GLOSSARY OF TECHNICAL TERMS

This glossary of technical terms contains explanations of certain terms used in this document in connection with our Group and our business. The meaning of these terms may not necessarily correspond to standard industry meaning or the usage of these terms.

"AD"	active directory
"AGI"	artificial general intelligence
"AI agent"	an intelligent system capable of autonomously performing specific tasks on behalf of a user to achieve a proposed goal
"AI Infrastructure"	the integrated set of hardware, software, data systems, or cloud resources necessary for training and inferencing AI models
"AI model"	mathematical algorithms which can take unstructured data as input and transform them into informative outputs through its "intelligence," namely, the capability of perceiving the world, transcribing and organizing information, enhancing or generating contents, or making decisions
"AI"	artificial intelligence, an area of computer science that focuses on simulating human intelligence by machines
"AIGC"	artificial intelligence generated content, content such as text, visual, and audio created automatically by artificial intelligence technologies without direct human creation
"AI-native"	refers to products, services, or systems that are built from the ground up with artificial intelligence as a core component, rather than integrating AI as an add-on. AI-native solutions are designed to leverage AI technologies in their fundamental architecture and functionality
"algorithm"	a procedure or formula for solving a problem, based on conducting a sequence of specific actions, especially by a computer
"AOI"	automated optical inspection
"API"	a set of rules and tools that allows different software systems to communicate and interact with each other
"app" or "application"	application software designed to run on smartphones and other mobile devices
"bucket"	a logical container in an object storage system used to store and manage objects

GLOSSARY OF TECHNICAL TERMS

"CAGR"	compound annual growth rate
"cloud"	a network of remote servers hosted on the Internet and used to store, manage, process data, and offer algorithms in place of local servers or personal computers
"Cold data"	data that is rarely accessed and is primarily stored for cost efficiency
"Data Lake Storage"	a storage architecture, used to store large volumes of structured and unstructured data in raw format (often built on object storage)
"DRAM"	dynamic random access memory
"EB"	exabyte
"EC"	erasure coding
"edge computing"	a distributed computing paradigm that brings computation, data storage, and analytics closer to the source of data generation
"ESG"	environmental, social and governance
"firmware"	an essential built-in software that manages the basic functions of hardware
"full-protocol access"	the capability of a distributed storage platform to support block, object and file storage protocols simultaneously within a single cluster, enabling unified data access and management across different storage types
"generative AI"	a type of artificial intelligence that creates new content by learning patterns from existing data and generating original outputs
"GHG"	greenhouse gas
"GPU"	graphic processing unit
"HBM"	high bandwidth memory, a high-speed computer memory interface for 3D-stacked synchronous dynamic random-access memory
"HDD"	hard disk drive
"inference cost"	the cost of computational resources needed to use a trained AI model to process inputs and generate outputs

"IOPS"	input/output operations per second, a performance metric measuring the number of input and output operations a storage system can process per second
"ISVs"	independent software vendors, independent third-party software developers that develop and integrate software solutions on or with the company's platform
"kernel"	an open-source, monolithic, Unix-like operating system kernel that serves as the core interface between a computer's hardware and its software applications, responsible for managing resources like CPU, memory, and peripheral devices
"KV Cache"	an optimization technique used in large language model inference that stores the computed vectors of previous tokens in GPU memory, which eliminates the need to re-compute the attention mechanism for past context at every generation step, significantly accelerating inference speed at the cost of relatively higher memory consumption
"LCD"	liquid crystal display
"LDAP"	lightweight directory access protocol
"LLM"	large language model, an AI language model that uses deep learning techniques and massively large data sets to understand, summarize, generate and predict new content
"NFS"	network file system, a file-level storage protocol that enables remote access to files over a network
"open platform"	a publicly accessible interface that allows third-party developers to integrate with a company's systems or services by providing standardized programming access, enabling the development of applications or services based on the platform's capabilities
"open-source"	a source code that is made freely available for possible modification and redistribution
"PACS"	picture archiving and communication system
"parameters"	internal numerical variables or configurations that are learned and adjusted by the foundation model during its training process
"PB"	petabyte
"PCB"	printed circuit board
"PMO"	product management office

"POSIX"	portable operating system interface, a set of standards defining file system interfaces and behaviours to ensure application compatibility
"PV"	photovoltaic
"QoS"	quality-of-service
"R&D"	research and development
"RAG"	retrieval-augmented generation
"RDMA"	remote direct memory access
"RTC"	run-to-completion
"S3"	simple storage service, an object storage interface that provides scalable access to data stored in buckets through a standardized application programming interface
"SAS"	serial attached SCSI, a high-performance storage interface standard used to connect enterprise storage devices, providing higher reliability and throughput than consumer-grade interfaces
"SATA"	serial advanced technology attachment, a storage interface standard used to connect storage devices such as hard disk drives and solid-state drives to host systems
"SCSI"	small computer systems interface, a set of standards that define how computers communicate with and control storage devices through commands for data transfer and management
"SDK"	software development kit, a collection of tools, libraries, documentation, and code samples that developers use to build applications for a specific platform, system, or service
"SLA"	service level agreements, formal contracts between a service provider and a customer that define the expected level of service, including performance metrics such as uptime, response time, and support quality, as well as remedies if standards are not met
"SSDs"	solid-state drives
"T&I"	training and inference
"TLC"	triple-level cell, a type of flash memory technology in which each memory cell stores three bits of data, commonly used in enterprise and data-center solid-state drives

"Vector DB"	vector database
"XEBS"	our enterprise block storage, a distributed block-storage system
"XEDP"	our enterprise data platform, a unified data storage platform that integrates file, object and block storage into a single system
"XEOS"	our enterprise object storage, a enterprise-grade object storage system designed to support the long-term retention, governance and distribution of unstructured datasets
"XGFS"	our global file system, a distributed file storage system designed to support training, big data processing and office/media workloads through a unified global namespace and broad protocol support
"XScale"	unstructured metadata engine
"XSEA"	eXtreme shared-everything architecture, a distributed storage architecture developed for the all-flash data center era
"I/O Wall"	a bottleneck in AI training where storage systems cannot deliver data fast enough to match AI chips' concurrent access demands, causing chips to remain idle and reducing training efficiency
"Memory Wall"	a bottleneck in AI inference where limited on-chip memory cannot hold the intermediate results required by complex models, forcing use of slower external storage or additional high-cost chips
"Data Gravity Wall"	the difficulty of migrating large volumes of enterprise data across multiple systems into high-performance storage for AI training, which is time-consuming and risks creating data silos or version inconsistency
"Hot, warm and cold data"	classifications of data by access frequency, with hot data actively accessed during AI training and inference, warm data retained for periodic reference, and cold data retained for long-term archiving