptation. This will be achieved by recruiting experienced professionals across distributed systems architecture, storage kernel development and testing, with approximately 20-25 personnel engaged in the development of high-performance distributed storage engines and AI storage products, and approximately 15-20 personnel focused on compatibility certification and ecosystem community development. See "Future Plans and Use of [**REDACTED**]" for details of our future R&D investment. Our R&D roadmap includes, among others, iterative releases and enhancements of AIMesh, the continued enhancement of MeshFS to support broader commercial deployment, and the ongoing upgrade of XSEA and XScale to support large-model, inference-intensive and multimodal workloads.

We retain key technical personnel through competitive compensation and benefits and defined technical career paths. To improve our R&D capability at the group level, we have established employee shareholding platform programs for our R&D employees. We invest in continuing education and sponsor industry conferences and standards engagement to keep our engineers current with best practices.

To mitigate personnel change risks and protect proprietary know-how, we have implemented: (i) non-compete and IP ownership agreements at onboarding to ensure patents, copyrights and proprietary code vest with our Company; (ii) centralized documentation and code repositories under tiered access controls; (iii) comprehensive resignation procedures with mandatory handover of artifacts, environments and client interfaces; and (iv) post-employment confidentiality undertakings and, where appropriate, enforcement of non-compete obligations on a case-by-case basis.

The salient terms of agreements with management and technical staff are set out below:

- ***Ownership of intellectual properties***. We hold all intellectual property rights, including patents, utility models, design rights, copyrights and related interests, in any proprietary technology, patented products or other works created or contributed to by an employee in the course of their employment duties.

- ***No conflict***. Our staff are required to confirm that their employment with us does not violate any confidentiality or non-compete obligations owed to their previous employers or third parties.

- ***Non-competition***. We have the right to unilaterally initiate a non-competition period of up to two years following the termination of employment. During the term of employment and the non-competition period initiated by us, employees shall not engage in any competitive behavior specified in the agreements. We pay monthly compensation during the post-employment non-compete period.

- ***Confidentiality***. During and after the term of employment, our staff could use our confidential information only for legitimate business purposes and must not disclose it to anyone — including other employees — unless required for their duties. This

obligation continues indefinitely until the information lawfully becomes public, and upon departure, our staff must promptly return or delete all materials (including from personal devices), transfer device ownership if needed (with fair compensation), and provide written confirmation of compliance.

During the Track Record Period and up to the Latest Practicable Date, there was no legal claim or proceeding that may have an influence on the R&D of our Specialist Technology Products.

Key Research Projects

The principal research themes and current project tracks are set out below.

Next-Generation Parallel File System

Our next-generation parallel file system is a core research and development initiative based on our over ten years of accumulated expertise in proprietary distributed storage technologies. It constitutes a generational upgrade within our existing technology framework. It is designed to deliver ultra-high performance data access and unified multi-protocol support, and serves as the core technical foundation for our MeshFS and MeshFusion products. It is intended to address the performance and scalability requirements of large-scale AI T&I workloads by providing a high-throughput, low-latency and highly concurrent data access layer.

To clearly distinguish our next-generation parallel file system from our existing technologies, we set forth below a comparison across functionality, performance and intended use cases:

- **Functionality**

  Our existing technologies (as represented by the engine underlying XGFS) primarily provide general-purpose distributed file services. In contrast, our next-generation parallel file system adopts a parallel file system architecture with unified multi-protocol support, enabling more efficient coordination of data access across distributed computing environments.

- **Performance**

  Our existing technologies support aggregated throughput at the level of tens of GB/s, with latency typically ranging from sub-millisecond to millisecond levels. Our next-generation parallel file system is designed to support significantly higher performance, with target aggregated throughput at the level of hundreds of GB/s per cluster, and to enable cache access latency as low as approximately 20 microseconds, depending on deployment configuration and workload conditions.

- **Intended use cases**

  Our existing technologies are primarily deployed in general enterprise scenarios, such as file sharing and big data processing. In contrast, our next-generation parallel file system is designed for large-scale T&I workloads and serves as a core underlying technology for AIMesh components such as MeshFS and MeshFusion.

The key development milestones of our next-generation parallel file system are as follows:

- In January 2026, we officially launched the MeshFS, which incorporates the architecture and capabilities developed through our next-generation parallel file system initiative, marking the successful validation of its architecture and performance objectives, and the beginning of its commercial availability.

- In June 2026, we launched MeshFS V7.3 for AI training storage, which incorporates the capabilities of our next-generation parallel file system and delivers enhanced production-grade AI file storage capabilities, with file performance of up to 4 million IOPS.

- In the fourth quarter of 2026, we plan to further upgrade our next-generation parallel file system with coordinated operation and tiering capabilities between MeshFS and MeshSpace, enabling efficient data movement across training, inference and data lake layers.

AIMesh

AIMesh is part of our Specialist Technology Products and is designed as a unifying software-defined layer that spans both our data lake storage solutions and our storage solutions for T&I. AIMesh is a unified AI data and memory fabric including the following portfolios: (i) MeshFS, a high-speed data layer for AI training that is engineered to overcome the I/O wall, (ii) MeshFusion, an inference memory layer that expands memory capacity for large-context AI applications, and (iii) MeshSpace, a global data lake for managing and moving massive datasets. MeshSpace is our next-generation object storage capability defined for the XEOS product line. It is designed to support single-bucket, exabyte-scale data capacity and to provide unified access and a single namespace for geographically distributed, cross-cloud and heterogeneous unstructured storage environments. Together, these portfolios form AIMesh, a platform that connects diverse resources into a consistent, software-defined system.

The key development milestones of AIMesh are as follows:

- In the fourth quarter of 2025, we completed the technical architecture prototype and validation, confirming feasibility of the global namespace and large-scale metadata design.

- In January 2026, we officially launched the MeshSpace, MeshFS and MeshFusion.

- In the fourth quarter of 2026, we plan to complete the development of upgraded AIMesh version, delivering full global namespace management capabilities across multiple regions and storage platforms.

R&D Process

Our in-house R&D process for software-defined data infrastructure runs from demand intake to release and post-release support, and is governed by company-wide engineering policies covering design reviews, change control, and security and compliance baselines.

We collect requests mainly from customer needs managed by our product management office (PMO) and from platform improvements planned jointly by engineering and the PMO. Each request is classified as ordinary or urgent: ordinary items follow a gated cycle, while urgent items take an expedited path with risk sign-offs. The PMO assesses each request for replicability and business value, and customized, non-replicable items proceed only after reconfirmation of necessity and scope.

- **Project Initiation.** We define the problem based on actual workloads and use cases, assess feasibility and cost-benefit, and set success criteria (including response time targets, data reliability, expected effective capacity and, where applicable, WORM and legal hold). A formal scoping review confirms scalability, correctness with relevant protocols and hardware, and fit with multi-tenant and access-control design, resources and budget. Only projects that pass this review are included in our product roadmap.

- **Project Planning.** Once approved, a cross-functional team is formed to prepare a comprehensive project plan detailing objectives, key design decisions, milestones, and acceptance criteria. The plan also defines a robust testing strategy spanning diverse

hardware configurations, operating systems, network protocols, and integration environments. Security and regulatory compliance requirements are identified at the outset, and development changes are scheduled in phases, using controlled rollouts and feature flags to mitigate implementation risk.

- **Product Development.** Engineering proceeds in short iterative cycles with modular components. Code changes undergo peer review, static checks and unit testing before merging through continuous integration pipelines, alongside performance testing to meet latency and throughput targets as data volumes and concurrency grow. Security controls such as encryption in transit and at rest, secure credential management, and role-based access control are built in, and documentation is maintained.

- **Product Validation.** We run protocol interoperability and system tests, failure and durability tests, and scale/soak testing on large datasets, and benchmark response time, throughput and checkpoint/clone behavior under realistic concurrency. Security validation covers hardening, access control and audit logs and, where appropriate, penetration testing; release candidates proceed only if predefined performance, reliability, security and compatibility standards are met.

- **Product Release.** We typically issue two major versions each year and interim releases for bug fixes and compatibility updates, with expedited fixes for critical defects subject to standard testing and approvals. Releases are rolled out in stages from internal labs to pilot customers and then deployment, with post-release monitoring, feedback and corrective updates, and lessons learned incorporated into internal standards.

PRODUCTION, DELIVERY AND INVENTORY MANAGEMENT

Hardware Production

Our self-operated manufacturing facility is located in Beijing, China. The hardware manufactured by us during the Track Record Period primarily comprised turnkey appliances. The key raw materials of our turnkey appliances include SSDs, HDDs, memory modules, CPUs and network interface cards (NICs). All of the above key raw materials and components are procured from third-party suppliers and are not manufactured in-house by our Group. We are responsible for the design and specification of our hardware products, while the assembly and integration are carried out through a combination of (i) procurement of customized server units from OEM suppliers and (ii) in-house assembly, system integration, software installation and testing. We divide our production process into two main stages, namely the material preparation and the assembly, inspection and packaging stage, which allows for flexible scheduling of production work orders based on customer requirements and production capacity.

The following table sets forth the details of the designed production capacity and utilization rates for our manufacturing facilities for the years/periods indicated:

	As of/For the Year Ended December 31,									As of/For the six months ended June 30,		
	2023			2024			2025			2026		
	Designed Production capacity[1]	Actual production volume[2]	Utilization rate (%)[3]	Designed Production capacity[1]	Actual production volume[2]	Utilization rate (%)[3]	Designed Production capacity[1]	Actual production volume[2]	Utilization rate (%)[3]	Designed Production capacity[1]	Actual production volume[2]	Utilization rate (%)[3]
Shenzhen plant[4]	500	429	85.8	500	351	70.2	/	/	/	/	/	/
Beijing plant	500	453	90.6	500	428	85.6	1000	982	98.2	500	494	98.8

Notes:

(1) Our production capacity is calculated based on a single-shift arrangement and our normal production schedule during the relevant year/period, taking into account the time required for assembly, inspection and packaging for each unit, as well as the parallel testing capacity allowed by safety power-load conditions and available manpower.

(2) The actual production refers to actual output for the relevant year/period.

(3) The utilization rate during the period is calculated by dividing the actual production by the designed capacity for the same year/period.

(4) The Shenzhen plant was closed in August 2024 for management efficiency purposes.

Our production process is designed to uphold high standards of quality while maintaining the agility to respond swiftly to customer demands. Leveraging our mass production capabilities and stringent quality control measures, we ensure the performance, reliability and consistency of our solutions. The process begins with the generation of production tasks and material requisition, followed by multiple verification steps to confirm consistency between material quantities, configuration lists and contractual specifications. During assembly, we integrate components and accessories. Before final dispatch, a comprehensive outbound inspection is conducted, covering appearance, quantity, labeling, packaging and accessories.

We are committed to the continual development of production process techniques to enhance our production capabilities and accelerate the automation and digitalization of our production lines. The diagram below illustrates the key steps of our production process.

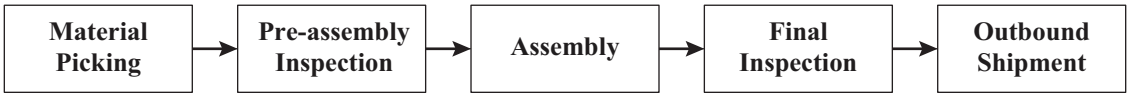

Solution Delivery

We have established a standardized customer service process that spans the entire lifecycle of solution delivery. From initial design to ongoing operations, this model ensures smooth deployment, rapid adoption and sustained value realization for our customers, which strengthens project success rates and accelerates ROI. Our standard service framework covers the following four stages:

- **Solution design and validation.** We aim to align precisely with our customers' business objectives. Our expert team conducts in-depth requirement analysis to identify pain points of our customers and frequently supports proof of concept testing during the pre-sales phase to confirm functionality and performance. This approach mitigates project risks and ensures solution viability. Ultimately, we collaborate closely with our customers to jointly develop a comprehensive solution and project blueprint, laying a solid foundation for successful delivery.

- **Deployment and system integration.** We endeavor to achieve full integration of our solutions with the existing IT environment and protocols such as ERP and CRM systems adopted by our customers. Following a proven deployment methodology that encompasses installation, configuration, data migration and interface development, we maintain transparent communication and provide detailed configuration documentation to ensure our solutions are implemented with precision and care.

- **Training and go-live support.** We work closely with our customers' technical teams to ensure full proficiency in using our solutions and to achieve a smooth transition that minimizes potential adaptation costs. We deliver structured, role-based training for administrators, key users and operators. During the go-live period, our experts provide intensive support both on-site and remotely, ensuring rapid issue resolution and uninterrupted business continuity during this critical phase.

- **Continuous operation and value optimization.** After our solutions successfully go-live, we focus on building a long-term partnership and helping our customers maximize the value of our solutions. We offer targeted and tailored technical support and maintenance to guarantee stability and security. Our customer success team performs regular health checks, business reviews and value assessments, offering optimization recommendations based on new product features and industry best practices to drive continuous improvement.

When delivering integrated hardware appliances, our delivery steps include contract-based inventory preparation, server assembly, software integration, outbound inspection and shipment to the customer. Software integration specifically involves configuring system disks, installing the operating system and management tools, deploying core software packages, performing network and hardware checks, and finally initializing the cluster, configuring storage resources, and activating licenses to enable the operation of the integrated hardware-software stack. These steps are completed at our production base. At the customer site, the process begins with cluster activation, followed by two parallel workflows: (i) application delivery track: completing application deployment and related configurations; and (ii) network setup track: rack mounting and network setup, modifying IP and network parameters, configuring Bond interfaces, setting time sources, and activating network cards.

If only software is delivered, hardware assembly and shipping are not required. The main steps at the customer site include network setup, operating system installation, management tools installation, deployment of core software packages, cluster initialization and license activation and application delivery.

Hardware-Software Customization

Our solution architecture is built on the principle that hardware and software are intrinsically inseparable. Core performance, stability and efficiency are achieved not by optimizing either layer in isolation, but through their coordinated design, tuning and validation as a single solution. To achieve deep integration of software and hardware, we apply four layers of customization at the hardware level to meet the specific needs of AI storage workloads:

- **AI-Specific Hardware Optimization:** We design and optimize hardware specifically for artificial intelligence workloads. For different T&I patterns, we build all-flash systems and hybrid nodes equipped with ultra-fast solid-state drives, large memory channels, and low-latency network connections. At the physical level, this includes customized chassis design, storage backplanes, and slot layouts to improve cooling, power delivery, vibration control, and maintenance efficiency. On the software side, we use our own high-speed data access technology and optimized data layouts to enable direct data transfer without unnecessary steps, maximizing speed and bandwidth.

- **High-Performance Network Customization:** We create high-speed network connections that allow data to move quickly and reliably. Our network engine supports multiple paths and routing, enabling extremely fast data transport. We carefully select and tune network interfaces and link configurations based on performance and cost goals. By fine-tuning for different graphics processors and network structures, we ensure that even under heavy workloads, the system maintains very low delay and stable throughput.

- **Converged Compute-and-Storage Design:** We integrate storage directly into computing servers so data can be accessed without extra switching, improving speed and efficiency. At the system level, we customize and validate settings to optimize processor performance, memory behavior, resource allocation, and cooling for storage-heavy tasks. At the server level, we reserve storage, memory, and bandwidth based on the ratio of graphics processors to CPUs, creating systems that balance computing and storage within a single machine.

- **End-to-End Engineered Hardware Systems:** We deliver fully engineered hardware systems at both appliance and cluster levels. Our end-to-end capabilities cover hardware selection, thermal and power design, EC verification and cluster assembly. This includes component-level firmware selection and long-duration compatibility testing, as well as system-level validation under high-load, power-failure and fault-injection scenarios. For AI factories, large-model inference clusters or AI data lake, we provide reference designs and deployments verified for consistency, latency, and reliability, creating a purpose-built, integrated compute-and-storage foundation.

Inventory Management

Our inventories primarily comprise raw materials, finished goods and goods in transit. To ensure effective inventory management and mitigate associated risks, we have established a comprehensive inventory control framework including procurement acceptance, warehousing, production requisition, outsourced processing, finished goods delivery, insurance, stocktaking, impairment testing, and disposal of obsolete items. Our inventory system clearly defines approval authorities and responsibilities and incorporates dual control through ERP systems and physical documentation, regular inventory counts, and insurance coverage to ensure accuracy, security and risk mitigation. In addition, we have implemented procedures for handling exceptions to achieve standardized and refined inventory management. At the end of each quarter, we conduct a full inventory count and perform ad hoc spot checks as necessary. Based on the results, we analyze discrepancies, identify underlying causes and implement corrective measures to maintain optimal stock levels.

INTELLECTUAL PROPERTY

Intellectual property rights are important to our business. Our future commercial success depends, in part, on our ability to obtain and maintain patents and other intellectual property and proprietary protections for commercially important technologies, inventions and know-how related to our business, defend and enforce our patents, preserve the confidentiality of our trade secrets, and operate without infringing, misappropriating or otherwise violating the valid, enforceable intellectual property rights of third parties. As of June 30, 2026, we had 103 registered patents in China. As of June 30, 2026, we also had 370 trademarks, 96 software copyrights and three domain names in China. During the Track Record Period and up to the Latest Practicable Date, we did not encounter any impediments in intellectual property rights applications.

The following table lists the portfolio of material patents, patent applications and software copyrights for our core technologies of which we were the registered owner as of the Latest Practicable Date:

	Patent Title	Specialist Technology Products	Functions	Application Date	Authorization Date
1.	Method, Device and System for Snapshot Based on Object Storage Bucket (基於對象存儲桶的快照方法、裝置和系統)	Data lake storage solutions	Provides high-performance snapshot creation for object storage buckets to support efficient data versioning and replication	2019/08/02	2021/02/19
2.	Method and Device for Network Sub-Health Detection in Distributed Storage System (分布式存儲系統網絡亞健康檢測方法及裝置)	Data lake storage solutions, storage solutions for T&I	Detects sub-healthy network states in distributed storage clusters to enable early fault identification and recovery	2019/12/26	2021/05/04
3.	Data Transmission Method and Device Based on FUSE File System and NFS Protocol (基於FUSE文件系統和NFS協議的數據傳輸方法及裝置)	Data lake storage solutions, storage solutions for T&I	Uses direct I/O with FUSE and NFS to bypass page cache and improve concurrent read-write performance	2020/06/28	2024/05/31

	Patent Title	Specialist Technology Products	Functions	Application Date	Authorization Date
4. . .	Method and Device for Object Merging (對象歸並方法及裝置)	Data lake storage solutions	Provides distributed QoS control for multi-gateway object storage by allocating resources across bucket queues and gateways	2021/09/28	2025/05/20
5. . .	Metadata Index Processing Method, Device, Electronic Equipment and Storage Medium (元數據索引的處理方法、裝置、電子設備及存儲介質)	Storage solutions for T&I	Caches hot metadata and optimizes index structures to improve read-write efficiency in distributed storage	2024/09/04	2024/12/13
6. . .	Data Storage Method, Device, Electronic Equipment and Program Product (數據的存儲方法及其裝置、電子設備及程序產品)	Storage solutions for T&I	Splits data into protected sub-blocks to enable integrity checking during storage	2025/07/08	2025/09/30
7. . .	Method and Device for Validity Verification of Snapshot Data, Electronic Equipment and Program Product (快照數據有效性的校驗方法及裝置、電子設備及程序產品)	Storage solutions for T&I	Verifies validity of snapshot data and recovers invalid snapshot blocks to improve snapshot reliability	2025/07/17	2025/11/07
8. . .	Data Shard Balancing Method and Device Based on Distributed Storage, Electronic Equipment (基於分布式存儲的數據分片均衡方法及其裝置、電子設備)	Storage solutions for T&I	Balances data shards across nodes in distributed storage by iteratively migrating unbalanced logical volumes	2025/07/08	2025/11/21

We confirm that all of the above listed patents are significant for carrying out the key functions of our Specialist Technology Products, and no other material patents are directly applied in our Specialist Technology Products.

Regarding the tenure of our intellectual properties: (i) for patents, according to the Patent Law of the PRC, the validity period of an invention patent is 20 years from the filing date; and (ii) for copyright, according to the Copyright Law of the PRC, except for the rights of authorship, modification and the protection of the integrity of the work, which are not subject to time limitations, the publication right of a legal entity's software copyright is protected for fifty years from the date of the work's first publication. The protection period for other copyrights is fifty years, ending on December 31 of the fiftieth year after the first publication.

Regarding the payment obligations in relation to our intellectual properties: (i) for issued invention patents, we are mainly required to pay the annual patent fee to competent authorities. We have kept track of the payment requirements for annual fees and made payment accordingly; and (ii) for pending patents, we are mainly required to pay the application fee, the substantive examination fee and the re-examination fee, depending on the examination progress, and we made the payment as required by the competent authorities as of the Latest Practicable Date. In 2023, 2024, 2025 and the six months ended June 30, 2026, our costs incurred in relation to the application, maintenance and protection of intellectual properties amounted to RMB58.6 thousand, RMB70.1 thousand, RMB59.9 thousand and RMB86.9 thousand, respectively. As the intellectual properties for each of our Specialist Technology Products are all self-developed, and have not been licensed or transferred from third parties, so there are no corresponding license or transfer fees that we are obligated to pay.

The term of an individual patent may vary based on the countries/regions in which it is granted. In China, the term of an issued patent for invention is generally 20 years from the filing date of the earliest non-provisional patent application on which the patent is based in the applicable country. The actual protection afforded by a patent varies on a claim-by-claim and country-by-country basis and depends upon many factors, including the type of patent, the scope of its

coverage, the availability of any patent term extension or adjustment, the availability of legal remedies in a particular country/region and the validity and enforceability of the patent. We cannot provide any assurance that patents will issue with respect to any of our owned or licensed pending patent applications or any such patent applications that may be filed in the future, nor can we provide any assurance that any of our owned or licensed issued patents or any such patents that may be issued in the future will be commercially useful in protecting our product candidates and methods of designing the same. See "Risk Factors — Risks Relating to Our Intellectual Property Rights — Our patents and other non-patented intellectual properties are valuable assets. Insufficient intellectual property protection may reduce our competitive advantage and lead to conflicts that adversely affect our business, financial condition and results of operations."

We rely in part on trade secrets, confidential information and contractual protections, including confidentiality, non-competition and intellectual property assignment agreements, to safeguard our proprietary technology, data and know-how. However, these measures may not provide adequate protection and may be breached, resulting in unauthorized disclosure, misappropriation, misuse or independent development of our proprietary information by third parties. In addition, despite maintaining physical and electronic security measures, unauthorized parties may obtain access to or use information that we regard as proprietary. Any failure to protect our trade secrets, confidential information, data or other intellectual property could materially and adversely affect our competitive position, business, financial condition and results of operations.

As of the Latest Practicable Date, we were not involved in any legal, arbitral or administrative proceedings or claims of infringement of any intellectual property rights, in which we may be a claimant or a respondent, nor were we subject to any legal claims or proceedings that may have an influence on the R&D for Specialist Technology Products. Our Directors confirm that they are not aware of any legal, arbitral or administrative proceedings of infringement of any third parties' intellectual property rights by us as of the Latest Practicable Date. See "Risk Factors — Risks Relating to Our Intellectual Property Rights."

DATA SECURITY AND PRIVACY

During the Track Record Period and up to the Latest Practicable Date, our users are enterprises rather than individual customers. Under our direct sales channel, we collect data directly from customers, primarily including (i) basic corporate and contact information, such as customer name, business address and contact person details; and (ii) transaction-related information, such as purchase orders, product specifications, pricing, billing and payment records. Under our distributorship channel, we collect similar categories of data from our distributors, including their corporate and contact information and transaction-related information. In addition, we may obtain limited information relating to the underlying end-customers on a need-to-know basis, such as end-customer name, location and order-related information for the purpose of facilitating transactions and service delivery. Other than the types of personal information mentioned above, our collected data do not involve any personal or sensitive data. Our AI storage solutions are deployed on our customers' own on-premises infrastructure and operate entirely within customer-controlled environments. We do not conduct AI model training as part of our current business operations. We do not otherwise obtain or store customers' business data, personal information or other sensitive content, nor do we provide any such data to third parties during the R&D stage or the provision of after-sales maintenance services. Save for limited business contact information collected for communication purposes, we do not directly or indirectly obtain or store customers' data in our non-core business activities. Such collection and use of data complied with all applicable laws and regulations relating to data privacy and data security during the Track Record Period and up to the Latest Practicable Date.

During our business operations, and only with the prior consent of customers, suppliers and other partners, we collect and maintain essential contact information necessary for business communication in accordance with applicable laws and regulations. We have implemented a comprehensive data security management framework, including an Information Security Risk Assessment and Remediation Procedure. Our measures include the following:

- We have implemented internal policies and guidelines in relation to data compliance management, classification and grading and information security incident management. Our internal control guidelines cover the full lifecycle of data processing including data ingestion, data encryption and transportation, data storage security, data backup and recovery, data processing, proper use of data, and data destruction and disposition.

- We implement internal authorization and authentication procedures and policies, which require that our employees only have access to data which is directly relevant and necessary for their responsibilities and for limited purposes and are required to verify authorization beyond such scope. We also implement corresponding authorization and authentication system to ensure that our confidential and important data can only be accessed for authorized use and by authorized personnel. In addition, we conduct periodic reviews of user access rights to ensure that granted permissions remain appropriate and aligned with job responsibilities, and promptly adjust or revoke access where necessary.

- We conduct mandatory training for new employees involved in data-related services and apply encryption, identity authentication and access control measures to safeguard data.

- Our legal department serve as our data compliance management team, which is responsible for formulating data and information security strategies, making decisions in important data and information security incidents, and overseeing, inspecting and guiding all data compliance related works. In addition, engineers from our R&D support department are responsible for the security management and daily maintenance of our network system, coupled with regular inspections with traceable records.

- We maintain firewalls and protective systems designed to prevent data loss or leakage caused by cyberattacks. To minimize the risk of data loss or leakage, we conduct regular data backup and data recovery tests.

- We monitor regulatory developments relating to cybersecurity and data protection and conduct routine internal compliance checks and rectification.

As a proof of the security and reliability of our data protection technologies, we have completed various information security, privacy and compliance certifications/validations.

We plan to develop an API ecosystem as part of our future business development. See "Future Plans and Use of [**REDACTED**]." Such APIs are designed to facilitate integration with third-party systems and do not, in themselves, involve the collection or processing of personal data. The operation of such APIs will be subject to our internal data governance policies and access control mechanisms. Accordingly, we do not expect the development of such API ecosystem to have any material implications for our compliance with applicable data protection and privacy laws and regulations.

During the Track Record Period and up to the Latest Practicable Date, (i) we had not received any claim from any third party against us on the ground of infringement of any third party's right to data and privacy protection as provided by any applicable laws and regulations, (ii) there had been no investigation or other legal proceeding pending or threatened against us initiated by competent government authorities or third parties with respect to cybersecurity, data and personal information protection, and (iii) we had not experienced material leakage of personal information

in relation to laws and regulations of cybersecurity, data protection and personal information protection. Based on the above, as well as that (i) we currently do not directly obtain or store customer's data in our primary business; (ii) we have implemented internal policies and technological measures on protecting cybersecurity, data security and personal information; and (iii) we will continuously pay close attention to the legislative and regulatory development in cybersecurity, data protection, personal information protection and the legality use of the training data, maintain ongoing communication with relevant government authorities and implement all necessary measures in a timely manner to ensure continuous compliance with the relevant laws and regulation, our PRC Legal Advisor is of the view that, during the Track Record Period and up to the Latest Practicable Date, (i) we have implemented compliance measures concerning cybersecurity, data protection and personal information protection in accordance with the requirements of relevant cybersecurity, data and personal information protection laws and regulations in all material aspects and (ii) has complied with the relevant data protection and privacy, and cybersecurity regulations in all material aspects. See "Risk Factors — We may be subject to complex and evolving laws and regulations regarding cyber security, personal information protection and data protection, which may lead to increased compliance costs, and adversely affect our reputation and business operations." Based on the due diligence conducted by the Sole Sponsor, nothing has come to the Sole Sponsor's attention that would reasonably cause the Sole Sponsor to disagree with the PRC Legal Advisor's view that the Company has complied with the relevant data protection and privacy, and cybersecurity regulations in all material aspects.

During the Track Record Period and up to the Latest Practicable Date, we had not experienced any material IT system failure or downtime.

SALES AND MARKETING

Our Sales Networks

We primarily serve our users through selected distribution channels and also operate a direct sales model. The following table sets forth our revenue breakdown by sales channel in absolute amount and as percentage of our total revenue for the years/periods indicated:

	Year ended December 31,						Six months ended June 30,			
	2023		2024		2025		2025		2026	
	RMB'000	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*
							(Unaudited)			
Distribution	98,223	58.9	102,905	59.7	190,357	71.3	62,529	62.8	121,883	65.6
Direct sales	68,618	41.1	69,579	40.3	76,656	28.7	37,000	37.2	63,877	34.4
Total	**166,841**	**100.0**	**172,484**	**100.0**	**267,013**	**100.0**	**99,529**	**100.0**	**185,760**	**100.0**

Direct Sales

We have built a sales and marketing team with relevant industry experience to identify market opportunities and communicate the value of our technologies and the performance of our solutions. As of June 30, 2026, our sales and marketing team comprised 162 employees. We have established sales offices in major cities in China. Through these sales offices, we have extended our reach to almost all provinces in China. Leveraging its deep industry experiences, our sales and marketing team identifies market trends and customers' demands thoroughly and simultaneously works closely with our R&D team to ensure that they can accurately address customer pain points and deliver products and services to the customers' satisfaction in a timely manner. The salient terms of our standard sales agreements or service agreements for our solutions with direct sales customers during the Track Record Period are set out below:

- ***Product specifications***. Our customers typically set forth specific product specifications for products ordered, such as solution name, configuration and features.

- ***Term***. For solutions and products, we generally enter into one-off sales agreement under which the contract automatically terminates upon both parties fulfilling their contractual obligations. For service agreements, the term is typically one year.

- ***Payment and credit term***. For our solutions and products, we typically charge a fixed total price for the products purchased and generally offer our direct sales customers a credit period of up to 60 days. In addition, we provide on-site operation and maintenance services for direct sales customers, for which we typically charge a pre-agreed service fee quarterly.

- ***Warranty and maintenance services.*** We usually provide a 36-month warranty and maintenance period for our products and solutions.

- ***Intellectual Property***. We retain full intellectual property rights to our products. We grant direct sales customers a license to use our intellectual property solely for purposes related to the agreement. Without our prior written consent, customers may not rent, copy, imitate or modify our products.

- ***Confidentiality***. All confidential information exchanged between the parties shall not be disclosed to any third party.

- ***Termination***. The agreement may be unilaterally terminated under circumstances set forth in the agreement, including prolonged non-payment by the customer, defects in the products we provide, or material breach of contract terms by either party.

During the Track Record Period and up to the Latest Practicable Date, we did not experience any material breach of agreement with our customers.

Distributors

During the Track Record Period, we also sold our AI storage solutions through third-party distributors. Revenue from our ten largest distributors in each year/period during the Track Record Period amounted to RMB83.5 million, RMB81.5 million, RMB125.2 million and RMB94.9 million in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively. Our distributors primarily engage in downstream sales and channel development in their respective regions and industry verticals, which primarily involve customer acquisition, identification of sales opportunities, industry promotion activities and facilitating introductions to potential channel partners. After purchasing our solutions, they may perform onward sales to their respective customers, which mainly comprise enterprises in intelligent manufacturing, financial institutions, telecommunications operators and technology enterprises, among others.

During the Track Record Period, a majority of our customers were distributors. This is primarily because distributors enhance market coverage and customer acquisition efficiency through their broader sales networks and stronger local presence. In addition, certain enterprise customers prefer or are required to procure IT solutions through certified or approved distributors under their internal procurement policies, or due to their established relationships and trust with such distributors, who are often able to provide integrated offerings and value-added services. Distributors also support localized delivery, including on-site and ongoing services, particularly in regions where our direct sales coverage is limited. Accordingly, customers may choose to purchase our solutions through distributors rather than directly from us as a matter of commercial convenience and operational efficiency. In performing onward sales, distributors had advantages in leveraging their established relationships with end customers and their understanding of customers' deployment environments and operational requirements to configure and on-sell different combinations of our turnkey appliances and software solutions. Such experience and know-how are primarily derived from distributors' established customer networks, their familiarity with procurement and deployment scenarios in their respective regions or industry verticals, and, in

certain cases, their prior system integration experience. Rather than relying on a single uniform go-to-market approach, we leverage distributors to translate our core products into deployable solutions that are suitable for different customer environments and operational constraints. At the same time, we retain control over our core technologies, intellectual property and key technical know-how to avoid potential misappropriation by distributors. Our proprietary technologies, including core software architecture, source code, product design and technical know-how, are not transferred to distributors; rather, they are only granted limited rights to market and resell our offerings. Such arrangements are governed by contractual safeguards, including confidentiality obligations and restrictions on the unauthorized use or disclosure of our proprietary information. Key technical support and deployment activities are typically provided directly by us. Thus, distributors primarily function as marketing and sales agent. While the scope of rights granted to distributors is limited, they play an important role in facilitating market access and expanding customer coverage in their respective regions and industry verticals, without requiring us to deploy extensive direct sales resources.

We believe that our distributor model enables us to (i) diversify our revenue sources without over-reliance on any single region or type of sales channel, by leveraging distributors as our commercial, technical and service extension across different geographies and industry verticals; (ii) tailor our sales and marketing strategies for different customer types, regions and sales channels; and (iii) strengthen our market presence by allowing our solutions to penetrate markets more quickly and effectively. We had not engaged any sub-distributors during the Track Record Period and up to the Latest Practicable Date. In particular, revenue generated from distributors as a percentage of our total revenue increased from 59.7% in 2024 to 71.3% in 2025, and increased from 62.8% in the six months ended June 30, 2025 to 65.6% in the six months ended June 30, 2026, primarily driven by the rapid expansion of market demand, during which distributor channels provided broader customer reach and faster market coverage. Leveraging their established customer networks and capabilities in delivering integrated AI infrastructure solutions, distributors were able to efficiently serve geographically dispersed end customers through one-stop procurement and delivery. As a result, we continued to expand our distributor partnerships, which contributed an increasing share of our revenue.

Our distributors independently determine the selling prices to end customers based on the specific project requirements. Notwithstanding the above, we promote consistency in pricing outcomes through various indirect mechanisms, including:

- *Unified pricing system:* We apply a unified pricing framework to both direct sales customers and distributors, and do not differentiate our pricing based on sales channels. Distributors' profit margins are derived from the value-added services they provide, rather than from any direct price concessions granted by our Group.

- *Regional management arrangements:* Distributors are assigned defined sales territories to minimize inter-regional price competition and cross-territory sales, thereby maintaining orderly pricing within each region.

- *Direct sales price anchoring:* Transaction prices under the direct sales channel serve as an important market reference for distributor pricing, as distributors' resale prices are indirectly constrained by the fact that end-users may, in appropriate circumstances, procure directly from us on comparable commercial terms.

- *Monitoring of overdue distributor orders:* Our sales operations team regularly monitors distributors' overdue orders and sets internal limits on overdue amounts and order counts. Distributors exceeding such thresholds may face reduced order allocation or termination of cooperation. This mitigates the risk of aggressive low-price selling strategies.

- *Market competition and tendering discipline:* A significant portion of end-customer projects are awarded through competitive bidding processes, where pricing is determined by market competition. In addition, the PRC distributed AI storage market is relatively fragmented (with the top five players accounting for approximately 47.9% of the market based on installed capacity in 2025), providing customers with multiple alternatives and further constraining pricing deviations. These external market forces operate in conjunction with our internal mechanisms.

During the Track Record Period and up to the Latest Practicable Date, we had not identified any instance where distributors materially disrupted or adversely affected end-customer pricing.

The following table sets forth the average customer value under distributorship model during the Track Record Period:

	Year ended December 31,			Six months ended June 30,
	2023	2024	2025	2026
Distributorship revenue (RMB)	98,223,495	102,905,049	190,356,749	121,882,740
Number of distributors	178	172	185	165
Average customer value (RMB)	551,817	598,285	1,028,955	738,683

In selecting distributors, we take into account a number of factors, including company size, professionalism and adequacy of personnel, historical business performance, business integrity, local market development capabilities, as well as technical expertise. We generally enter into standard distribution agreements with our distributors and monitor their performance through periodic reviews. On an annual basis, we conduct dynamic evaluations to determine whether to continue, strengthen or reduce our business relationships with distributors, considering factors such as sales contribution, channel expansion, payment collection capabilities and the competence of their sales and technical personnel. The selling price for each of our transactions with distributors is determined on a case-by-case basis by reference to a combination of factors, including the deployment scale of the project, technical configuration and complexity, prevailing market conditions, underlying cost structure and the distributor's cooperation level and historical performance. In addition, we implement a tiered channel management approach, under which we categorize distributors into different tiers based on factors such as their industry focus and strengths, business scale, sales capabilities, technical and service capacities, and credit profile. Distributors in different tiers are subject to different commercial terms and incentive arrangements, such as differentiated credit terms. This mechanism is designed to allocate channel resources more effectively and incentivize distributors to expand their business coverage and improve service quality. Furthermore, we provide distributors with comprehensive technical and sales training to enhance their operational, delivery and after-sales service capabilities. Our Directors confirm that we have no ownership or management control over any of our distributors, and they operate independently from our Group. All of our distributors during the Track Record Period and up to the Latest Practicable Date were our Independent Third Parties. According to CIC, it is consistent with industry norm for providers of AI storage solutions to engage distributors to leverage their knowledge of local marketing practices and consumer preferences.

We have a buyer-seller relationship with distributors. We sell our solutions to distributors in accordance with our distribution agreements and the specific purchase orders placed by them, and we do not impose minimum purchase or sales targets. Title to, and the risk of damage of our AI storage solutions passes to the distributors upon delivery. Our distributors' subsequent sales to downstream customers, including customer selection, sales strategy and inventory management are entirely at their discretion. Our distributors are solely responsible for managing their inventories, and any risks arising from which are borne by them. Under the distributorship model, revenue is

recognized upon receipt of acceptance confirmation from the distributor. During the term of the distribution agreements, our distributors are generally not allowed to return unsold products to us. Product return policies do not apply to our distributors, except that in accordance with relevant laws and regulations, our distributors may negotiate with us for the return of, or indemnification in respect of products that are defective due to our fault. During the Track Record Period and up to the Latest Practicable Date, we had not received any product return requests, and the amount of returned products was nil. We had not received any indemnification claims from our distributors, nor had we made any provisions for potential product returns or indemnification.

During the Track Record Period, certain end customers purchased our solutions through different channels, including via multiple distributors and, in some cases, both directly from us and through distributors. This primarily reflects differences in project context, procurement processes and service requirements. Generally, distributorship is adopted where customers require localized service, while direct sales are used where customers require direct technical engagement or contracting arrangements. Considering the above factors in adopting distributorship and direct sales by end customers, we do not believe that there is a material risk of cannibalization. Taking into account all sales channels, the largest end customer in each year/period during the Track Record Period, all of whom were served through our direct sales channel, accounted for approximately 14.4%, 10.2%, 7.8% and 10.4% of our total revenue in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, while the top five end customers in aggregate accounted for approximately 30.9%, 28.5%, 20.1% and 36.0% of our total revenue for the same years/periods, respectively. While the level of customer concentration has fluctuated during the Track Record Period, we generally do not rely on any single end customer, and our revenue base has become more diversified during the Track Record Period.

Under the distributorship model in particular, we note that a portion of our revenue was generated from sales of our solutions to the same underlying end customers through multiple distributors. Such revenue amounted to approximately RMB46.8 million, RMB39.4 million, RMB95.1 million and RMB61.4 million in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, representing approximately 28.1%, 22.8%, 35.6% and 33.1% of our total revenue for the respective years/periods. Despite the above, the revenue attributable to each individual end customer served through multiple distributors remained limited. In particular, the largest such end customer accounted for no more than 7.0% of our total revenue in each year/period during the Track Record Period. This arrangement primarily reflects the project-based nature of our business and the involvement of different distributors in separate projects, procurement batches or regional deployments for the same end customer, rather than any reliance on a single distributor-end customer pair. Despite such overlaps, the relevant revenue is typically fragmented across multiple contracts and distributors, and no single end customer accounted for a material portion of our revenue during the Track Record Period. Accordingly, we did not have material revenue concentration in any individual end customer under the distributorship model.

To manage potential overlap and avoid channel conflict or cannibalization, we have established a structured customer management mechanism supported by our customer relationship management (CRM) system. Under this mechanism, all sales channels, including direct sales and distributors, are required to register and obtain internal approval for customer engagement through the CRM system. Customer allocation is generally granted on a first-registration basis, and once a customer is registered under a particular channel, other channels are restricted from engaging with the same customer without prior approval. We also maintain a dynamic allocation mechanism. For example, where no substantive follow-up is recorded within a specified period, the relevant customer may be released for reallocation, allowing other channels to apply for engagement through the CRM system. In addition, we may adjust channel allocation based on customer preferences at different stages of engagement. Through these measures, we aim to ensure clear customer ownership across channels and effectively mitigate potential conflicts or cannibalization.

During the Track Record Period, our designated sales territories were generally divided into five major geographical regions in China. The following table sets forth a breakdown of our revenue generated from sales to distributors by designated sales territory during the Track Record Period:

	Year ended December 31,						Six months ended June 30,			
	2023		2024		2025		2025		2026	
	RMB'000	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*	*RMB'000*	*%*
							(Unaudited)			
East China	22,899	23.3	29,190	28.4	51,482	27.0	21,147	33.8	40,122	32.9
North China	43,301	44.1	25,612	24.9	65,413	34.4	20,567	32.9	41,913	34.4
Central China	8,347	8.5	8,233	8.0	17,315	9.1	4,306	6.9	7,544	6.2
South China	8,949	9.1	15,265	14.8	33,309	17.5	9,674	15.5	23,508	19.3
West China	14,727	15.0	24,605	23.9	22,838	12.0	6,835	10.9	8,796	7.2
Total	**98,223**	**100.0**	**102,905**	**100.0**	**190,357**	**100.0**	**62,529**	**100.0**	**121,883**	**100.0**

We maintain a centralized inventory control and management policy under which inventory is fully controlled, recorded and monitored through our ERP system prior to delivery, and each shipment to distributors is subject to internal approval, outbound documentation and receipt confirmation. We generally prohibit speculative stocking by distributors without underlying orders. Upon delivery, title and risk pass to the distributors.

During the Track Record Period and up to the Latest Practicable Date, to the best knowledge of our Directors, there was no material channel stuffing or cannibalization issue among our distributors. Considering that we do not impose minimum purchase requirements on distributors, and that our distributors are not allowed to return any unsold products to us, our Directors are of the view that the channel stuffing issue is not applicable to us. To further minimize the potential cannibalization risks, we (i) set internal limits on the number and value of overdue orders from distributors, which refer to orders placed by distributors for which payment remains outstanding beyond the payment due date; if a distributor exceeds these thresholds, our collaboration with such distributors may be reduced or even terminated, thereby discouraging distributors from inflating sales through excessive inventory purchases that cannot be promptly sold or paid for; (ii) only permit distributors to operate within their designated authorized regions; and (iii) employ a stringent and dynamic channel qualification and certification process; distributors must consistently meet our standards for payment collection to maintain or enhance their partnership status with us.

Set forth below is a summary of the key terms that are common to most of our agreements with distributors:

- ***Geographic or other exclusivity.*** Distributors are only permitted to sell within their designated sales territories.

- ***Minimum purchase amounts or sales target.*** We do not set minimum purchase amounts for our distributors. We negotiate annual recommended sales targets with them, and their long-term unsatisfactory sales performance may result in a reduction of the level of our cooperation with them (including the discount levels or priority in product allocation).

- ***Logistics.*** We are typically responsible for delivering the goods to the location designated by the distributors.

- ***Product return.*** We generally do not accept product returns. The risk transfers to the distributors upon the delivery of the products. If the goods fail inspection or quality issues arise during the warranty period, we will be responsible for negotiating a resolution and arranging repair or replacement of qualified products.

- ***Payment.*** We generally offer distributors a credit period of up to 60 days.

- ***Termination.*** We may terminate a distribution agreement upon expiry without renewal, by mutual written agreement, or for an unremedied breach after notice. We may also terminate immediately for specified serious misconduct, such as leakage of confidential information and violation of anti-bribery obligations.

Movement of Distributors

The following table sets forth the total number of our distributors, the number of new distributors and the number of distributors whose distributorship relationships were ended during the year/period indicated:

	Year ended December 31,			Six months ended June 30,
	2023	2024	2025	2026
Number of distributors as of the beginning of year/period	305	340	364	404
Addition of new distributors[(1)]	35	24	40	17
Number of terminated distributors for the year/period	–	–	–	–
Number of distributors at the end of year/period	340	364	404	421
Number of inactive distributors for the year/period[(2)]	162	192	219	256

Notes:

(1) Refers to distributors from whom we had generated revenue for the first time during the corresponding year/period.

(2) Refers to distributors from whom we did not generate any revenue for the year/period, but with whom we maintained business relationship for potential future sales.

Pricing

We price our solutions with reference to multiple factors, including (i) the scale of deployment, including the number of petabytes and overall system capacity involved; (ii) the technical complexity; (iii) our cost structure (software and hardware components and associated R&D expenditure); (iv) prevailing market prices in the relevant competitive landscape; and (v) the nature of depth of our relationship with the customer. We maintain open communications with customers on pricing and typically contract on a fixed total, project-based basis for a bundled scope that combines standard modules with any approved custom elements, taking into account customer requirements, negotiated volume discounts and tender conditions. Our data lake storage solutions and storage solutions for T&I generally follow a consistent pricing logic.

Although our current model is not subscription-based, we generate substantial follow-on revenue from capacity expansions as customers' datasets and use cases scale. Under a single customer ID, multiple incremental purchase orders may be placed over time for new workloads or growing datasets, creating a recurring pattern of add-on sales. In addition, for limited scenarios we provide time-and-materials support services priced on a per-person-day basis, which are quoted separately from the product price.

Our contracts commonly include maintenance and support for three years as standard, with optional multi-year terms (for example, three to five years) where required by tenders or customer policy. Standard maintenance services under our AI storage services are typically bundled with product sales and are priced at approximately 6% to 10% per annum of the contract value. Upon expiry of the standard warranty and maintenance services, customers are notified through our management interface that the warranty has lapsed; thereafter, maintenance services (including ticket handling and remote/video support) are offered on a paid basis. Post-warranty maintenance services under our AI storage services are typically priced as a percentage of the original selling price or list price, and is determined through negotiation having regard to market conditions, competition, service scope and service levels. Such services are typically provided for one year, with optional multi-year terms (for example, two to three years) at the request of customers.

We monetize our data lake storage solutions and storage solutions for T&I primarily through fixed-price project-based contracts. Revenue is recognized when delivery is completed and the relevant customer acceptance confirmation is obtained. We monetize our AI storage services primarily through annual service fees, which are generally recognized over the service period on a time-apportioned basis. Depending on the sales channel, the paying customer is either the end customer, in the case of direct sales, or the distributor, in the case of distributor sales.

The following charts illustrate our transaction and fund flows between stakeholders:

Direct sales model

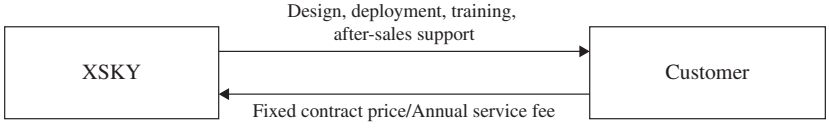

Distributor model

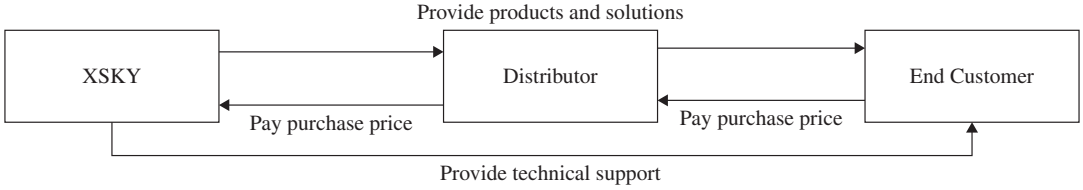

The average customer value for direct sales ranged from approximately RMB0.7 million to RMB0.8 million during the Track Record Period. According to CIC, the pricing of our products, services and solutions generally falls within the prevailing market price range for comparable offerings in the industry.

Marketing and Branding

We pursue an integrated marketing and branding strategy that combines offline and online channels to reach enterprise buyers and developers. Offline, we participate in industry events to highlight customer success and technical innovation; online, we drive awareness through content marketing, social media, and developer platforms.

Our go-to-market model prioritizes industry-led, head-list coverage over traditional regional segmentation. We focus on leading enterprises in intelligent manufacturing, finance, telecommunications, and technology, expanding into adjacent sectors as product-market fit strengthens. Marketing operations are structured into three teams: (i) Strategic Partnerships, which cultivates benchmark relationships and ecosystem alliances; (ii) Industry and Key Accounts, which delivers tailored engagement in priority verticals; and (iii) Regional Marketing, which supports local penetration and channel enablement.

As brand recognition has grown, customer acquisition costs have declined. We continue to invest in high-impact activities while selectively expanding coverage to AI-intensive accounts, balancing lead generation efficiency with deeper vertical penetration.

CUSTOMERS

We primarily sell our solutions to enterprise customers in China (including Hong Kong), with a small portion of our revenue derived from overseas markets such as Singapore and Indonesia, across sectors such as intelligent manufacturing, financial institutions, telecommunications operators and technology enterprises, among others.

A substantial portion of our five largest customers in each year/period during the Track Record Period were distributors, and sales to such distributors reflect their downstream distribution to multiple end customers rather than demand from a single end customer. We typically enter into standard sales agreements or service agreements with our direct sales customers. See "— Sales and Marketing — Our Sales Networks — Direct Sales." We typically enter into standard distribution agreements with our distributors. See "— Sales and Marketing — Our Sales Networks — Distributors."

Although the contractual terms governing our customer arrangements are generally short in duration, our customer relationships typically extend over a long period in practice. Although our business model is not subscription-based, we generate recurring follow-on revenue from existing customers as their datasets, workloads and use cases expand. Customers may place multiple incremental purchase orders over time under a single customer ID for capacity expansions or new deployments, resulting in a recurring pattern of add-on sales. Our contracts also generally include three-year maintenance and support, with optional multi-year terms where required by customer policies or tenders. Our customer stickiness is reflected in our operating metrics. Our overall net dollar retention rate reached 160.0% in the six months ended June 30, 2026.

According to CIC, a certain level of customer concentration is common for independent vendors in our sector. Looking ahead, while we expect our solutions to continue to represent a meaningful share of customers' AI storage architectures, we plan to gradually reduce customer concentration by expanding our addressable customer base through continued product iteration, new product launches and enhanced industry-focused go-to-market capabilities to acquire new customers across additional verticals and use cases.

The tables below set forth the information of our top five customers in each year/period during the Track Record Period:

Year ended December 31, 2023

Customers	Customer type	Solutions/ services provided	Revenue	% of total revenue	Year of commencement of business relationship	Credit terms	Payment method
			(RMB'000)				
Customer A[(1)]. .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	28,576	17.1	2021	40 days	Telegraphic transfer
Customer B[(2)]. .	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	24,014	14.4	2021	40 days	Telegraphic transfer

Customers	Customer type	Solutions/ services provided	Revenue	% of total revenue	Year of commencement of business relationship	Credit terms	Payment method
			(RMB'000)				
Customer C[3]	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	13,075	7.9	2020	30 days	Telegraphic transfer
Customer D[4]	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	12,395	7.4	2020	30 days	Telegraphic transfer
Customer E[5]	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	10,136	6.1	2021	30 days	Telegraphic transfer
Total			88,196	52.9			

Notes:

(1) Customer A is a company based in Sichuan, China, primarily engaged in ICT consumer product distribution business and ICT corporate product distribution business. It is a subsidiary of a company listed on the Hong Kong Stock Exchange. Its ultimate end customers primarily consisted of enterprises and institutions in the energy, healthcare, financial services and technology sectors.

(2) Customer B is a company based in Shanghai, China, primarily engaged in software and information technology services.

(3) Customer C is a company based in Beijing, China, primarily engaged in information system integration, including data center solutions, data storage, cloud computing management platforms, and network security products. Its ultimate end customers primarily consisted of educational institutions, transportation companies, research organizations and technology companies.

(4) Customer D is a company based in Chongqing, China, primarily engaged in consumer electronics, enterprise systems and cloud computing. It is a subsidiary of a company listed on the Hong Kong Stock Exchange. Its ultimate end customers primarily consisted of healthcare institutions, utility service providers, financial institutions, insurance companies and telecommunications companies.

(5) Customer E is a company based in Nanjing, China, primarily engaged in software and information technology services.

Year ended December 31, 2024

Customers	Customer type	Solutions/ services provided	Revenue	% of total revenue	Year of commencement of business relationship	Credit terms	Payment method
			(RMB'000)				
Customer A	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	36,671	21.3	2021	40 days	Telegraphic transfer
Customer B	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	17,486	10.1	2021	40 days	Telegraphic transfer
Customer E	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	16,616	9.6	2021	30 days	Telegraphic transfer

Customers	Customer type	Solutions/ services provided	Revenue	% of total revenue	Year of commencement of business relationship	Credit terms	Payment method
			(RMB'000)				
Customer D . . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	9,898	5.8	2020	30 days	Telegraphic transfer
Customer F[(1)] . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	7,954	4.6	2020	30 days	Telegraphic transfer
Total			88,625	51.4			

Notes:

(1) Customer F is a company based in Shanghai, China, primarily engaged in IT product sales, IT technical services, software development, and information system integration solutions. Its ultimate end customers primarily consisted of utility service providers, technology companies, financial institutions and industrial companies.

Year ended December 31, 2025

Customers	Customer type	Solutions/ services provided	Revenue	% of total revenue	Year of commencement of business relationship	Credit terms	Payment method
			(RMB'000)				
Customer A . . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	45,684	17.1	2021	40 days	Telegraphic transfer
Customer E . . .	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	20,722	7.8	2021	30 days	Telegraphic transfer
Customer B . . .	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	14,807	5.5	2021	40 days	Telegraphic transfer
Customer D . . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	11,029	4.1	2020	30 days	Telegraphic transfer
Customer F . . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	10,217	3.8	2020	30 days	Telegraphic transfer
Total			102,459	38.3			

Six months ended June 30, 2026

Customers	Customer type	Solutions/ services provided	Revenue	% of total revenue	Year of commencement of business relationship	Credit terms	Payment method
			(RMB'000)				
Customer A . . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	41,057	22.1	2021	40 days	Telegraphic transfer
Customer E . . .	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	19,701	10.6	2021	30 to 150 days	Telegraphic transfer
Customer B . . .	Direct customer	Data lake solutions, storage solutions for T&I, AI storage services	17,103	9.2	2021	90 days	Telegraphic transfer
Customer C . . .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	13,703	7.4	2020	30 days	Telegraphic transfer
Customer G[(1)] .	Distributor	Data lake solutions, storage solutions for T&I, AI storage services	11,596	6.2	2024	30 days	Telegraphic transfer
Total			103,160	55.5			

Notes:

(1) Customer G is a company based in Guangzhou, China, primarily engaged in providing data lifecycle management and digital transformation solutions and services. It is a subsidiary of a company listed on the Taiwan Stock Exchange. Its ultimate end customers primarily consisted of enterprises and institutions in the intelligent manufacturing, embodied intelligence and life sciences sectors.

As of the Latest Practicable Date, none of our Directors, their respective close associates or any of our shareholders (who, to the knowledge of our Directors, owned more than 5% of our issued share capital) had any interest in any of our five largest customers in each year/period during the Track Record Period.

After-Sales Services

We provide structured and responsive after-sales services to support the deployment and operation of our AI storage solutions. Services are delivered through a standardized incident handling framework that spans the full lifecycle of customer engagement. Customers may submit service requests via designated portals, hotlines, email, or account representatives. Upon receipt, our personnel collect relevant operational context and confirm service level entitlements. Verified cases are logged by Level-1 technical support, which conducts initial diagnostics and proposes solutions based on internal knowledge bases. Customers may be asked to provide logs, system data, or on-site access to facilitate resolution. If unresolved, cases escalate to higher-level support and, if necessary, further escalate to functional teams for root cause analysis and final resolution. Closure is confirmed upon system recovery. We offer multiple service tiers, including standard three-year business-hour support, with optional upgrades to three- or five-year 24/7 support and on-site services. Support is delivered through online guidance, documentation, and field services. This structure enables timely resolution, minimizes operational disruption, and ensures sustained system performance.

We normally provide a standard warranty period of 36 months from acceptance for our products, which primarily covers defects in hardware components of the turnkey appliances. Within the product warranty period, warranties associated with the turnkey appliances are provided free of charge and serve as an assurance that the products sold comply with agreed-upon specifications. We also provide software maintenance and upgrade services which primarily cover version upgrade and defect rectification of the delivered software embedded in the turnkey appliances or sold separately. Software maintenance and upgrade services provided by us generated revenue of RMB29.1 million, RMB34.4 million, RMB38.1 million, RMB18.1 million and RMB20.4 million in 2023, 2024, 2025 and the six months ended June 30, 2025 and 2026, respectively. During the Track Record Period, we recorded no product returns in 2023, 2024, 2025 and the six months ended June 30, 2026. Revenue generated from software maintenance and upgrade services is included in our AI storage services revenue. During the Track Record Period and up to the Latest Practicable Date, we had no material disputes with customers regarding product returns or refunds.

Third-party Payment

During the Track Record Period, we had one instance where the payment due from a customer (the "**Relevant Customer**") was settled by a third-party payor. In 2019, the Relevant Customer purchased turnkey appliances from us and subsequently failed to settle the outstanding balance. After multiple unsuccessful attempts to collect payment, we initiated litigation and obtained a court order freezing the Relevant Customer's bank account. We reached a settlement during litigation proceedings, as reflected in a court-approved mediation decision (the "**Settlement**") in 2025, under which the Relevant Customer was obligated to pay us RMB1.3 million.

To satisfy its payment obligations under the Settlement, the Relevant Customer proposed that the Payor make payment to the Company on its behalf (the "**Third-Party Payment Arrangement**"). The Company did not select, designate, or recommend the Payor; rather, the arrangement was proposed unilaterally by the Relevant Customer. The Relevant Customer and the Payor entered into two written delegation letters (the "**Delegation Letters**") under which the Payor was entrusted to pay a total of approximately RMB0.8 million. Further, the remaining RMB0.5 million was paid directly by the Relevant Customer's legal representative (the "**Representative**"). Both the Relevant Customer and the Payor are Independent Third Parties.

The Third-Party Payment Arrangement arose because the Relevant Customer encountered difficulty making payment directly from its own corporate account, which had been frozen during the litigation process and otherwise lacked sufficient funds.

In order to comply with the settlement timeline and avoid breaching the Settlement, the Relevant Customer proposed that the Payor and Representative discharge the outstanding amounts on its behalf. We agreed to this proposal in order to expedite receipt of the long-overdue payments.

As advised by our PRC Legal Advisor, in light of the Settlement which gave rise to the Third-party Payment Arrangement, (i) the Third-party Payment Arrangement was not in breach or contravention of mandatory requirements of applicable laws or regulations in China (including but not limited to anti-money laundering laws); (ii) the content of the Delegation Letter is lawful and valid, and constitutes a legally binding obligation on the part of the Payor; and (iii) we had confirmed with the Representative on his intention to pay on behalf of the Relevant Customer, and the Representative made special notation to the wire transfer indicating that the payment from the Representative was to satisfy the Settlement. Accordingly, as advised by our PRC Legal Advisor, we are entitled to receive payments made by the Payor and the Representative.

The third-party payment was a one-off arrangement solely for settling the relevant litigation with the Relevant Customer, and we have not allowed any other third-party payments during the Track Record Period and do not expect to allow any third-party payments in the ordinary course of business going forward.

SUPPLIERS

Our suppliers comprise original component manufacturers in China. Original components primarily refer to the key hardware components that directly form part of the core configuration of our solutions and are incorporated into the final deliverables to customers, such as SSDs, HDDs, memory modules, CPUs and NICs. The tables below set forth the basic information of our Group's top five suppliers in each year/period during the Track Record Period:

Year ended December 31, 2023

Suppliers	Products/ services purchased	Purchase amount	% of total purchases	Year of commencement of business relationship	Credit terms	Payment method
		(RMB'000)				
Supplier A[(1)]	Servers and hard drives	32,895	22.6	2018	60 days	Bank transfer
Supplier B[(2)]	Servers	16,623	11.4	2022	N/A	Bank transfer
Supplier C[(3)]	Servers and hard drives	6,639	4.6	2021	60 days	Bank transfer
Supplier D[(4)]	Servers	5,899	4.0	2021	60 days	Bank transfer
Supplier E[(5)]	Servers and hard drives	3,821	2.6	2022	N/A	Bank transfer
Total		65,877	45.2			

Notes:

(1) Supplier A is a company based in Beijing, China, primarily engaged in the research, development, manufacturing and sales of information technology equipment. It is a subsidiary of a company listed on the Shenzhen Stock Exchange.

(2) Supplier B is a company based in Guangdong, China, primarily engaged in providing information and communication technology solutions.

(3) Supplier C is a company based in Guangdong, China, primarily engaged in the research, development, manufacturing and sales of servers. It is a company listed on the NEEQ.

(4) Supplier D is a company based in Beijing, China, primarily engaged in the research, development, manufacturing and sales of information technology equipment. It is a Company listed on the Shenzhen Stock Exchange.

(5) Supplier E is a company based in Jiangsu, China, primarily engaged in the wholesale and retail of electronic hardware.

Year ended December 31, 2024

Suppliers	Products/services purchased	Purchase amount	% of total purchases	Year of commencement of business relationship	Credit terms	Payment method
		(RMB'000)				
Supplier C	Servers and hard drives	15,904	21.0	2021	60 days	Bank transfer
Supplier D	Servers and hard drives	11,529	15.2	2021	60 days	Bank transfer
Supplier A	Servers and hard drives	8,265	10.9	2018	60 days	Bank transfer
Supplier F[(1)] . . .	Hard drives	4,086	5.4	2024	N/A	Bank transfer
Supplier G[(2)] . . .	Servers and hard drives	2,714	3.6	2023	60 days	Bank transfer
Total		42,498	56.1			

Notes:

(1) Supplier F is a company based in Shanghai, China, primarily engaged in the research, development and sales of information technology equipment.

(2) Supplier G is a company based in Beijing, China, primarily engaged in the research, development and sales of information technology equipment. It is a subsidiary of a company listed on the Shenzhen Stock Exchange.

Year ended December 31, 2025

Suppliers	Products/services purchased	Purchase amount	% of total purchases	Year of commencement of business relationship	Credit terms	Payment method
		(RMB'000)				
Supplier C	Servers and hard drives	45,378	35.5	2021	60 days	Bank transfer
Supplier D	Servers and hard drives	11,435	8.9	2021	60 days	Bank transfer
Supplier H[(1)] . . .	Servers and hard drives	8,355	6.5	2025	N/A	Bank transfer
Supplier A	Servers and hard drives	8,088	6.3	2018	60 days	Bank transfer
Supplier I[(2)] . . .	Servers and hard drives	7,900	6.2	2025	90 days	Bank acceptance bill
Total		81,156	63.5			

Notes:

(1) Supplier H is a company based in Guangdong, China, primarily engaged in the research, development and production of storage devices and server-related hardware.

(2) Supplier I is a company based in Fujian, China, primarily engaged in providing IT information system services.

Six months ended June 30, 2026

Suppliers	Products/services purchased	Purchase amount	% of total purchases	Year of commencement of business relationship	Credit terms	Payment method
		(RMB'000)				
Supplier C	Servers and hard drives	37,955	22.5	2021	60 days	Bank transfer
Supplier J[(1)] . . .	Servers and hard drives	23,231	13.8	2025	60 days	Bank transfer
Supplier D	Servers and hard drives	16,259	9.7	2021	60 days	Bank transfer
Supplier K[(2)] . . .	Hard drives	9,176	5.4	2026	N/A	Bank transfer
Supplier L[(3)] . . .	Hard drives	8,229	4.9	2026	0 to 30 days	Bank transfer
Total		94,850	56.3			

Notes:

(1) Supplier J is a company based in Beijing, China, primarily engaged in the research, development, manufacturing and sales of servers, computer hardware and information technology equipment. It is a subsidiary of a company listed on the Shenzhen Stock Exchange.

(2) Supplier K is a company based in Beijing, China, primarily engaged in the research, development, design and sales of solid-state drives and the provision of related storage solutions.

(3) Supplier L is a company based in Beijing, China, primarily engaged in the research, development and sales of solid-state drives, embedded storage products and integrated circuits.

As of the Latest Practicable Date, none of our Directors, their respective close associates or any of our shareholders (who, to the knowledge of the Directors, owned more than 5% of our issued share capital) had any interest in any of our five largest suppliers in each year/period during the Track Record Period.

Procurement

We procure software, hardware and services from qualified suppliers to maintain product quality and to optimize our cost structure, which are the key raw materials and components used in our turnkey appliance hardware. Our procurement scope covers (i) the key raw materials and components of our turnkey appliance, including SSDs, memory modules, HDDs, and network interface cards (NICs), (ii) hardware warranty extensions, and (iii) outsourced services such as data labeling, solutions testing and on-site deployment. We collaborate with multiple suppliers for each major product to ensure continuity of supply. We are actively working with more suppliers of our key software, hardware and services to further diversify our supplier base and shore up the stability, timeliness and quality of our procurement. During the Track Record Period and up to the Latest Practicable Date, we did not experience shortages of software, hardware or services that materially affected our operations.

We source performance-critical components (for example SSDs, memory, NICs) directly from original manufacturers or first-tier channels under framework arrangements that combine pricing, supply assurance and technical support. Non-core parts (such as chassis, motherboards, power supplies and cabling) are generally integrated by server OEM/ODM partners. For on-site deployment, which typically involves installation, environment integration and acceptance testing, we are capable of delivery in-house. For standardized, low-customization projects we may engage qualified third parties given the labor-intensive nature of these tasks. In all cases, our staff supervise and manage the on-site process to ensure quality and schedule control.

The design, assembly and production of our hardware products (including turnkey appliances) are conducted under a hybrid model combining OEM server manufacturers and our own facilities. We specify hardware configuration requirements, embedded software, production standards and testing criteria, and engage OEM server manufacturers to produce customized or semi-finished servers based on such specifications. In parallel, we also procure components at the part level and perform in-house assembly, software installation, integration and testing at our own facilities.

We maintain stringent supplier selection and management procedures. Supplier qualification is assessed based on factors including technical capability, production capacity, quality standards, certifications, compliance, financial strength, delivery performance, business reputation and alignment with our strategic needs. Potential suppliers must undergo screening, qualification reviews, testing and audits before being admitted to our qualified supplier list. Suppliers are further classified into different categories based on their qualifications and performance and are subject to ongoing performance evaluations covering commercial management, product quality, delivery fulfilment and service support. We regularly review and update supplier classifications to ensure continued compliance with our procurement and supply chain management requirements.

To support supply assurance, we maintain frequent communication with suppliers on lead times and price trends, coordinate internally between sales and operations for rolling demand forecasts and issue advance notices for components that may face extended lead times. This diversified and balanced sourcing strategy has enabled us to avoid material delivery disruptions caused by capacity constraints. We have implemented various measures to mitigate the impact of raw material price fluctuations. First, we enter into framework agreements with core suppliers. Through long-term cooperation and volume-based procurement, we are able to secure relatively stable pricing arrangements. Second, for certain key components, we adopt an authorized vendor and pricing arrangement, under which original manufacturers authorize OEM suppliers to supply such components to us at pre-agreed prices. This mechanism helps reduce procurement costs and mitigate short-term price volatility. Third, we generally maintain multiple qualified suppliers for key components. This multi-supplier strategy allows us to conduct competitive sourcing and flexibly select suppliers, thereby diversifying procurement risks and reducing reliance on any single supplier. During the Track Record Period and up to the Latest Practicable Date, while the prices of storage devices increased due to the increase in market prices of storage chips, we did not experience any other major price fluctuations, nor did we experience any overstock, delays or shortages in the supply of performance-critical components.

Salient Terms of Agreements with Suppliers

The salient terms of our standard purchase contract with component manufacturers during the Track Record Period are set out below:

- ***Specifications.*** The agreement specifies the product name, specifications, configuration, price, quantity and other detailed requirements.

- ***Delivery.*** The supplier is responsible for delivering the goods in accordance with the agreed schedule, delivery location and method, and shall bear all shipping costs and in-transit risks.

- ***Payment.*** We typically make a one-time payment for the amount stipulated in the procurement contract and are generally granted a credit period of up to 60 days.

- ***Transfer of Risk.*** The supplier assumes all risks associated with the goods until delivery and successful acceptance by us. Upon delivery and acceptance, ownership and associated risks transfer to us.

- ***Quality Guarantee.*** We reserve the right to inspect the procured goods within a specified period after receipt. If the goods fail to meet the agreed requirements, we may reject them and require the supplier to provide replacements, with the supplier bearing any resulting losses.

- ***Termination.*** We reserve the right to terminate the agreement if the supplier fails to deliver the goods on time and the delay becomes excessive. Conversely, if we fail to make payment for an extended period, the supplier may terminate the agreement.

OVERLAPPING OF CUSTOMERS AND SUPPLIERS

Customer B, one of our five largest customers in 2023, 2024, 2025 and the six months ended June 30, 2026, was also a supplier during the Track Record Period. Customer B contributed to 14.4%, 10.1%, 5.5% and 9.2% of our revenue in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and 1.4%, 1.5%, 0.8% and 0.2% of our purchase amounts in the same year/period, respectively. We sold data lake storage solution, storage solution for T&I and AI storage services to Customer B during the Track Record Period. As a separate and independent matter, we procured computing software from Customer B during the Track Record Period.

Customer C, one of our five largest customers in 2023 and the six months ended June 30, 2026 was also a supplier during the Track Record Period. Customer C contributed to 7.9%, 0.6%, 3.7% and 7.4% of our revenue in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and nil, 0.9%, nil and 0.0% of our purchase amounts in the same year/period, respectively. We sold data lake storage solution, storage solution for T&I and AI storage services to Customer C during the Track Record Period. As a separate and independent matter, we procured on-site operation and maintenance support services from Customer C in 2024 and the six months ended June 30, 2026, including routine inspections, troubleshooting, emergency response and real-time technical support, primarily because Customer C is able to provide on-site operation and maintenance services that meet specific requirements of underlying end customers, thereby improving service responsiveness and contributing to overall customer satisfaction.

Customer D, one of our five largest customers in 2023, 2024 and 2025 was also a supplier in 2023 and 2024. Customer D contributed to 7.4%, 5.8%, 4.1% and 0.5% of our revenue in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and 0.0%, nil, nil and nil of our purchase amounts in the same year/period, respectively. We sold data lake storage solution, storage solution for T&I and AI storage services to Customer D during the Track Record Period. As a separate and independent matter, we procured laptops from Customer D in 2023.

Customer F, one of our five largest customers in 2024 and 2025 was also a supplier in 2023 and 2024. Customer F contributed to 1.3%, 4.6%, 3.8% and 0.4% of our revenue in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and 0.1%, 0.0%, nil and nil of our purchase amounts in the same year/period, respectively. We sold data lake storage solution, storage solution for T&I and AI storage services to Customer F during the Track Record Period. As a separate and independent matter, we procured marketing and promotional event services from Customer F in 2023 and 2024.

Supplier A, one of our five largest suppliers in 2023 and 2024 was also a customer in 2023, 2024, 2025 and the six months ended June 30, 2026. Supplier A contributed to 22.6%, 10.9%, 6.3% and 4.0% of our purchase amounts in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and 0.1%, 0.2%, 0.6% and 0.3% of our revenue in the same year/period, respectively. We procure servers and hard drives from Supplier A during the Track Record Period. As a separate and independent matter, we sold data lake storage solution and storage solution for T&I to Supplier A during the Track Record Period.

Supplier D, one of our five largest suppliers in 2023, 2024, 2025 and the six months ended June 30, 2026, was also a customer in 2023, 2024, 2025 and the six months ended June 30, 2026. Supplier D contributed to 4.0%, 15.2%, 8.9% and 9.7% of our purchase amounts in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and 0.7%, 0.9%, 0.9% and 0.1% of our revenue in the same year/period, respectively. We procure servers and hard drives from Supplier D during the Track Record Period. As a separate and independent matter, we also sold data lake storage solution and storage solution for T&I to Supplier D during the Track Record Period.

Supplier J, one of our five largest suppliers in the six months ended June 30, 2026, was also a customer in 2023, 2024, 2025 and the six months ended June 30, 2026. Supplier J contributed to 0.1%, 0.1%, 5.6% and 13.8% of our purchase amounts in 2023, 2024, 2025 and the six months ended June 30, 2026, respectively, and 0.3%, 1.7%, 0.3% and 0.8% of our revenue in the same year/period, respectively. We procure [●] from Supplier J during the Track Record Period. As a separate and independent matter, we also sold [●] to Supplier J during the Track Record Period.

The reason for the above overlapping of customers and suppliers is largely because certain of our customers are distributors with broad product coverage. In addition to distributing our data lake storage solutions and storage solutions for T&I, these distributors also supply technology products such as laptops and computing software that we procure for our daily operations. As advised by CIC, such arrangements are common in our sector, where distributors often engage in both the procurement and resale of a wide range of technology products. All of our sales and purchases from Customer B, Customer C, Customer D, Customer F, Supplier A, Supplier D and Supplier J were conducted in the ordinary course of business under normal commercial terms and on arm's length basis.

QUALITY CONTROL

We are committed to maintaining the highest level of quality in our solutions. We have designed and implemented a quality management system that provides the framework for continuous improvement of products and processes. We have also implemented a management review control process to conduct regular systematic reviews of our quality management system, in order to closely monitor the implementation of our quality management system. During the Track Record Period and up to the Latest Practicable Date, we had not been involved in any material complaints, product liability claims or product returns in relation to our solutions.

COMPETITION

According to CIC, the market size of China's AI infrastructure reached RMB381.3 billion in 2025, accounting for approximately 15% of the global AI infrastructure market. It is expected to expand to RMB1,528.6 billion by 2030, representing a CAGR of 32.0%, at which point China's share of the global market is projected to rise to approximately 25%. Within this market, the size of on-premise AI infrastructure is expected to grow from RMB192.2 billion in 2025 to RMB802.5 billion in 2030, representing a CAGR of 33.1%.

The market size of on-premise distributed AI storage in China, measured by revenue, reached RMB17.1 billion in 2025 and is expected to grow to RMB90.8 billion by 2030. The installed capacity of distributed AI storage in China reached 17.5 EB in 2025 and is expected to grow to 89.1 EB by 2030, representing a CAGR of 38.5%. In 2025, AI data lake storage accounted for approximately 78.5% of the market in terms of installed capacity, while AI training and inference storage represented 21.5%. By 2030, these proportions are projected to shift to approximately 72% and 28%, respectively.

Based on distributed AI storage installed capacity, the top five providers in China collectively held a 47.9% market share in 2025. Although we believe that our technological capabilities position us competitively in this rapidly expanding market, we may face competition from established industry players with more substantial R&D resources and marketing capabilities. See "Risks Factors — Risks Relating to the R&D of Our Solutions — The industry in which we operate is

characterized by constant development. If we fail to continuously innovate our technology and provide innovative solutions that meet the expectations of our customers, our business, financial condition and results of operations may be materially and adversely affected." We are China's second largest distributed AI storage provider and the largest independent provider, based on distributed AI storage installed capacity in 2025.

The top five distributed AI storage solution providers in China accounted for a combined market share of 47.9% by installed capacity in 2025, and we ranked as China's second-largest distributed AI storage solution provider and the largest independent solution provider based on installed capacity in 2025, with a market share of 10.6%. See "Industry Overview — Competitive Landscape" for details.

EMPLOYEES

We believe that our workforce is the driving force of our long-term growth. As of June 30, 2026, we had 328 full-time employees. The following table sets forth the number of our employees by function as of the dates indicated:

	As of December 31,						As of June 30,	
	2023		2024		2025		2026	
Employee Function	Number of employees	% of total	Number of employees	% of total	Number of employees	% of total	Number of employees	% of total
Sales and Marketing	162	45.1	150	47.5	152	48.4	162	49.4
Research and Development	156	43.5	133	42.1	129	41.1	132	40.2
General and Administration	29	8.1	26	8.2	26	8.3	33	10.1
Supply Chain Management	12	3.3	7	2.2	7	2.2	1	0.3
Total	359	100.0	316	100.0	314	100.0	328	100.0

Our success depends on our ability to attract, retain and motivate qualified personnel, and we believe that our talent pool is one of our core strengths. We adopt high standards and strict procedures in our recruitment to ensure the quality of new hiring and use various methods for our recruitment, including campus recruitment, online recruitment and internal recommendation to satisfy our demands for different types of talents. We enter into employment contracts and agreements regarding confidentiality, intellectual property and non-competition with our executive officers, managers and employees. In addition, we usually enter into proprietary information and inventions agreement with our core employees, under which we have all right, title and interest relating to any and all inventions by such employee during the term of his/her employment. Further, when employees are hired, we give them an employee handbook, which informs them of our policies and their rights in all material respects, from recruitment, compliance, salary, benefits, performance assessment to training and development.

We provide our employees with competitive remuneration and benefits. We place significant emphasis on investing in our employees and have established a well-rounded talent development system. Our training programs are categorized into company-level, department- level and function-level training. Prior to commencing their roles, new employees are required to complete relevant training and pass examinations. We offer a wide range of specialized training aimed at enhancing the skills of our employees. In addition, we have cultivated a number of internal training courses and developed a series of targeted courses to effectively implement our talent development strategy, foster the growth of key talents and enhance the managerial proficiency of our team. Our employees are currently represented by our internal labor union. We maintained a good relationship with our employees and did not have any material labor dispute during the Track Record Period and up to the Latest Practicable Date.

As required under PRC regulations, we participate in various employee social security plans that are organized by applicable local municipal and provincial governments, including housing, pension, medical, work-related injury, maternity and unemployment benefit plans. During the Track Record Period, we engaged third-party agencies to pay social insurance premium and housing provident funds for certain of our employees. For details, see "Risk Factors — Risks Relating to Our General Operation — We have relied on third-party agencies to handle certain employee benefit contributions, and any deficiencies in their practices may result in liabilities for us."

INSURANCE

We maintain insurance policies that are required under PRC laws and regulations, and based on our assessment of our operational needs and industry practice. We maintain insurance coverage including property insurance and all-risk cargo insurance. We believe that the amount of our insurance coverage is in line with the customary standard in the industry and is adequate for our operations. During the Track Record Period and up to the Latest Practicable Date, we did not make any material insurance claims in relation to our business. See "Risk Factors — Risks Relating to Our General Operations — Our insurance coverage may be insufficient to cover all of our potential losses."

INFORMATION TECHNOLOGY SYSTEM

IT is fundamental to our competitive edge and operational efficiency. We primarily utilize our OA system that evolves in tandem with our business growth, ensuring they meet our varied operational demands. Our OA system underpins key areas such as sales management, supply chain management, customer management, employee management, financial management and project management.

During the Track Record Period and up to the Latest Practicable Date, we had not experienced any material IT system failure or downtime that had a material adverse effect on our business operations.

ENVIRONMENTAL, SOCIAL AND GOVERNANCE

We believe that integrating Environmental, Social, and Governance (ESG) factors into our decision-making processes not only enhances our overall business performance but also creates long-term value for our stakeholders. In line with our commitment to sustainable development and responsible business practices, we recognize the importance of ESG considerations as fundamental components of our corporate strategy.

ESG Governance Framework

The governance of ESG matters is an integral part of our overall corporate governance framework, which also includes the unified management of climate-related considerations. Our Board is responsible for overseeing the Company's ESG strategies and initiatives, and takes the responsibility of:

(i) setting the direction of our sustainability efforts and ensuring that these initiatives align with the Company's long-term business objectives;

(ii) management of material ESG-related issues;

(iii) developing, monitoring and adjusting ESG strategies and objectives;

(iv) establishing ESG framework, policies;

(v) guiding ESG key performance indicators and overseeing the progress of ESG initiatives; and

(vi) reviewing and approving ESG reports.

Our Board comprises directors with diverse backgrounds in industry operations, finance, risk management and corporate governance, and is strengthening its capability to oversee ESG matters through ongoing measures. The Company has engaged external consultants to support ESG risk assessments, materiality reviews and climate related issues. Senior executives are accountable for delivering ESG initiatives and presently report ESG matters directly to the Board.

To facilitate effective oversight, we plan to establish an ESG Committee upon [**REDACTED**], which will be responsible for implementing ESG work plan, developing ESG policies and monitoring performance against established targets. Once established, the ESG Committee will assume a reporting and oversight role. Our Board also actively engages with key stakeholders to ensure that their views are considered in the Company's ESG strategies. Our Board plans to review ESG progress at least annually after [**REDACTED**] and to approve externally facing ESG disclosures prior to publication.

We have established measurable ESG goals, targets and policies aligned with recognized international frameworks and industry practice. These targets are set against baseline data, time-bound milestones and clear ownership, and are monitored regularly by management and the ESG Committee.

Our ESG policies cover multiple topics and are developed with reference to ISSB and relevant ISO standards. Where appropriate, we benchmark our performance and targets against industry peers and sectoral norms to ensure comparability and competitiveness.

Material ESG Issues

We recognize the importance of managing material ESG issues and have completed a comprehensive materiality assessment using a structured identify, prioritize and address process. Issues are assessed on their impacts on business performance and stakeholder interests and prioritized under a double materiality approach. Assessment criteria include severity of impact, likelihood of occurrence, scope of affected business units and supply-chain exposure, regulatory and compliance risk, market trends, financial exposure, and alignment with corporate strategy. We draw on internal operational and financial condition, stakeholder engagement, industry benchmarking and external advisor input.

We have established and implemented a Risk Management Framework (《風險管理規程》), which is designed to identify, analyze, and address risks, including those related to ESG and climate change. This comprehensive approach enables us to manage risk in a structured and systematic manner.

Material issues	Potential risks and impacts	Actions
GHG emissions	Increased operating expenses, enlarged carbon footprint, and regulatory compliance burdens	– Measure and report GHG emissions – Establish targets and reporting cadence
Climate-related risk and opportunity management .	Service disruptions, revenue loss, higher capex for resilience, asset impairment risk, and customer shifts toward low-carbon providers	– Integrate ESG risks into Company risk management – Prioritize using clean energy and green products

Material issues	Potential risks and impacts	Actions
Information and data security	Legal penalties, customer churn and remediation costs, and long-term reputational damage	– Enforce least-privilege access and continuous monitoring/logging – Run penetration tests and vulnerability management
Supply chain management	Disrupted deliveries, increased procurement costs and working capital needs, and impair service availability	– Implement supplier due diligence, ESG screening and risk-based segmentation – Secure multi-source procurement, long-term agreements and strategic spares inventory – Plan to include ESG and quality clauses in procurement contracts
Human resources	Reduced service quality, increased hiring/training costs, constrained growth and operational delivery	– Implement competitive compensation and multiple welfare – Provide continuous technical and security training and certification
Innovation	Worse competitiveness, higher operating/environmental costs	– Continue investing in R&D – Prioritize differentiated high-performance, low-carbon product tiers

Climate-Related Risks and Opportunities

Climate-Related Risks

Physical Risks. Our operational facilities, as well as those of our partners, may face threats from extreme weather events, such as typhoons, heavy rain, and floods, which can disrupt logistics and damage critical infrastructure, ultimately impacting our service delivery. In response to these risks, we have implemented a comprehensive environmental management system, developed contingency plans to effectively address climate-related emergencies.

Transition Risks. As the shift toward a lower-carbon economy unfolds, our business operations, market demand and regulatory environment may be impacted to varying degrees.

- **Regulatory Changes** — we may face compliance costs and operational adjustments to align with more stringent climate policies and emissions regulations.
- **Technological risk** — we may have to incur significant costs in transitioning to lower-emission technologies.
- **Market and reputation risk** — growing awareness and concern for the environment among clients can shift demand toward more sustainable and eco-friendly products.

To mitigate these risks, we are closely monitoring relevant policies and regulations and are prepared to revise our current ESG measures to ensure compliance in the event of regulatory changes. We carefully select qualified suppliers to minimize climate-related risks in our supply chain. We are also actively investing in R&D for innovative solutions and engaging with our customers to ensure our offerings align with market trends and preferences.

Climate-Related Opportunities

We have identified climate-related opportunities to enhance the operation of our business. By investing in R&D for green technologies, we can create advanced products that emphasize energy efficiency and low carbon emissions, aligning with regulatory requirements and attracting environmentally conscious clients. By committing to ESG practices, we can enhance our brand reputation and build customer stickiness, positioning ourselves as a leader in the sustainable market.

Environmental Protection

We strictly adhere to the PRC Environmental Protection Law and other applicable laws and regulations and have established internal processes to ensure the identification of legal requirements and compliance evaluation. Additionally, we have achieved ISO 14001 Environmental Management System certification, which ensures that our operations consistently meet high standards for environmental protection.

Greenhouse Gas Emissions

We are committed to reducing our carbon footprint and have made substantial progress in this area, including the decision to eliminate all company vehicles by 2025 and to encourage the use of public transportation. The following table presents our GHG emissions for the years/periods indicated:

		Year ended December 31,			Six months ended June 30,
Scope of GHG emissions	**Unit**	**2023**	**2024**	**2025**	**2026**
Scope 1 emissions[(1)] . .	tCO2e	2.4	1.1	0.6	0.5
Scope 2 emissions[(2)] . .	tCO2e	1,227.3	1,348.4	911.2	361.3
Total	tCO2e	1,229.7	1,349.5	911.8	361.7

Notes:

(1) Scope 1 GHG emissions cover direct emissions from the Company's business fleet (operations ceased in mid-2024) and emissions from processing of fresh and wastewater; and

(2) Scope 2 GHG emissions cover indirect emissions from purchased electricity consumption.

We have begun identifying and accounting for Scope 3 emissions. The primary sources are Purchased Goods and Services, Upstream Transportation and Distribution, and Business Travel. Based on DEFRA 2025 emission factors and current data availability, our six months ended June 30, 2026 Scope 3 emissions are estimated at approximately 100.7 tonnes; this estimate presently covers business travel by air and rail only.

In alignment with the 1.5°C temperature goals set out in the Paris Agreement and the global transition toward carbon neutrality, we have established emissions management objectives. Using 2024 data as the baseline, we aim to reduce our GHG emissions intensity by 10% by the year 2030.

Resource Consumption

Since our business does not involve the production of hardware or electronic devices, our resource consumption primarily consists of municipal electricity and domestic water. The following table presents our resource consumption for the years/periods indicated:

		Year ended December 31,			Six months ended June 30,
Resource Consumption	**Unit**	**2023**	**2024**	**2025**	
Purchased electricity	kWh	2,287,176	2,512,819	1,698,062	673,304
Water	t	2,544	2,344	1,694	1,240

Our electricity consumption comes from both office and data center operations, while water consumption is limited to office premises. Due to adjustments in office space and implementation of resource saving measures, our overall resource consumption declined significantly in 2025, while increased use of the employee gym led to a rise in water consumption in 2026.

Responsible resource stewardship forms a fundamental pillar of our environmental management strategy. Our efforts focus on the sustainable use of electricity and water, the reduction of paper consumption, and a systematic approach to waste management. Using 2024 data as the baseline, we aim to reduce the intensity of municipal electricity and domestic water consumption by 10% by the year 2030.

To achieve these goals, we have implemented specific measures focused on electricity conservation water saving and paper reduction. Besides, energy-efficient bulbs are utilized for lighting, and lights are turned off in unoccupied areas, with air conditioning permitted only when indoor temperatures exceed 25°C.

Waste Management

Because our business model does not involve manufacturing or production activities, we do not generate production related hazardous or non hazardous waste. Our company vehicle use was minimal, and the fleet was discontinued in mid 2024. Accordingly, the Company does not produce significant operational waste, nor does it generate additional process related air emissions. Office and IT equipment at end of life are disposed of or recycled in accordance with local regulations. If material waste streams arise in the future, we will quantify and disclose the relevant metrics in our ESG reporting.

Social Matters

Our approach to social governance is founded on a framework that prioritizes ethical leadership, responsibility, and transparency. We aim to extend our positive influence beyond our immediate operations by actively engaging with the broader community and contributing to the enhancement of overall social well-being.

Anti-Corruption

We prohibit any form of bribery, acceptance of bribes or unfair competition, and require employees to politely decline gifts from suppliers, partners, or other third parties. Employees must disclose any personal relationships that could create conflicts of interest and inform us of any outside employment or business activities. Violations of these policies will be treated as serious misconduct, and we reserve the right to terminate employment contracts when necessary.

We encourage employees and relevant stakeholders to report any suspected misconduct and are committed to protecting whistleblowers. Any retaliation against individuals who report concerns in good faith is considered a serious violation of company policies.

Supply Chain Management

Our supply chain management is governed by the Procurement Management System (《採購管理制度》), which establishes clear guidelines for personnel responsibilities, procurement processes, and supplier management. We conduct regular assessments of our suppliers to ensure they meet our quality and performance standards. The Internal Audit Department is responsible for overseeing the procurement activities through process supervision and post-activity audits.

We encourage our suppliers to adhere to high standards of ethical conduct and to comply with relevant environmental, labor and safety requirements. We plan to gradually incorporate ESG performance into our supplier evaluation system, prioritizing vendors with outstanding ESG metrics under equivalent conditions.

Information and Data Safety

We have obtained ISO 20000 certification for our IT Service Management System and ISO 27001 certification for our Information Security Management System. In accordance with our information security policies, we classify and categorize data based on its sensitivity and establish core data access permissions according to job roles, all while strictly adhering to the principle of least privilege. Our protection measures include data encryption, access control, and log auditing.

Human Resources and Innovation

Employment. We strictly adhere to all relevant labor laws and regulations with no employment of child labor or forced labor. In addition to mandatory social insurance, we offer a range of benefits, including support and care for birthdays, weddings, the birth of children, and during periods of illness. Dismissals are conducted in accordance with applicable laws, following fair and documented procedures, including advance notice and provision of statutory entitlements. All employment policies and procedures are documented in the employee handbook.

We invest significant resources in employee development and provide regular, structured training programs tailored to different roles and career stages, with the aim of enhancing skills, technical capabilities, and management competence. We have established an internal information network that facilitates knowledge exchange among employees.

Diversity. We are committed to providing equal employment opportunities to all employees and candidates, regardless of race, gender, religion, age, nationality, social background or other characteristics. We strive to cultivate an inclusive workplace culture that enables individuals to develop rewarding careers and contribute to the Company's growth based on merit and capability.

Occupational health. We prioritize occupational health and workplace safety and strictly comply with applicable occupational health and safety laws and regulations. We have established and enhanced our safety management system and the safety responsibility framework for all personnel and functional departments. We have achieved ISO 45001 certification for our Occupational Health and Safety Management System. During the Track Record Period, and up to the Latest Practicable Date, we did not record any work-place-related accidents or incur any work-related fatalities.

Social Responsibility

Since 2021, we have established a comprehensive strategic partnership with Southwest Jiaotong University, which has led to the creation of the XSKY Elite Class Program. As of the Latest Practicable Date, we have invested RMB300,000 in these initiatives, supporting the cultivation of professionals in the field of data storage.

We actively engage in initiatives that promote rural economic development and support agricultural communities. In collaboration with a working group of Haidian District, we donated books to the Huangyangwa School in Inner Mongolia's Aohan Banner. Additionally, we sponsored the 2024 China Rural Revitalization Entrepreneur Support Program, aimed at promoting local agricultural products and providing support to rural entrepreneurs.

BUSINESS SUSTAINABILITY

Historically, we incurred net losses primarily because we prioritized technology development and market expansion. We made substantial investments in research and development to strengthen our leadership in AI storage, expand our product portfolio and enhance our underlying technologies. Our research and development expenses amounted to RMB110.0 million, RMB88.0 million, RMB77.2 million, RMB38.0 million and RMB41.4 million in 2023, 2024, 2025 and the six months ended June 30, 2025 and 2026, respectively. We also incurred significant selling and marketing expenses to expand our customer base and deepen industry penetration, which amounted to RMB124.1 million, RMB86.5 million, RMB73.6 million, RMB36.2 million and RMB36.9 million in 2023, 2024, 2025 and the six months ended June 30, 2025 and 2026, respectively. In addition, our administrative expenses amounted to RMB56.2 million, RMB32.1 million, RMB24.9 million, RMB11.0 million and RMB13.5 million in the same years/periods. These investments were made while our revenue base was still developing and had not yet been sufficient to cover our operating expenses, resulting in net losses of RMB180.7 million in 2023, RMB84.2 million in 2024 and RMB10.5 million in the six months ended June 30, 2025.

We achieved a turnaround in 2025 and recorded profit for the year of RMB7.1 million in 2025 and RMB0.3 million in the six months ended June 30, 2026, primarily as a result of improvements in our operating performance. This was primarily attributable to the continued commercialization of our data lake storage solutions and storage solutions for T&I, which drove our revenue to increase by 54.8% from RMB172.5 million in 2024 to RMB267.0 million in 2025, and increased from RMB99.5 million in the six months ended June 30, 2025 to RMB185.8 million in the six months ended June 30, 2026. Our gross profit increased by 55.1% from RMB109.9 million in 2024 to RMB170.4 million in 2025, and increased by 38.0% from RMB69.8 million in the six months ended June 30, 2025 to RMB96.3 million in the six months ended June 30, 2026. The increase in profitability also reflected improved operating efficiency and greater economies of scale. As our solutions became more mature and our industry adaptation and unified data foundation capabilities strengthened, we were able to improve scalability, reduce technology migration and deployment costs and enhance delivery efficiency. This is also reflected in the decline of our research and development expenses, selling and marketing expenses and administrative expenses as a percentage of revenue from 65.9%, 74.4% and 33.7% in 2023, respectively, to 28.9%, 27.6% and 9.3% in 2025, respectively. Our research and development expenses, selling and marketing expenses and administrative expenses as a percentage of revenue decreased from 38.2%, 36.3% and 11.1% in the six months ended June 30, 2025, respectively, to 22.3%, 19.8% and 7.3% in the six months ended June 30, 2026, respectively.

Going forward, we aim to maintain profitability through a combination of revenue growth, continued operating discipline and greater recurring revenue contribution. We plan to accelerate commercialization of our AIMesh by expanding deployment in large-scale AI storage scenarios and deepening collaboration with ecosystem partners. We plan to continue enhancing our overall product features through ongoing R&D investment, focusing on higher-performance distributed storage, broader protocol compatibility and improved integration with diverse AI compute

environments. At the same time, we aim to acquire new customers across additional verticals while strengthening relationships with existing clients through superior delivery, service quality and long-term support. To sustain scalable growth, we also plan to enhance operating efficiency through optimized resource allocation, refined implementation processes and improved cost control.

- **Revenue Quality Enhancement:** On the revenue side, we intend to place greater emphasis on the quality and sustainability of our revenue as our business continues to grow. We expect demand for our data lake storage solutions and storage solutions for T&I to remain supported by ongoing enterprise adoption of AI workloads and increasing data processing requirements. Our solutions are typically deployed as part of customers' core AI data infrastructure and are integrated into their key operational workflows, which enhances customer stickiness and supports repeat purchases as customers expand their usage over time. We also offer annual maintenance services and value-added offerings such as capacity expansion. As a result, we expect our revenue to sustain steady growth as our customer base deepens and deployment scale expands, which is expected to support improved revenue visibility over time.

- **Cost Discipline and Operating Leverage:** We will continue to focus on disciplined cost management and improving operational efficiency across our business. We aim to enhance coordination across product development, solution design and delivery functions to streamline implementation processes and improve execution efficiency. We also plan to optimize procurement and supply chain arrangements through more standardized configurations and scale-driven purchasing, thereby improving cost visibility and reducing volatility. As our revenue base expands, we expect to achieve greater operating leverage while maintaining prudent control over operating expenses. We will continue to invest in R&D to support product innovation, while improving investment efficiency through the reuse of core technologies and architectures. Selling and marketing expenses are expected to benefit from increasing brand recognition and a more focused sales approach, while administrative expenses are expected to scale more efficiently as internal processes mature, supporting gradual improvements in operating margins.

PROPERTIES

Our corporate headquarters is located in Beijing, China. As of the Latest Practicable Date, we did not have any self-owned property, and leased seven properties in the PRC with an aggregate gross floor area of 4,343.15 sq.m. Our leased properties are primarily used for the purpose of office and R&D.

As of the Latest Practicable Date, none of the properties leased by us had a carrying amount of 15% or more of our consolidated total assets. According to Chapter 5 of the Hong Kong Listing Rules and section 6(2) of the Companies (Exemption of Companies and Prospectuses from Compliance with Provisions) Notice, this document is exempt from the requirements of section 342(1)(b) of the Companies (Winding up and Miscellaneous Provisions) Ordinance to include all interests in land or buildings in a valuation report.

As of the Latest Practicable Date, we have not completed the lease registration procedures for seven leased properties in China. According to relevant PRC regulations, the relevant government authorities may require us to complete the registration within a prescribed period, and failure to do so may result in a fine of RMB1,000 to RMB10,000 for each unregistered lease agreement. See "Risk Factors — Risks Relating to Our General Operation — Deficiencies in some of our leased properties may adversely affect our business operations and financial condition."

LICENSES, APPROVALS AND PERMITS

We have obtained various permits, approvals and certificates for our business. As advised by our PRC Legal Advisor, we had obtained the requisite approvals and permits from, and have made all the filings with the applicable authorities which were material to our operations, and such approvals and permits are valid and subsisting as of the Latest Practicable Date. In addition, our products and R&D procedures passed various industry-recognized certifications and tests for liability. The following table sets out the main standards, certifications or requirements that we were compliant with as of the Latest Practicable Date:

Standards, certifications or requirements	Publication date	Effective date	Definition of the standards, certifications or requirements	Granting authority	Our compliance with the standards, certifications or requirements
GB/T 44271-2024 Information Technology – Cloud Computing – General Technical Requirements of Edge Cloud	August 23, 2024	March 1, 2025	GB/T 44271-2024 standard specifies the overall architecture of edge cloud and specifies the technical requirements for edge cloud infrastructure, infrastructure services, platform services, and application services.	Subcommittee on Cloud Computing and Distributed Platforms, National Technical Committee for Information Technology Standardization	We participated in drafting this standard, and our cloud infrastructure meets the requirement.
GB/T 17966-2024 Information Technology – Microprocessor Systems – Floating-Point Arithmetic .	September 29, 2024	April 1, 2025	GB/T 17966-2024 standard adopts international standards from organizations such as ISO and IEC and specifies the fundamental definitions and technical requirements for floating-point arithmetic in microprocessor systems, including number formats, arithmetic operations, rounding modes, and exception handling.	National Information Technology Standardization Technical Committee	We participated in drafting this standard, and our microprocessor systems meet the requirement.
T/CESA 1312-2024 High Performance Computing-Technical Requirements for Distributed Storage System.	January 30, 2024	January 30, 2024	T/CESA 1312-2024 standard emphasizes the support and application of distributed storage systems in a broad range of high-performance computing (HPC) scenarios, including traditional HPC, with a focus on data storage, processing, and analytical capabilities. With increasing awareness of energy costs and environmental protection, the energy efficiency and sustainability of HPC distributed storage systems are also key areas of concern and research addressed in this document.	China Electronics Standardization Association	We participated in drafting this standard, meet the requirement in high-performance computing scenarios.

Standards, certifications or requirements	Publication date	Effective date	Definition of the standards, certifications or requirements	Granting authority	Our compliance with the standards, certifications or requirements
GB/T 46448-2025 Information Technology — Storage Networking — Requirements for Ethernet-Based Device Discovery .	October 5, 2025	May 1, 2026	GB/T 46448-2025 standard specifies the device discovery mechanism for Ethernet-based storage networks, including network architecture, procedures, interface protocols, messages and other related matters.	National Information Technology Standardization Technical Committee	We participated in drafting this standard, and our distributed storage systems meet this requirement.
T/GCC 7009-2025 Server Application Scenario Performance Test Methods — Hyper-Convergence . .	December 1, 2025	December 1, 2025	T/GCC 7009-2025 standard specifies the methods for testing server performance in hyper-converged scenarios and applies to the performance testing and evaluation of servers.	Global Computing Consortium	We participated in drafting this standard, meet the requirement in hyper-converged scenarios.

LEGAL PROCEEDINGS AND COMPLIANCE

Legal Proceedings

We may from time to time be subject to various legal or administrative claims and proceedings arising from the ordinary course of business. Litigation or any other legal or administrative proceeding, regardless of the outcome, is likely to result in substantial cost and diversion of our resources, including our management's time and attention. See "Risk Factors — Risks Relating to Our General Operations — We may from time to time become a party to litigation, other legal and contractual disputes, claims and administrative proceedings that may materially and adversely affect our business and reputation."

During the Track Record Period and up to the Latest Practicable Date, there were no legal proceedings pending or threatened against us or our Directors that could, individually or in the aggregate, have a material adverse effect on our business, financial condition and results of operations.

Compliance

During the Track Record Period and up to the Latest Practicable Date, we had not been and were not involved in any material non-compliance incidents, nor have we been subject to fines, enforcement actions, or other penalties for non-compliance that, individually or in the aggregate, could have a material adverse effect on our business, financial condition, and results of operations. As confirmed by our PRC Legal Advisor, we were in compliance with all applicable laws and regulations in all material respects in China during the Track Record Period and up to the Latest Practicable Date.

During the Track Record Period, we engaged third-party human resources agencies to pay social insurance premium and housing provident funds on our behalf for certain employees in 15 cities where we did not have local subsidiary or branch. Under the agreements entered into between the third-party agencies and us, the third-party agencies have the obligation to pay social insurance premiums and housing provident funds for our relevant employees. If local authorities in the future deem such arrangement non-compliant, or if the third-party agencies that we engage fail to meet certain obligations, such as paying social insurance and housing provident funds for our employees in line with PRC laws and regulations, we may need to take corrective measures. This could include paying late fees or penalties to address any non-compliance with our social insurance and housing

provident fund obligations. As of the Latest Practicable Date, we had not received any notices of rectification or penalty related to the engagement of third-party agencies to pay social insurance. To prevent the recurrence of similar non-compliances, we have implemented internal control measures, including, among others, monitor our ongoing compliance with the social insurance and housing provident fund contributions regulations and to ensure the timely implementation of any necessary remedial action.

As of the Latest Practicable Date, registration and filing procedures of seven leased properties used in our operations remain incomplete. These properties were used by us as offices. As advised by our PRC Legal Advisor, the non-registration and filing of the relevant property lease will not affect the validity of the lease contracts or our legal right to use the leased properties, however, relevant local housing authorities may require us to complete the filing within the prescribed timeframe, and we may be subject to penalties of RMB1,000 to RMB10,000 per lease as a result of delay in filing for each lease agreement of such properties, with an aggregate maximum penalty of RMB70,000. Our failure to complete the lease registration and filing was mainly due to our lessors being unwilling to cooperate with us to complete the lease registration and filing. As of the Latest Practicable Date, we had not received any registration request from or become subject to any such fine levied by the relevant government authorities. We will take all practicable and reasonable steps to ensure that the relevant leases are registered and continue to communicate with such lessors to seek their cooperation to complete the registration and filing process. We have enhanced our internal control measures with respect to our leased properties, including, among others, requiring the relevant employees to conduct the lease registration and filing in a timely manner and recording detailed information of the lease agreements.

Based on the foregoing, our Directors are of the view that the above-mentioned remedies and internal control measures our Group implemented are adequate and effective.

RISK MANAGEMENT AND INTERNAL CONTROL

We have established a risk management and internal control framework covering key aspects of our business operations, including operational risk, compliance risk, information security and data privacy risk, and intellectual property risk. We regularly review and enhance our policies and procedures to ensure their effectiveness and compliance with applicable laws and regulations. Our legal department is responsible for compliance management, contract review and monitoring regulatory developments, while dedicated personnel oversee information security and data protection measures. To monitor the effectiveness of our risk management and internal control systems, we have established an Audit Committee and an internal audit function, which are responsible for reviewing internal controls, identifying deficiencies and reporting significant matters to the Board on a timely basis.

TRADE RESTRICTIONS AND EXPORT CONTROLS

During the Track Record Period and up to the Latest Practicable Date, we did not engage in transactions with entities subject to blocking sanctions. Although we have engaged in transactions involving a limited number of customers listed on non-blocking sanctions lists, since such sanction lists did not restrict non-U.S. person from engaging in sales activities with entities designated on the lists, our sanctions legal advisor is of the view that our dealings did not pose any primary or secondary sanctions risks.

We had no direct exports to the United States during the Track Record Period and up to the Latest Practicable Date. Accordingly, U.S. tariffs on imports from the PRC did not have any direct impact on our sales. We did not identify any material indirect impact on demand from end customers (including sales by distributors to end customers) attributable to U.S. or PRC tariff measures during the Track Record Period and up to the Latest Practicable Date. We did not have any procurement attributable to tariff measures imposed by the U.S. or the Chinese Mainland during the Track Record Period and up to the Latest Practicable Date, as we did not import U.S.-origin hardware or

other tariff-sensitive goods. As a result, PRC retaliatory tariffs on U.S.-origin imports did not lead to any increase in our cost of sales, and our overall financial performance was not materially affected by such tariffs. Our sanctions legal advisor is of the view that the reciprocal tariffs imposed by U.S. and retaliatory measures imposed by China had limited impact on our procurement and sales activities during the Track Record Period and up to the Latest Practicable Date.

Based on the nature of our solutions, our self-developed software, technology and turnkey appliances were not subject to U.S. export control restrictions. Our products and technology are developed outside the U.S., do not contain U.S.-origin controlled content above the applicable de minimis thresholds, and do not fall within the scope of applicable foreign direct product rules. Although we had dealings with certain customers subject to U.S. export control restrictions or customers subject to "BIS 50% rule", since our self-developed products and technology were not subject to the EAR pursuant to the de minimis rule and applicable foreign-direct product rules and thus we did not violate the U.S. export control law. We have not experienced any material impact from export controls or other trade-related measures imposed by governmental authorities in the U.S. during the Track Record Period and up to the Latest Practicable Date, according to our sanctions legal advisor. Our Directors are of the view that, the applicable trade restrictions and export controls are not expected to have a material adverse impact on our Group's business and development plans, including any future business expansion outside the PRC. During the Track Record Period and up to the Latest Practicable Date, neither our Group nor our business activities has been subject to any sanctions.

We do not engage in any of the "covered activities" under the Outbound Investment Rule related to semiconductors and microelectronics, quantum information technologies or restricted AI activities. None of us or our subsidiaries is listed on any restricted lists under the Outbound Investment Rule which could potentially render us or any of our subsidiaries a "covered foreign person." In particular, although we engage in the development and sales of distributed AI storage solutions, we do not develop AI models or algorithms, nor do we provide AI model training or inference services. Furthermore, our products do not contain proprietary AI models or AI decision engines. Based on the foregoing, and as advised by our sanctions legal advisor, we were not a "covered foreign person" and we did not engage in activities covered by the Final Rule as of the Latest Practicable Date. Upon completion of the [**REDACTED**], U.S. persons are allowed to purchase our publicly traded shares.

Having considered the foregoing, our Directors are of the view that (i) we are not subject to any material tariffs or sanctions in our ordinary course of business; (ii) U.S. tariffs on the PRC and PRC retaliatory tariffs on U.S.-origin imports have not had a material adverse impact on our sales, cost of sales or overall financial performance; (iii) U.S. export control restrictions and the U.S. Final Rule have not had, and are not expected to have, a material adverse impact on our existing business operations; and (iv) the Final Rule will not have a material impact on our operations, financial performance or [**REDACTED**] prospects. Based on the independent due diligence conducted by the Sole Sponsor, nothing has come to the Sole Sponsor's attention that would reasonably cause them to disagree with the Directors' foregoing view.

AWARDS AND RECOGNITIONS

During the Track Record Period and up to the Latest Practicable Date, we received awards and recognition in respect of our products, technology and innovation, significant ones of which are set forth below:

Award/Recognition	Award year	Awarding Institution/Authority
Beijing Digital Economy Benchmark Enterprise (北京市數字經濟標桿企業)	2026, 2025	Beijing Software Industry Association, Beijing Digital Economy Promotion Center, China Academy of Information and Communications Technology (CAICT), and Deloitte China (北京軟件協會、北京市數字經濟促進中心、中國信通院及德勤中國)
China Unicorn Enterprise (中國獨角獸企業)	2026, 2025	Zhongguancun Unicorn Enterprise Development Alliance (中關村獨角獸企業發展聯盟)
2025 Information Technology Application Innovation Solutions (Typical Solution) (2025年信息技術應用創新解決方案(典型解決方案))	2026	Network Security Industry Development Center of the Ministry of Industry and Information Technology (工業和信息化部網絡安全產業發展中心)
AI Storage Laboratory KV Cache Specialized Testing Certificate for Storage Software Systems (AI儲存實驗室·KV Cache專項測試儲存軟件系統檢測證書)	2026	China Telecommunication Technology Labs and CAICT (中國泰爾實驗室、中國信通院)
National High-Tech Enterprise (國家高新技術企業)	2025	Beijing Municipal Science and Technology Commission, Beijing Municipal Finance Bureau, and Beijing Municipal Taxation Bureau (北京市科委、北京市財政局及北京稅務局)
AI Cloud 100 China 2025 – AI Infrastructure Benchmark Enterprise (AI Cloud 100 China 2025—AI基礎設施—標桿企業)	2025	Eminence Ventures (靖亞資本)
2025 Typical Practices in Data Storage (2025年度數據存儲典型實踐案例)	2025	Data Storage Committee of China Electronics Standardization Association (中國電子工業標準化技術協會數據存儲專業委員會)
Catalog of Large Model Application Delivery Service Providers (大模型應用交付供應商名錄)	2025	China Academy of Information and Communications Technology (中國信通院)
Top 20 Leading Artificial Intelligence Companies (人工智能領軍企業TOP 20)	2025	Synced (機器之心)
National "Specialized, Sophisticated, Distinctive, and Innovative" (SSDI) "Little Giant" Enterprise (國家專精特新"小巨人"企業)	2024	Ministry of Industry and Information Technology (工業和信息化部)