

NEXAVM

— ENTERPRISE STORAGE DIVISION —

TECHNICAL WHITE PAPER

NexaVM N5000 Storage Array

High-Performance · Independently Innovated · Enterprise-Class Storage

Version 1.0 · October 2026

NexaVM Storage Technology Division

For internal review and partner distribution only

About this Document

Purpose. This document presents the technical architecture, design principles, and feature set of the NexaVM N5000 enterprise-class storage array. It is intended for system architects, pre-sales engineers, integration partners, and end-customer technology evaluators.

Scope. The white paper covers hardware architecture, independent storage IP, the NexaVM SAN OS software stack, advanced data services, and key software specifications of the N5000 platform.

Audience. IT infrastructure architects, storage administrators, value-added resellers, managed service providers, and OEM / ODM partners evaluating enterprise SAN / NAS deployments.

Conventions. Hexadecimal numbers are prefixed with 0x. Acronyms such as iSCSI, SAS, RAID, and BBU are defined in the Abbreviations section at the end of this document.

Contents

1 Overview	—
2 Technical Advantages	—
2.1 Fully Independent Storage System Logic IP	—
2.2 Independently Innovated SAS Control Chip	—
2.3 Independently Developed Hard RAID Engine	—
2.4 ASIC-Level iSCSI Protocol Parsing	—
2.5 Independent Storage Operating System — NexaVM SAN OS	—
3 Hardware Design	—
3.1 Hardware Architecture	—
3.2 Hardware Specifications	—
3.3 Hardware Advantages	—
4 Software Design	—
4.1 Software Features	—
4.2 Software Specifications	—
5 Abbreviations and Terms	—

1

CHAPTER

Overview

Product overview, market context, and key value proposition.

Driven by the wave of informatization, the widespread application of technologies such as IoT, big data, cloud computing, artificial intelligence, 5G, edge computing, and blockchain across various fields has driven explosive growth in data volume, and the potential value of data has become increasingly prominent. However, during the high-speed growth period of digital transformation, enterprises not only face the challenge of surging data volumes but also need to cope with the resulting management complexity and cost pressure.

The NexaVM N5000 storage array is a high-performance enterprise-class disk array launched by NexaVM based on more than ten years of R&D experience and technical accumulation in the storage field. The product adopts multi-core CPUs, the first domestically independently innovated enterprise-class storage system NexaVM SAN OS, and combines NexaVM's independently developed ASIC-level SAS engine, RAID engine, and iSCSI engine acceleration technologies. Compared with similar storage arrays from other vendors, performance is improved by 50% and power consumption is reduced by 60%, truly achieving the goal of independent innovation. It meets various application requirements such as database OLTP/OLAP, Exchange, server virtualization, and video surveillance, and can provide storage capability support for critical businesses of users in multiple important industries including government, enterprise, healthcare, and education.

It supports data compression and deduplication, thin provisioning, dynamic expansion, LUN snapshot, LUN copy, remote synchronization, remote mirroring, and other data protection and advanced disaster recovery features; supports unified management of all features in the user interface. These excellent hardware designs and software features not only improve the scalability and reliability of the storage system but also greatly reduce system management difficulty and operational costs, enabling enterprise users to focus more on core business rather than storage management.

2

CHAPTER

Technical Advantages

Core technology pillars and architectural innovations.

The NexaVM N5000 storage array is an enterprise-class storage array independently developed by NexaVM. Through the complete integration of independently controllable software and hardware core technologies, the product excels in performance, reliability, and flexibility, providing users with more efficient, scalable, and easy-to-maintain storage solutions, helping enterprises achieve higher-level breakthroughs in data management, storage efficiency, and security.

The NexaVM N5000 storage array can provide flexibly customized storage services, reduce dependence on external suppliers, ensure data transmission security, redundancy protection, and fault tolerance, and enhance the reliability and stability of the storage system. At the same time, the deeply integrated hardware and software system optimizes storage performance, reduces latency, and improves data access speed and throughput, meeting the needs of enterprise users in high-load and large-scale data environments.

2.1 Fully Independent Storage System Logic IP, Building a Trusted and Controllable Foundation for Data Storage

Relying on NexaVM's hardware technology research capabilities far exceeding the industry average, the N5000 storage array possesses fully independent storage system logic IP, covering mainstream storage interfaces (such as NVMe, SAS, SATA), data transmission protocols (such as RDMA, iSCSI), data protection and management technologies (such as RAID), and high-speed network interfaces (such as Ethernet and Fibre Channel). This enables the N5000 storage array not only to have higher performance and reliability but also to provide more customized options in product development and upgrades, meeting the storage needs of different enterprises; at the same time, it reduces dependence on external suppliers and enhances the controllability and innovation capability of product technology.




Figure 2-1 NexaVM Fully Independent Storage System Logic IP

The fully independent storage system logic IP of the NexaVM N5000 storage array involves key aspects such as the storage system architecture, data processing, and management. It is not only the foundation

of the storage system but also ensures the reliability, controllability, flexibility, and security of the system, constituting the core support of the trusted and controllable foundation for data storage.

2.2 Independently Innovated SAS Control Chip, Creating a Secure and Reliable Engine for Data Transmission

SAS (Serial Attached SCSI) is a high-performance serial interface standard mainly used to connect computer storage devices such as hard disk drives and solid-state drives. Compared to traditional parallel SCSI, SAS provides higher data transfer rates, longer transmission distances, and greater reliability. It supports full-duplex communication, enabling simultaneous data reading and writing, improving data throughput; and supports multi-device connections, suitable for large-scale storage systems. SAS is also backward compatible with SATA hard drives, providing more flexible storage solutions, and is therefore widely used in enterprise-class storage and data center environments.

Independent SAS control refers to the storage device or array system completely implementing and managing SAS protocol control and data transmission processes through independently developed control chips and firmware, enabling the storage device to achieve complete control of the SAS protocol at both hardware and software levels. Through independent SAS control, data transmission efficiency can be optimized, system performance improved, latency reduced, and error detection and recovery capabilities in data transmission enhanced. In addition, independent SAS control enables storage devices to be customized according to specific requirements, enhancing flexibility and scalability, avoiding dependence on external hardware suppliers, and improving product reliability and stability.

The NexaVM N5000 storage array possesses a fully independently developed SAS control chip. Through SAS drivers and SAS firmware code, it implements the complete SAS protocol from the ground up, enabling the NexaVM N5000 storage array to exercise finer control over various issues in storage devices and data transmission processes, reducing the risk of communication failures and data loss, and ensuring that the storage system can still provide stable and secure services when facing challenges such as large-scale data transmission, hardware failures, and network attacks, building a trusted and controllable data transmission engine.

2.3 Independently Developed Hard RAID Engine, Enhancing Storage Data Redundancy and Fault Tolerance

RAID (Redundant Array of Independent Disks) is a technology that combines multiple hard drives together, using different data distribution and redundancy strategies to improve storage performance and data protection. Common RAID levels include RAID 0, RAID 1, RAID 5, RAID 6, and RAID 10. Among them, RAID 0 divides data into small blocks and writes them to multiple hard drives separately, providing high performance but no redundancy (if one hard drive fails, all data will be lost); RAID 1 (mirroring) completely copies data to two hard drives, providing data redundancy and improving data security (if one hard drive fails, data remains safe); RAID 5 (striping with parity) distributes data and parity information across multiple hard drives, providing high performance and redundancy capability (can tolerate one hard drive failure without data loss); RAID 6 (dual parity) is similar to RAID 5 but uses two parity drives, capable of tolerating two simultaneous hard drive failures, enhancing data protection capability; RAID 10 (RAID 1 + RAID 0) combines RAID 1 mirroring and RAID 0 striping, providing data redundancy and high performance, but requires at least four hard drives. RAID is widely used in storage systems requiring high performance and high reliability, especially suitable for enterprise-class storage solutions.

Independent RAID control refers to the storage array completely implementing RAID protocol management and control from the ground up through an independently developed hardware RAID engine. Through hardware-accelerated processing of data redundancy management, error recovery, and data verification operations, it can significantly improve the performance and stability of the storage system. Independent RAID control not only optimizes data transmission and processing efficiency but also enhances the system's fault tolerance and data protection, meeting enterprise-class storage requirements for high performance and high reliability.

The NexaVM N5000 storage array possesses a fully independently developed ASIC-level hard RAID engine, implementing the complete RAID protocol from the ground up, fully leveraging hardware characteristics, improving data processing performance, system reliability, and fault tolerance, while reducing dependence on external hardware, ensuring that data remains safe, complete, and continuously accessible even under hard drive failures or system pressure, greatly enhancing the reliability and availability of the storage system, thereby providing users with more stable and reliable storage services.

2.4 ASIC-Level iSCSI Protocol Parsing, Accelerating Storage System Data Read/Write

iSCSI (Internet Small Computer Systems Interface) is a storage area network (SAN) protocol based on IP networks, used to connect storage devices with servers through Ethernet for data access. iSCSI encapsulates traditional SCSI commands in IP packets, allowing data to be transmitted over local area networks (LAN) or wide area networks (WAN), thereby enabling remote sharing and management of

storage resources. iSCSI provides high performance and low latency similar to SCSI, but because it uses standard network infrastructure, it has lower deployment costs and higher flexibility compared to traditional storage connection methods (such as Fibre Channel), and is widely used in enterprise storage, data centers, and virtualization environments.

iSCSI acceleration refers to optimizing the data transmission of the iSCSI protocol in the storage array through TOE hardware acceleration technology. Under this technology, the storage array parses and processes the iSCSI protocol through dedicated ASIC chips, reducing the CPU burden and achieving faster data read/write and transmission. This acceleration technology can effectively improve the performance of iSCSI storage systems, especially when processing large amounts of data, significantly reducing latency and improving throughput.

The NexaVM N5000 storage array, relying on TOE technology, parses the iSCSI protocol at the ASIC level, capable of accelerating read/write processing, significantly improving storage system performance, especially in high-load environments, significantly accelerating data read/write speeds, reducing latency, and ensuring faster response times. In addition, because hardware acceleration can share the task of processing protocols, it reduces the storage CPU's computing burden, allowing storage to concentrate resources on other tasks, thereby improving overall system efficiency. At the same time, iSCSI acceleration technology improves the stability and reliability of data transmission, enabling the storage array to better support large-scale concurrent access.

2.5 Independent Storage Operating System NexaVM SAN OS, Deeply Integrated with Storage Hardware

A storage operating system is a software system specifically used to manage and control storage devices. It is responsible for coordinating data read/write operations between storage hardware and computer systems, providing functions such as storage resource management, data protection, performance optimization, and security assurance. Through the storage operating system, users can effectively manage the storage array, achieve data backup, redundancy, recovery, and disaster recovery, while improving the performance, scalability, and reliability of the storage system.

The storage operating system NexaVM SAN OS developed by NexaVM with independent intellectual property rights can be customized and optimized according to specific hardware and user requirements, effectively improving system performance, reliability, and flexibility. Compared to other storage vendors relying on third-party operating systems, the NexaVM N5000 storage array, through its proprietary operating system NexaVM SAN OS, can better combine hardware acceleration functions, providing more efficient data processing capabilities and lower latency. In addition, NexaVM SAN OS features a simple

and easy-to-use Web management platform, enabling users to easily configure and manage the storage system, reducing operational complexity and maintenance costs. Therefore, the NexaVM N5000 storage array, based on the independently developed storage operating system NexaVM SAN OS, can improve the overall performance, scalability, and ease of use of the storage array, meeting enterprise users' needs for efficient, stable, and easy-to-manage storage solutions.

The N5000 storage array, by integrating the above technical advantages, can optimize data transmission, improve storage efficiency, reduce latency, and provide customized storage services for users, meeting enterprise needs in data security, management, and expansion. As digital transformation advances, NexaVM N5000 will continue to drive the development of the storage industry with its independently controllable technological innovation, bringing enterprise customers more efficient, secure, and easy-to-manage storage experiences.

3

CHAPTER


Hardware Design

Chassis design, controller architecture, and key hardware specs.

The NexaVM N5000 storage array is a 19-inch, 2U, 12-bay domestically independently controllable storage disk array, adopting multi-core CPU processors, supporting Gigabit network, 10 Gigabit network, and FC network. It adopts a dual-controller architecture, with all modules having redundant design and no single point of failure; and power supplies, controller boards, hard drives, etc. all support hot-swapping. The entire machine provides 2 power module slots, 2 storage controller slots, and 12 2.5/3.5-inch mixed-use hard drive slots; the multi-core CPU processor on the main control board runs the self-developed NexaVM SAN OS storage operating system. The back end uses SAS technology to access hard drives, and the front end uses TOE technology-based 10 Gigabit network and FC network to provide storage services.

3.1 Hardware Architecture

The hardware architecture of the NexaVM N5000 storage array is shown in Figure 3-3. Its hardware system consists of multiple hardware components including controllers, chassis, backplane, power supplies, and hard drives. All components work together to provide efficient and reliable storage services for the storage system, meeting enterprise needs in data management, security, and storage performance.



The role of each core component is as follows.

- ▶ **Controller**

The controller is the core component of the N5000 storage array, responsible for the management and scheduling of the entire storage system. It has functions such as processing data access requests, managing data read/write operations, storage policies, data redundancy, and load balancing. The

controller interacts with hard drives and other components to achieve data storage and protection, ensuring efficient data transmission and high reliability.

► Chassis

The chassis is the physical enclosure of the N5000 storage array, mainly used to install and protect various hardware components. It provides physical space to accommodate hard drives, controllers, power supplies, etc., and has heat dissipation functions to ensure stable temperature control of the storage system. The chassis also provides a reasonable structural design for easy replacement and maintenance of hardware components.

► Backplane

The backplane is the connection platform inside the chassis of the N5000 storage array, responsible for physically connecting various hardware components (such as hard drives, controllers, power supplies, etc.) through circuits. It provides a transmission channel for data flow and is an important medium for data exchange between hard drives and controllers.

► Power Supply

The power supply provides stable power supply for the N5000 storage array, ensuring the normal operation of all system components. The NexaVM N5000 storage array adopts a redundant power supply design. When one power supply fails, the other power supply will automatically take over, ensuring uninterrupted operation of the storage system. At the same time, the NexaVM N5000 storage array power supply also has efficient energy conversion and protection functions, capable of preventing power fluctuations from damaging storage devices.

► Hard Drive

The hard drive is the most basic storage medium in the NexaVM N5000 storage array, responsible for persistent data storage. The N5000 storage array supports multiple hard drive types such as SAS, SATA, and SSD. Through these hard drives, it provides large-capacity, high-speed data storage space. And through interaction with the controller, it externally implements data reading and writing to ensure data security and efficiency.

3.2 Hardware Specifications

The hardware specifications of the NexaVM N5000 storage array are shown in Table 3-1.

Table 3-1 NexaVM N5000 Storage Hardware Specifications Table

3.3 Hardware Advantages

Through excellent hardware design and powerful hardware configuration, the NexaVM N5000 storage array has the following hardware advantages.

- ▶ Equipped with high-performance hardware components, significantly improving overall device performance

The NexaVM N5000 storage array is equipped with high-performance hardware components. The CPU adopts a domestic processor, supporting high-speed data computation and transmission, improving overall system performance; the backplane uses domestic high-speed connectors and domestic power connectors to link various hardware modules, obtaining power from power modules to provide to the main control board and hard drives, providing high-speed data channels for connecting the main control board with hard drives and for dual-controller data transmission and management between main control boards, capable of optimizing data transmission speed, ensuring efficient and low-latency data transmission between the main control board and hard drives, and between main control boards, thereby improving the overall performance of the storage array.

- ▶ Optimizing internal structure and data transmission paths, effectively accelerating data flow

The NexaVM N5000 storage array fully utilizes the performance potential of hardware, optimizing the internal structure and data transmission paths of the storage system through multi-level cache architecture, parallel processing, and hardware acceleration engines, effectively accelerating data flow and achieving high-speed data read/write operations.

- ▶ Redundant design for all core modules, effectively avoiding single points of failure

The NexaVM N5000 storage array not only adopts a redundant dual-controller architecture, but other core modules such as power supplies are also redundantly designed, effectively avoiding single points of failure, improving the reliability and fault tolerance of the hardware system, enabling the N5000 storage array to be suitable for critical business systems, data centers, and cloud computing application scenarios with strict requirements for high availability and business continuity in finance, healthcare, government, and large enterprises.

- ▶ Important components support hot-swapping, ensuring system business continuity

Important components such as power supplies, controller boards, and hard drives support hot-swapping, enabling the NexaVM N5000 storage array to perform hardware replacement or upgrades without shutting down the system power, improving system maintainability and business continuity, and reducing the risk of production or service interruption caused by downtime.

4

CHAPTER

Software Design

Software features, advanced data services, and specifications.

4.1 Software Features

The NexaVM N5000 storage array supports multiple software features designed to provide high storage system availability, data protection capability, read/write performance, and manageability, capable of meeting the needs of users of different scales and different businesses.

4.1.1 Dual-Controller Active-Active Working Mode

Dual-controller active-active working mode means that there are two controllers in the storage system, and both controllers are in an active state, capable of simultaneously processing read/write requests and data management tasks. Each controller has complete access permissions and can independently process I/O operations and data flows. This mode is different from the traditional dual-controller active-passive mode, in which one controller is the active controller and the other is in standby state, only activated when the primary controller fails. In dual-controller active-active mode, the two controllers share the load between them, improving the system's processing capability and throughput, ensuring efficient data operations. At the same time, if one controller fails, the other controller can seamlessly take over, ensuring high system availability, avoiding business interruption, and reducing the risk of single points of failure.

The main advantages of this working mode include:

- Load balancing: Two controllers process data requests simultaneously, improving overall performance and throughput.
- High availability: Two controllers work simultaneously. Once one controller fails, the other can immediately take over, ensuring continuous data availability.
- Fault tolerance: Enhances the system's fault tolerance, reduces potential service interruption time, suitable for mission-critical and high-reliability enterprise-class applications.

The NexaVM N5000 storage array includes dual management and business channels between dual controllers. Among them, the management channel is mainly used for management-related data communication such as synchronizing configuration modifications and dual-controller negotiation; the business channel is mainly used for user data transmission such as memory mirroring and read/write request forwarding. The dual-controller channel adopts PCIE interconnection, with the management channel prioritized, ensuring immediate response to various events. If one controller fails, the other controller automatically takes over all functions.




Figure 4-1 Storage Active-Active Internal Structure Diagram

The NexaVM N5000 storage array dual controllers have a master-backup distinction in management, namely Backup (management backup) and Master (management master). All management requests are initiated by the master controller. If a management request from a management terminal is sent to the backup controller, the backup controller will forward it to the master controller, and the master controller will process it. The dual controllers can automatically determine the presence status and master-backup priority through hardware interconnection IO lines. Each controller simultaneously transmits heartbeat messages and status messages through the management channel, ensuring their respective operating states and timely handling of abnormal events.

Figure 4-5 Storage Active-Active Write IO Flow

The NexaVM N5000 storage array dual controllers adopt active-active working mode for data access. Both controllers simultaneously respond to user IO requests. The application server can determine which controller each IO is dispatched to based on the multipath configuration. Each IO request is only sent to a single controller, and the controller that receives the IO request is the access controller. In addition, all LUNs have an owning controller, and all LUN access needs to be processed through the owning controller.

4.1.2 SAN/NAS Unified Storage

SAN (Storage Area Network) is a high-speed network specifically used to connect storage devices (such as disk arrays) with servers. SAN provides block-level storage, meaning storage devices directly provide storage volumes (LUNs) through the network for servers to access. SAN typically uses protocols such as

FibreChannel or iSCSI. It can provide high throughput and low-latency data transmission, suitable for enterprise-class applications requiring large-scale data processing, such as databases and virtualization.

NAS (Network-Attached Storage) is a solution for sharing file-level storage through the network. Storage devices connect to clients (usually computers or servers) through standard networks (such as Ethernet). Unlike SAN, NAS provides file-level access through network file protocols (such as NFS or SMB/CIFS). NAS is suitable for shared storage, suitable for file storage and management, such as document management, file sharing, and backup application scenarios.

SAN/NAS unified storage refers to integrating the functions and features of SAN and NAS into one storage system, thereby achieving unified storage management and data access.

SAN/NAS unified storage has the following benefits:

- ▶ Flexible data access

Users can access storage through file-level access (NAS) or block-level access (SAN) as needed, so that the most suitable access method can be selected according to specific application requirements.

- ▶ Simplified storage management

Through a unified storage platform, users can more conveniently manage different types of storage requirements, simplifying storage architecture and management.

- ▶ Improved resource utilization

The unified storage platform can efficiently allocate and schedule storage resources, avoiding excessive waste or shortage of storage resources, and improving the efficiency of hardware resource usage.

- ▶ Enhanced scalability

Through the combination of SAN and NAS, storage capacity or performance can be expanded as needed, supporting multiple workloads and meeting the storage needs of different enterprises.

- ▶ Improved data security and availability

Unified storage solutions typically provide data protection, backup, disaster recovery, and other functions, ensuring data integrity and high availability.

SAN/NAS unified storage applicable scenarios include:

- ▶ Mixed workload scenarios

When enterprises need both block-level storage (such as databases, high-performance applications) and file-level storage (such as file sharing and backup), SAN/NAS unified storage can meet both needs.

- ▶ Virtualization environment scenarios

Many virtualization environments require efficient block-level storage and file-level storage. The unified storage platform can support these different storage requirements and provide a unified management interface.

- ▶ Big data and multimedia storage scenarios

Suitable for environments requiring high throughput, low latency, and simultaneously sharing large amounts of files, such as big data processing, video editing, and archiving.

The NexaVM N5000 storage array software platform implements the integration of SAN and NAS functions, providing native file sharing interfaces and block mapping interfaces in one software platform, providing customers with a unified storage product that simultaneously satisfies file access and block access, protecting customer investment.

In the NexaVM N5000 storage array SAN/NAS unified storage architecture, LUNs and file systems run in parallel on top of the intelligent tiered storage pool. Below the storage pool is the RAID subsystem. LUNs can implement block storage mapping through FC or iSCSI protocols, and file systems can implement file sharing through CIFS, NFS, or FTP protocols.

4.1.3 Intelligent Tiered Storage

Intelligent tiering is a data storage technology that automatically allocates different types of data to storage media of different performance and cost tiers based on factors such as access frequency, importance, and storage cost, to optimize the use of storage resources and improve the overall performance and efficiency of the system.

The NexaVM N5000 storage array, based on the intelligent tiered storage pool, intelligently analyzes the access frequency of data, supports storing frequently accessed data on high-performance hard drives and storing infrequently accessed historical data on lower-performance hard drives. The intelligent tiered storage pool can simultaneously meet user performance and capacity requirements while reducing user costs.

The intelligent tiered storage pool divides storage space into several data blocks according to 4M, and LUN space allocation is also allocated according to block size. As shown in Figure 4-6, the intelligent tiered storage pool divides each tier of RAID into fixed-size physical blocks, and each LUN is also divided into logical blocks of the same size. Each logical block points to a physical block, and migrating the physical block corresponding to a logical block can change the data storage location. During data access, the storage pool statistically analyzes data access frequency in units of blocks, and through data migration, migrates frequently accessed data to hard drives with higher performance, and migrates infrequently accessed data to hard drives with large capacity and lower capacity cost.

Figure 4-6 *Intelligent Tiered Storage*

Optimizing storage space in the intelligent tiered storage pool requires three steps: statistics of IO access frequency, access frequency sorting, and data migration.

- ▶ Statistics of IO access frequency

IO access frequency is statistically based on the blocks of the storage pool. During the lifecycle of the storage pool, every issued IO is recorded for subsequent access frequency sorting. IO access statistics information is all recorded in memory. After device startup, IO access frequency is re-statisticized, and IO access frequency before device restart is not counted. After each data migration is completed, IO access frequency is weighted to weaken the impact of historical access frequency on subsequent statistical sorting.

- ▶ Access frequency sorting

Access frequency is sorted based on IO access frequency. In the same storage pool, all data block access frequencies are sorted from high to low, and the data migration threshold is determined according to the capacity size of each tier of the storage pool. Data blocks with higher access frequency are stored on higher-performance hard drives, determining which type of hard drive each block should be mapped to.

- ▶ Data migration

According to IO access frequency, migrate the storage location of data blocks in the tiered pool, so that data with higher access frequency is stored in higher tiers as much as possible, and data with lower access frequency is stored in lower tiers as much as possible. Data migration supports two trigger methods: manual trigger data migration and scheduled trigger data migration. Manual trigger migration can be manually triggered according to user needs, and scheduled trigger can set a fixed migration start time.

Therefore, the NexaVM N5000 storage array intelligent tiered storage is often applied in scenarios requiring efficient management of large amounts of data, especially in environments with large differences in data access frequency. Through intelligent analysis and automatic migration mechanisms, intelligent tiered storage can optimize the use of storage resources, improving performance while reducing costs.

4.1.4 RAID

RAID (Redundant Array of Independent Disks) is a storage technology that combines multiple hard drives to provide higher performance, data redundancy, fault tolerance, and higher capacity. RAID technology achieves different performance and data protection requirements through different levels (RAID levels). Common RAID levels are shown in Table 4-2.

Table 4-2 Common RAID Level Comparison

RAID Level	Characteristics	Advantages/Disadvantages	Application Scenarios
RAID 0	Divides data into multiple small blocks (stripes) and writes these stripes to multiple hard drives separately. No data redundancy.	Advantages: Provides extremely high read/write performance because all hard drives participate in data reading and writing simultaneously. Disadvantages: No redundancy; failure of any hard drive will cause the entire data loss.	Applications with high performance requirements but less important data security, such as video editing, temporary storage, etc.
RAID 1	Completely copies data to two or more hard drives, forming a mirror.	Advantages: Provides data redundancy; if any hard drive fails, data can still be read from another hard drive. Disadvantages: Low storage space utilization efficiency because each data has a mirror copy.	Environments with high data security requirements and relatively tight storage space, such as small databases, critical applications, etc.
RAID 10	Combines RAID 1 mirroring and RAID 0 striping technology. Data is first mirrored (RAID 1), then striped on the mirror (RAID 0).	Advantages: Provides high performance and high redundancy capability, supports parallel read/write operations, strong fault tolerance. Disadvantages: Requires more hard drives, low storage space utilization.	Enterprise-class storage systems requiring high performance and high reliability, such as databases, high-load applications.
RAID 5	Data and parity information are distributed across multiple hard drives, requiring at least three hard drives. Each data block has a corresponding parity block, and the parity block can be used to recover data on a failed hard drive.	Advantages: Provides high performance and redundancy capability, high storage space utilization (higher than RAID 1). Can tolerate one hard drive failure without data loss. Disadvantages: Lower write performance because parity data needs to be calculated and written.	Application scenarios requiring a balance of performance, capacity, and redundancy, such as enterprise-class data storage.
RAID 6	Similar to RAID 5 but uses two independent parity	Advantages: Higher redundancy capability	Large enterprise data storage and critical

	drives, capable of tolerating two simultaneous hard drive failures. Requires at least four hard drives.	than RAID 5, can tolerate two simultaneous hard drive failures. Disadvantages: Lower write performance because two parity data need to be calculated and written, and storage capacity is lower than RAID 5.	applications with extremely high data reliability requirements.
RAID 50	Combines the advantages of RAID 5 and RAID 0. First creates multiple RAID 5 arrays, then stripes these RAID 5 arrays through RAID 0.	Advantages: Provides higher performance than RAID 5 because data striping is distributed across multiple RAID 5 arrays, with faster read/write speeds. Good redundancy capability. Disadvantages: Requires more hard drives, and storage space efficiency is lower than RAID 5.	Large-scale storage systems requiring high performance and high redundancy, such as high-concurrency applications, large data processing.
RAID 60	Combines the advantages of RAID 6 and RAID 0. First creates multiple RAID 6 arrays, then stripes these RAID 6 arrays through RAID 0.	Advantages: Provides higher redundancy, stronger data protection capability than RAID 5 and RAID 50, capable of tolerating multiple hard drive failures. Performance is also better than RAID 6. Disadvantages: Requires more hard drives, lower storage efficiency.	Enterprise-class large-scale storage environments with extremely high data reliability requirements.
Linear RAID	Directly combines multiple hard drives into a large storage pool. Data is written to each hard drive sequentially (linearly), with no redundancy or striping between hard drives. Linear RAID is not a standard RAID level; it is simply a simple way to use multiple hard drives	Advantages: Simple to implement, no need to calculate redundant data or striping, high storage space utilization. Disadvantages: No redundancy; once a hard drive fails, data loss is unrecoverable.	Scenarios that do not require redundancy protection and only need simple storage, suitable for some low-cost storage needs such as temporary storage or backup data.

	as a single storage unit.		
--	---------------------------	--	--

The benefits of RAID technology include:

- ▶ Data redundancy and protection

RAID technology protects data by distributing data across multiple hard drives, so that even if some hard drives fail, data can be protected. Different RAID levels provide different levels of redundancy and fault tolerance.

- ▶ Improved performance

Through data striping and parallel processing, RAID can significantly improve the read/write performance of storage systems, especially RAID 0, RAID 10, and RAID 50.

- ▶ Increased storage capacity

RAID allows multiple hard drives to be combined into one logical storage unit, thereby increasing available storage space to accommodate growing enterprise data storage needs.

The NexaVM N5000 storage array supports all common RAID levels, specifically including:

4.1.5 Storage Methods

The NexaVM N5000 storage array supports two storage methods: block storage and file storage. Users can choose the appropriate storage method according to their needs.

- ▶ Block Storage

Block storage is a storage method that divides data into fixed-size blocks and stores them independently on storage devices. Each block has a unique identifier (usually a block address), and the storage system disperses data storage across these blocks.

Operating systems and applications access specific data blocks through addresses. Block storage stores and manages data in units of blocks, supports LUN access, and is suitable for high-performance, low-latency application scenarios such as database storage, virtualization environments, and large-scale data processing. It can provide fine-grained storage control and high throughput, suitable for scenarios requiring frequent read/write and large data volume operations.

It has the following characteristics:

- ① Direct access: Block storage provides a more direct access method. The operating system can directly read and write storage blocks without needing to know the specific content of the data.
- ② Flexibility: Applications can read and write any data block arbitrarily according to requirements, without needing to access data in a specific order or structure.

③ High performance: Because block storage is managed at the lowest disk level, system access speed and performance are higher, suitable for scenarios requiring frequent data read/write.

► File Storage

File storage organizes data into files and manages them through a file system. Each file has a path to identify it. The file system is responsible for managing file reading, writing, storage, and indexing operations. Users access files through file system protocols without needing to care about the specific details of the underlying storage. File storage is suitable for high-performance, low-latency application scenarios such as database storage, virtualization environments, and large-scale data processing. It can provide fine-grained storage control and high throughput, suitable for scenarios requiring frequent read/write and large data volume operations.

It has the following characteristics:

- ① File system management: File storage relies on the file system to manage stored data. The file system handles tasks such as file creation, modification, deletion, and permission management.
- ② Ease of use: Because it is file-based storage, users can access and manage data through paths just like operating local files, simplifying operations.
- ③ Sharing and collaboration: File storage is very suitable for team and organizational use because file storage provides support for multiple users to share and access files.

Table 4-3 Block Storage and File Storage Comparison

Characteristic	Block Storage	File Storage
Data Management	Stored in units of blocks, each block has an independent address	Stored in units of files, files accessed through paths
Storage Method	No hierarchical structure, only independent data blocks	Has a hierarchical structure of files and directories, suitable for storing files and folders
Performance	High performance, low latency, suitable for frequent read/write operations	Lower performance, suitable for simpler file access
Application Scenarios	Database storage, virtualization, transaction processing, high-performance computing, etc.	File sharing, backup storage, large-capacity storage, media management, etc.
Flexibility	High, application layer can finely control storage allocation	Lower, file access depends on file system protocols and access methods
Management Complexity	Higher, requires precise management of blocks, volumes, and storage devices	Lower, depends on file system, users access and manage data through paths

Sharing	Difficult to share, requires specific software or protocols	Strong, multiple users can share access to the same file through network protocols
---------	---	--

4.1.6 Hot Spare Disk

Hot spare disk functionality is a redundancy design in storage arrays used to improve system fault tolerance. A hot spare disk is an additional hard drive. When any hard drive in the storage array fails, the hot spare disk will automatically replace the failed hard drive, quickly taking over data storage tasks, thereby maintaining data integrity and normal system operation.

After enabling the hot spare disk function, when a hard drive in the storage array fails, the system will automatically connect the hot spare disk to the array, replacing the failed hard drive's storage tasks. The hot spare disk itself does not actively participate in data storage, only activating when a failure occurs. Through redundant data in RAID (such as parity information), the array will quickly rebuild lost data, writing the hot spare disk's data to the failed hard drive's location, ensuring continuous operation of the storage system and data integrity. After the hard drive is repaired or replaced, the hot spare disk returns to standby status.

Because some RAID mechanisms do not cause data loss when some hard drives fail (such as RAID5 when one hard drive fails, data is not lost; this state is called RAID degraded state), when RAID is in degraded state, a new hard drive can be added to the RAID group to replace the failed hard drive's position, and redundant data can be used to calculate and recover the original hard drive's data. In order to start reconstruction faster, devices typically deploy one or more hot spare disks. After RAID degradation, the hot spare disk will immediately join the RAID group and begin data reconstruction, aiming to complete data reconstruction as soon as possible.

The NexaVM N5000 storage array's hot spare disk function includes dedicated hot spare disks and global hot spare disks. Dedicated hot spare disks are configured for a specific RAID group and only take effect when that RAID fails. Global hot spare disks are configured globally for the device and will automatically join any RAID group when it fails. The NexaVM N5000 storage array's dedicated hot spare disks and global hot spare disks do not conflict. If a RAID group already has a dedicated hot spare disk configured, the dedicated hot spare disk will be used first after RAID degradation. If a RAID group does not have a dedicated hot spare disk configured, the global hot spare disk will be used.

4.1.7 Dual-Machine Active-Active

Dual-machine active-active (Active-Active) is achieved by deploying two storage array devices to distribute the load across both devices. Each device is in an active state and can simultaneously process data access requests, improving the performance, reliability, and fault tolerance of the storage system.

In a dual-machine active-active architecture, the two storage devices do not depend on each other for a single point of failure, but each has complete read/write processing capability. They share data and can automatically switch when a failure occurs, ensuring business continuity; at the same time, they can continuously provide efficient storage services. Even if one device fails, the other device will immediately take over its tasks, ensuring data is not lost and the system continues to operate.

The benefits of dual-machine active-active include:

- ▶ High performance: Two storage devices jointly undertake data processing tasks, improving system throughput and processing capability.
- ▶ High availability: Through redundant design, when one device fails, the other device immediately takes over, ensuring continuous system availability.
- ▶ Seamless switching: When a failure occurs, the switching process is transparent, with no impact on users and applications.
- ▶ Load balancing: Through parallel processing of requests, the system can achieve higher load balancing, avoiding single device overload.

Dual-machine active-active storage arrays are widely used in scenarios requiring high availability and high reliability, such as data centers, financial institutions, and telecommunications operators.

The NexaVM N5000 storage array's dual-machine active-active implements AA active-active. The LUNs of the two arrays are mirrored and synchronized in real time, and both ends can simultaneously process application server requests, improving resource utilization and system efficiency and performance, allowing customers to obtain maximum value from the dual-machine active-active system.

When a link fails or a single storage array fails, the server multipath software and the storage array's heartbeat check is interrupted, and the server will automatically switch all requests to the other storage array. When a single storage array fails, the other storage array can provide services independently. After fault recovery, data consistency incremental synchronization is automatically performed to quickly restore data consistency.

Inter-array interconnection supports IP and FC links. Users can choose link types and connection methods according to application scenarios. Devices support single or multiple links. It is recommended to use dual-machine dual-controller cross-interconnection to maximize security. There are two cross-interconnection diagrams as follows: direct cross-interconnection and switch cross-interconnection. As

long as there is any link between the two machines, normal data synchronization between the two machines can be achieved.

Figure 4-7 *Dual-Machine Active-Active Interconnection Methods*

Text in figure: Array A, Controller 0, Controller 1; Array B, Controller 0, Controller 1; Switch.

When the link between the two machines fails, the arrays cannot synchronize in real time. To ensure data consistency, the NexaVM N5000 storage array has a comprehensive arbitration mechanism to determine which array provides services, solving the split-brain problem. The dual-machine active-active networking diagram is shown in Figure 4-8.

Figure 4-8 *Dual-Machine Active-Active Schematic Diagram*

Text in figure: Server; Virtual Array; Storage Array A; Dual-Controller Dual-Machine Active-Active; Storage Array B; Dual-Machine Channel; Switch.

The multipath software on the server side automatically virtualizes the dual-machine mapped hard drives into a single hard drive on the server. Under normal circumstances, when the server writes data, each IO write request only sends a write data command to a single path. The array that receives the write command writes locally and simultaneously sends the write command to the other array. Only after both machines successfully execute will the server be notified that the command was successfully executed, ensuring real-time data consistency; when the server reads data, the array that receives the read data command only reads from its own device to respond to the read request, providing faster response speed.

When one array is not allowed to access, when the server writes data, the read/write command is only sent to the normal array, and data is only written to the normal array. The normal array simultaneously records the difference positions between the two machines. After the abnormal array recovers, incremental synchronization is performed by reading the differences to quickly restore data consistency.

When the link between the two machines fails, data cannot be mirrored in real time. To ensure data consistency, only one array can provide services to the server. To prevent dual-machine split-brain, our company provides two arbitration mechanisms: static priority mechanism and arbitration server mechanism.

The static priority mechanism does not require a third-party arbitration server. Users configure one array as the preferred site and the other array as the non-preferred site. In this mode, when the inter-array link fails, the non-preferred site is not allowed to access. The arbitration server mechanism allows both machines to access normally when the inter-array link is normal. When the inter-array link is abnormal, to ensure data consistency, only one array can provide services, and the other array is not allowed to access. When the inter-array link is abnormal, the arbitration server performs arbitration. The array that wins the arbitration provides services. The array that cannot communicate with the arbitration server loses the

arbitration. When both machines can communicate with the arbitration server, the preferred site wins the arbitration.

4.1.8 Host Multipath Support

Host multipath support is an important function of storage arrays, allowing hosts (servers) to connect to storage arrays through multiple physical paths, ensuring high data availability and better performance. Its functions are as follows:

- ▶ Redundant paths

By configuring multiple independent physical paths between the host and the storage array (such as different Fibre Channel or iSCSI connections), if one path fails, the host can automatically switch to another path, ensuring uninterrupted storage access.

- ▶ Load balancing

Multiple paths can be used for data transmission simultaneously. The storage array can intelligently allocate load based on conditions such as traffic and bandwidth, thereby improving data transmission throughput and response speed.

- ▶ Fault tolerance

If a path fails (such as Fibre Channel connection interruption or network failure), multipath support can automatically switch to a backup path, ensuring continuous business operation. This automatic switching mechanism helps improve the reliability and fault tolerance of the storage system.

Each LUN in the NexaVM N5000 storage array generates a unique WWN when created, and this WWN will not change until the LUN is deleted. By checking whether the WWN of the mapped hard drives is the same, the multipath software virtualizes two or more path-mapped hard drives into one hard drive. As shown in Figure 4-9, LUN1 maps DISK1 and DISK2 on the server through path 1 and path 2. The multipath software checks that the WWN of the two hard drives is the same and virtualizes the two hard drives into mpatha. Similarly, LUN2 and LUN3 are virtualized into mpathb and mpathc respectively.

Figure 4-9 Multipath Mapping

Text in figure: mpatha, DISK1, DISK2; mpathb, DISK3, DISK4; mpathc, DISK5, DISK6; Path 1, Path 2; LUN1, LUN2, LUN3.

The server can map SCSI devices from multiple paths into one virtual hard drive through multipath software. However, when the virtual hard drive reads and writes data, each read/write command only needs to be sent to one path. The multipath software needs to select the controller for each execution. As shown in Figure 4-10, server access to LUN includes 4 paths in total: "Server Port 1 <-> Switch 1 <->

Controller 1 Port 1", "Server Port 1 <-> Switch 1 <-> Controller 2 Port 1", "Server Port 2 <-> Switch 2 <-> Controller 1 Port 2", "Server Port 2 <-> Switch 2 <-> Controller 2 Port 2".

Figure 4-10 *Server Access Paths*

Text in figure: Server; Port 1, Port 2; Switch 1, Switch 2; Port 1, Port 2; Controller 0, Controller 1; LUN.

The NexaVM N5000 storage array controllers support Active-Active mode. Each path can access any LUN, but the efficiency of accessing through different controllers is different. Usually, a LUN has a preferred controller, and accessing the LUN through this controller is most efficient. Accessing the LUN through other controllers will result in efficiency loss.

If a path accesses the LUN's owning controller, the server is notified that the path is Active/Optimized (AO path). If a path accesses a LUN's non-owning controller, the server is notified that the path is Active/NonOptimized (AN path). The multipath software groups all AO paths into one path group and all AN paths into one path group.

When a path group contains multiple paths, read/write commands are load-balanced within the path group by the multipath software, fully utilizing the bandwidth of multiple links and improving the overall throughput of the system. There are three load balancing methods: round-robin algorithm, current IO minimum path algorithm, and IO service time shortest algorithm.

If an AO path group exists, the AO path is accessed preferentially. As long as an AO path exists, the multipath software will not access the AN path. When all AO paths fail, the multipath software will automatically switch to the AN path for data access.

When the AO path recovers, the multipath software will automatically restore to the AO path for data access.

Therefore, the host multipath support of the NexaVM N5000 storage array is mainly applied in scenarios requiring high availability, load balancing, and redundant storage environments.

4.1.9 Data Snapshot

A data snapshot is a technology that creates a point-in-time copy of the current state of data, allowing the storage system to record and save the data state at a specific moment without copying all data. This technology provides an efficient way to protect data and ensure rapid system recovery in the event of a failure. It has the following characteristics:

- ▶ Incremental storage

Snapshots only save data changes since the last snapshot (incremental data), so the storage space occupied is relatively small.

- ▶ Fast creation

Because there is no need to copy the entire data set, snapshots can be created very quickly, usually within seconds.

- ▶ No interruption

Snapshots are usually performed online and do not affect the normal business operation of the storage array.

- ▶ Convenient recovery

Once data loss or corruption occurs, data can be quickly recovered by restoring to a previous snapshot version, minimizing downtime.

When creating a storage pool, the NexaVM N5000 storage array divides the included RAID into several physical blocks according to a fixed size. When creating a LUN based on the storage pool, the LUN divides several physical blocks from the storage pool according to user configuration and records the relationship between physical blocks and the LUN's logical positions. When users access the LUN, they search for the corresponding physical block for reading and writing based on the logical position. The snapshot of the NexaVM N5000 storage array clones the relationship between the source LUN's logical blocks and physical blocks, and applies for new physical blocks when subsequent data is written to the LUN, ensuring that the original physical block data does not change. When creating a snapshot, the NexaVM N5000 storage array first creates a new logical volume in the storage pool, then copies the relationship between the source LUN's logical blocks and physical blocks to this logical volume, making this logical volume share all physical blocks with the source LUN. Because the logical volume and source LUN mapping relationship is exactly the same, this logical volume contains the same data as the source LUN. This logical volume is the snapshot of the source LUN at the current moment. The snapshot creation process does not require copying any user data, only copying logical relationships, and the snapshot can be created instantly. When the source LUN of the NexaVM N5000 storage array snapshot performs data writing, the storage system checks whether the physical block is shared. If the physical block is a shared block, it applies for a new physical block for the source LUN, copies the old physical block data to the new physical block, then writes the new data to the new physical block, and finally modifies the correspondence between the source LUN logical block and physical block, making the source LUN logical block position point to the new physical block position. If the physical block is not shared with the snapshot, data is written directly.

Because traditional storage arrays' storage pools do not have a relationship between LUN logical blocks and physical blocks, when copying on write, they cannot apply for new physical blocks for the source LUN and can only apply for new physical blocks for the snapshot. Therefore, when there are many snapshots, traditional storage arrays need to apply for a physical block for each snapshot during copy-on-write and

execute multiple copy-on-write operations, wasting more storage space and resulting in slow writes. However, the NexaVM N5000 storage array snapshot technology can avoid the above problems. Even when multiple snapshots exist, only one block application and one data copy are required, effectively reducing storage space occupation and improving data write efficiency.

Compared to traditional storage arrays that need to copy snapshot data to the source LUN during snapshot rollback and can only complete rollback after copying is complete, the NexaVM N5000 storage array snapshot rollback only needs to copy the correspondence between the snapshot's logical blocks and physical blocks to the source LUN to restore the original data. The rollback process does not require copying any user data, only copying logical relationships, and the snapshot can also be completed instantly.

Figure 4-11 *Snapshot Mapping Relationship Diagram*

Text in figure: Source LUN logical block; Physical block; Create snapshot; Snapshot logical block; Copy-on-write; Snapshot rollback.

4.1.10 Data Clone

Cloning refers to the technology of accurately copying data from one logical unit (such as a LUN, volume, or disk) in a storage array to another target storage unit. Cloning is a method of data replication that can create a complete copy of the source data, and the copy is consistent with the original data in content and format.

During the cloning process, the storage system scans the data to be cloned and generates a complete copy of that data. Then, the storage system copies the data blocks one by one to the target storage unit. After cloning is complete, the source and target storage units have the same data copy, typically used for data backup, migration, or testing scenarios. The advantage of data cloning is that backup copies can be created quickly without interfering with the original data.

The NexaVM N5000 storage array cloning adopts an incremental synchronization backup method, capable of generating a complete copy based on the source data at a specified point in time. Among them, “complete” means completely copying the source data to generate a data copy; “incremental synchronization” means the data copy can dynamically synchronize the changed parts of the source data.

The cloning function advantages of the NexaVM N5000 storage array are as follows:

- ▶ Complete data backup

Achieves complete data backup. Data recovery does not depend on the source data, providing reliable data protection.

- ▶ Continuous data protection

Source data and copies can be synchronized in real time, providing continuous protection and achieving zero data loss.

- ▶ Reliable business continuity

Backup and recovery processes can be performed online without interrupting business, achieving zero backup window.

4.1.11 Data Mirroring

Data mirroring is a data protection technology that achieves data redundancy and high availability by creating identical data copies on multiple storage devices.

The NexaVM N5000 storage array data mirroring supports creating two physical copies for one LUN, strengthening continuous redundant backup protection for the LUN without application-side perception.

The NexaVM N5000 storage array data mirroring requires that the two mirrored physical copies of a LUN be allocated in two storage pools. In this case, if one storage pool fails, the data on the LUN can still be guaranteed not to be lost and can work normally. The application side can work normally without perception. After the failed mirror physical copy recovers, it can automatically incrementally synchronize through the normal mirror copy. As shown in Figure 4-12, the data mirroring function includes one mirror LUN and two copy LUNs. The three LUNs have the same logical capacity, but the mirror LUN does not occupy actual storage space. The copy LUNs occupy the same amount in two storage pools respectively.

Figure 4-12 Data Mirroring

Text in figure: LUN; Copy LUN1; Copy LUN2; Storage Pool 1; Storage Pool 2.

The data mirroring read/write process and fault recovery process of the NexaVM N5000 storage array are as follows.

- ▶ Read/write process

When writing data, the server write is sent to the mirror LUN. The storage system data mirroring function module first checks whether the two copy LUNs are normal. If both copy LUNs are normal, data is written to both copy LUNs simultaneously. If both copy LUNs write successfully, the server is directly notified of successful writing. If one copy LUN write fails, the mirror copy difference position is recorded and the failed copy LUN is marked as faulty, and the server is notified of successful writing; if one copy LUN fails and the other copy LUN is normal, data is written to the normal copy LUN and the difference position between the two mirror copy LUNs is recorded. After successful writing, the server is notified of successful writing.

When reading data, the server read command is sent to the mirror LUN. The storage system data mirroring function module checks the status of the two copy LUNs, selects a copy LUN with complete

data to read data. If data reading is successful, the read data is returned to the server. If data reading fails, the current copy LUN is marked as faulty, and an attempt is made to read data through the other normal copy LUN.

► Fault recovery process

When a mirror copy LUN is marked as faulty, the copy LUN status check mechanism is immediately started. The storage system check mechanism periodically copies data from the normal copy LUN to the faulty copy LUN to check whether the faulty copy LUN can successfully write data. If data cannot be written normally, the copy is considered still in a faulty state. If the faulty copy LUN can write data normally, the copy LUN is considered to have recovered to normal. After the faulty copy LUN recovers to normal, the storage system will start the data synchronization mechanism to incrementally copy data from the normal copy LUN to the faulty copy LUN, restoring data consistency between the two copies as soon as possible.

Therefore, the NexaVM N5000 storage array data mirroring, through real-time creation of data copies, provides high availability and disaster recovery capabilities, suitable for various critical task and data protection scenarios, especially those with strict requirements for data consistency and business continuity.

4.1.12 LUN Copy

LUN copy functionality is a technology used for data migration and backup in storage arrays, allowing data from one LUN to be copied to another LUN, providing flexible data management, backup, and migration solutions.

During the LUN copy process, the source LUN is the original storage location of the data, and the target LUN is the new location where data will be copied. The LUN copy operation copies all data from the source LUN to the target LUN. Once the LUN copy is complete, the target LUN can continue to be used as a new storage unit, often used for data backup, disaster recovery, etc. System administrators can configure and monitor the LUN copy process through the N5000 storage array Web interface.

The NexaVM N5000 storage array LUN copy supports copying data from a LUN in one storage array to a LUN in another remote storage array, enabling fast data transmission between two devices, meeting users' needs for rapid data migration, data distribution, and centralized backup.

4.1.13 Remote Replication

Remote replication is a data protection technology designed to synchronize data from a local storage device to a remote storage device by establishing a replication relationship between a pair of LUNs

between two devices, ensuring data security and availability. This technology is mainly used for disaster recovery and 容灾 backup. By creating a replication relationship between the primary device and the remote device (usually a LUN pair), it achieves distributed backup of data across geographical locations. When the local storage fails or a disaster occurs, the remote storage device can provide data recovery services, ensuring business system continuity and no data loss. It has the following advantages:

- ▶ Improved data security

Remote replication can protect data from natural disasters, hardware failures, and other unexpected events, ensuring data durability and reliability.

- ▶ Enhanced disaster recovery capability

Once the local storage device fails or a disaster occurs, the remote storage device can quickly take over, ensuring continuous operation of the business system.

- ▶ Flexibility and scalability

Users can choose different replication modes according to actual needs, optimizing performance and storage resources.

The NexaVM N5000 storage array remote replication supports three modes: synchronous replication, manual replication, and scheduled replication.

Table 4-4 Remote Replication Mode Comparison

Mode	Description	Advantages	Application Scenarios
Synchronous Replication	Data is synchronized in real time between the primary storage device (primary LUN) and the remote storage device (secondary LUN), ensuring high data consistency between the two sites.	Synchronous replication can minimize data loss, especially suitable for application scenarios with extremely high data consistency requirements, such as finance and transactions.	Typically used in environments with relatively close geographical locations and high bandwidth, because synchronous replication requires high network bandwidth to ensure real-time data transmission.
Manual Replication	Users need to manually trigger data synchronization operations, i.e., manually synchronize primary LUN data to the secondary LUN when needed.	Suitable for scenarios that do not require real-time synchronization. Users can flexibly control when to perform data synchronization.	Suitable for data synchronization needs that can tolerate certain time delays, or for manual control of the replication process.
Scheduled Replication	The system automatically performs data synchronization according	Reduces manual intervention, ensuring data is synchronized at	Suitable for scenarios with regular data protection needs but no

	to a preset schedule.	regular intervals.	real-time requirements.
--	-----------------------	--------------------	-------------------------

Through remote replication, the NexaVM N5000 storage array can help enterprises achieve efficient data protection, disaster recovery, and business continuity, especially in environments with strict requirements for data consistency and reliability.

4.1.14 Data Compression

Data compression refers to the technology of reducing the size of stored data by using compression algorithms, thereby effectively saving storage space. This technology identifies and removes redundant information in data, so that the same amount of data occupies less physical space, thereby improving storage efficiency.

When performing data compression, the disk array uses compression algorithms to identify redundant parts in the data and replaces them with shorter representations, thereby reducing the required storage space; when reading data, the controller first decompresses the data, then provides it to the requesting application.

The benefits of data compression are as follows:

- ▶ Save storage space

Data compression can significantly reduce storage requirements, especially when storing large amounts of repetitive or redundant data. In this way, more data can be stored under the same hardware configuration, extending the service life of hard drives or storage media.

- ▶ Improve storage efficiency

By reducing storage space requirements, compression technology enables storage systems to utilize available resources more efficiently. For enterprises, this can reduce the cost of storage expansion.

- ▶ Reduce storage costs

With the implementation of data compression, enterprises do not need to frequently expand storage devices, thereby saving hardware procurement and maintenance costs, especially when processing large amounts of data, compression can reduce the required hardware investment.

- ▶ Improve data transmission efficiency

Compressed data transmission speed will be faster because the compressed data volume is smaller and the bandwidth required for transmission is lower. This is especially important for environments with limited network bandwidth.

After the NexaVM N5000 storage array enables data compression, the storage system automatically handles the compression and decompression process, ensuring data integrity and availability. The data

compression function of the NexaVM N5000 storage array is widely used in enterprise storage environments, especially in processing large amounts of documents, backup data, virtualization environments, and storage-intensive application scenarios, effectively improving the capacity and usage efficiency of storage devices.

4.1.15 Data Deduplication

Data deduplication is the abbreviation of “duplicate data deletion” technology, a technology that optimizes storage space usage by identifying and removing duplicate data. Its main purpose is to eliminate redundant data, storing only one actual data copy, and replacing other duplicate data with references pointing to that unique copy, thereby significantly saving storage space (especially in backup and archiving where large amounts of duplicate data exist). The benefits of data deduplication include:

- ▶ Save storage space

Reduces redundant data storage, greatly saving storage capacity.

- ▶ Improve storage efficiency

Effectively reduces the demand for storage devices, reducing the need for new hardware expansion.

- ▶ Reduce costs

By reducing the amount of data that needs to be stored, the purchase and maintenance costs of storage devices can be reduced.

Figure 4-13 Comparison of Data Write Process Without and With Duplicate Data Deletion Enabled

The NexaVM N5000 storage array duplicate data deletion is achieved by dividing data into fixed-size blocks and then checking for duplicates between these blocks. When data is written to the storage system, it is first divided into small blocks. Each data block calculates a unique hash value (fingerprint information). The system compares this hash value with the hash values of existing data blocks. If the same hash value is found, it indicates that this data block is duplicate. The system then only stores one copy of this data block, and other duplicate blocks only retain the hash value, avoiding storing multiple copies of the same data.

4.1.16 Automatic Thin Provisioning

Automatic thin provisioning is a technology for dynamically managing storage resources, aimed at improving the efficiency of storage space usage. Under this configuration, the storage system does not allocate all fixed, preset storage space for applications at once, but allocates storage capacity on demand according to actual needs. Specifically, automatic thin provisioning creates “thin LUNs,” allowing

applications to only see the storage space already allocated to them, but this space may actually be much smaller than the total physical storage capacity of the system.

Initially, the system only allocates a small, physically existing storage space for the application. As the application's demand for storage space increases, the system dynamically supplements more storage space from the "thin pool" (i.e., storage resource pool). The space in the thin pool is reserved in idle physical resources and can be automatically provided when needed.

Figure 4-14 Automatic Thin Provisioning Working Principle

Text in figure: Host; Write operation; Thin LUN; Actual storage space; RAID5; Allocate when used.

The advantage of automatic thin provisioning is that it avoids allocating too much storage capacity in advance for each application, thereby saving storage resources and improving storage utilization. It also allows flexible adjustment of storage capacity. As application requirements change, the system can automatically allocate space, ensuring sufficient storage space is provided without wasting too many idle resources.

The tiered storage pool of the NexaVM N5000 storage array supports both thin LUNs and non-thin LUNs. Non-thin LUNs allocate all storage space according to user-configured capacity when created, improving storage resource usage efficiency and meeting business performance requirements to a greater extent. Thin LUNs do not pre-allocate storage space when created, but present the user-configured capacity form to the user, so that the storage space seen by the user equals the user-configured storage space. The capacity configured for thin LUNs is logical space, and the actual occupied space is actual storage space.

Thin LUNs are allocated on demand during use. Users are allocated as much as they use. Unused space in the storage pool can be allocated to any LUN that needs storage space. Because thin LUNs do not allocate space when created, users are allowed to create LUNs exceeding the storage pool capacity. If the data capacity exceeds the storage pool capacity, the storage pool capacity can be dynamically expanded.

Because the actual storage space of thin LUNs is not fully allocated, when reading and writing data, it is necessary to first query whether the LUN's read/write location has already allocated storage space.

Different processing methods are used based on the query results.

- ▶ Data reading

After a thin LUN receives a server read IO request, if no allocated storage space is found, the area is set to all zeros and returned to the server. This process only returns all-zero data to the server and does not allocate storage space; if allocated storage space is found, the data at that storage location is read and returned to the server.

- ▶ Data writing

After a thin LUN receives a server write IO request, if no allocated storage space is found, storage space is first allocated from the storage pool, with a minimum allocation of 4MB each time. Then data is written to the allocated storage space, and the relationship between the LUN's logical blocks and physical blocks is recorded; if allocated storage space is found, the data at that storage location is read and returned to the server.

The automatic thin provisioning of the NexaVM N5000 storage array is widely used in scenarios requiring flexible and efficient management of storage resources, especially in virtualization environments, cloud storage, big data storage, and database management. In these scenarios, storage requirements are often dynamically changing. Traditional pre-allocated storage space methods may lead to resource waste or shortage. Through automatic thin provisioning, enterprises can dynamically allocate storage capacity according to actual needs, avoiding excessive advance reservation of storage space, thereby saving storage costs while ensuring system efficiency.

4.1.17 NFS Permission Management

NFS (Network File System) is a network protocol that allows computers with different operating systems to share files and directories through the network. NFS, by providing a simple and efficient file sharing mechanism, enables different systems to easily share files and access them remotely, and is one of the indispensable protocols in modern network storage systems.

NFS permission management refers to controlling access permissions to NFS shared files and directories, ensuring that only authorized users or computers can access, modify, or manage shared file systems. NFS permission management mainly relies on host, user, and group identity authentication and access control mechanisms to ensure data security and integrity.

NFS has multiple versions, each version improving and enhancing on the previous version, providing higher performance, security, and functionality. The differences between common NFS versions are shown in Table 4-5. Different versions of the NFS protocol have different emphases in performance, security, scalability, and management functions. Enterprises can choose the appropriate version according to specific needs to implement network storage.

Table 4-5 Common NFS Version Comparison

NFS Version	Description	Characteristics
NFSv2	NFSv2 is the initial version, released in 1989. It is the first standardized version of the NFS protocol.	Simple protocol: NFSv2 uses a simple RPC (Remote Procedure Call) mechanism, transmitting data based on UDP (User Datagram Protocol). Stateless protocol: NFSv2 is stateless, meaning the server

		does not maintain client state information. Each request is independent. 32-bit file size limit: Maximum supported file size is 2GB, limiting NFSv2's ability to handle large files. Performance: Because it relies on UDP protocol, NFSv2 does not provide reliability guarantees and is susceptible to packet loss and network failures.
NFSv3	Released in 1995, it is an improved version of NFSv2.	Supports large files: NFSv3 removes the 2GB file size limit of NFSv2, supporting maximum file sizes of 64GB. Improved transport protocol: NFSv3 still uses UDP protocol but also supports TCP (Transmission Control Protocol), providing more reliable connections, especially in unstable network environments. Supports asynchronous write operations: Introduces asynchronous write operations (Write-Behind), reducing client waiting time and improving performance. Introduced statefulness: Although still a stateless protocol, NFSv3 introduces file locking mechanisms and partial state management, making it more reliable in concurrent access. More efficient performance: NFSv3 provides higher performance, especially in read/write operations and large file processing.
NFSv4	Released in 2000, it provides stronger security, performance improvements, and richer features. NFSv4 integrates the features of multiple versions and adds many new functions, improving the stability and availability of the NFS protocol.	Security: NFSv4 mandates the use of TCP, avoiding data loss that may be caused by UDP transmission. Stateful protocol: NFSv4 becomes a stateful protocol. The server can remember client session information, thereby providing more efficient fault recovery and lock management. Built-in security mechanisms: NFSv4 has built-in

		support for Kerberos authentication, client identity verification, and encryption, enhancing data security. Supports file locking: NFSv4 supports stronger file locking mechanisms, avoiding data conflicts during concurrent file access. Path names and file handles: Improved path name resolution and file handle management, making directory tree operations more efficient. Cross-platform support: NFSv4 enhances support for Windows and Unix/Linux systems, improving cross-platform interoperability. Delegation mechanism: Supports client caching mechanisms, reducing network latency through the delegation model.
NFSv4.1	An extended version of NFSv4, released in 2010, providing more optimized performance and enhanced functionality, mainly focusing on multipath operations and performance improvement.	Multipath I/O support: NFSv4.1 introduces multipath I/O (Multipath I/O) mechanism, supporting access to shared storage through multiple network paths, improving network reliability and performance. Distributed file system support: Enhanced support for distributed file systems, improving data management and storage efficiency. Asynchronous operation improvements: NFSv4.1 optimizes asynchronous operations, reducing client waiting time and further improving performance. Enhanced connectivity: Through enhanced session support, improves connection stability between clients and servers, especially performing better in high-latency networks.

The NexaVM N5000 storage array supports NFSv2, NFSv3, NFSv4, and NFSv4.1 versions. Users can use specific versions of the NFS protocol according to their needs. To ensure data security, the NexaVM N5000 storage array supports NFS multiple permission management, including “access network segment

control,” “access read/write permission control,” “whether ordinary users are anonymous control,” and “whether root users are anonymous control.” Among them, “access network segment control” can restrict the IP or network segment of clients allowed to access, restricting other network segments or IPs from accessing NFS; “access read/write permission control” can restrict certain clients to access with read/write permissions, while other clients can only access with read-only permissions; “whether ordinary users are anonymous control” is used to control the user identity information for accessing NFS. If configured as non-anonymous access, NFS uses the client’s original user ID and user group ID to check shared file permissions. Otherwise, users use anonymous user ID and user group ID to check shared file permissions; “whether root users are anonymous control” can restrict whether root users access anonymously. If configured as non-anonymous access, NFS uses the root user to check shared file permissions. Otherwise, NFS uses an anonymous user to check shared file permissions.

4.1.18 Storage Configuration Management

The NexaVM N5000 storage array supports multiple configuration management methods such as Web interface, command line, and SNMP, and supports providing external access interfaces such as WebService and Restful, which can conveniently and flexibly integrate with third-party storage management platforms or business systems.

The NexaVM N5000 storage array Web interface provides a bilingual Chinese-English graphical management interface. Users can manage all functions and features of SAN and NAS simultaneously through the same interface. The interface is clear and intuitive, providing detailed descriptions for system use and operation. Function operations are accompanied by help and comprehensive prompt information, enabling users to easily get started. In addition, the NexaVM N5000 storage array Web interface displays array configuration, status, mapping relationships, and performance monitoring information through charts and lists, facilitating users to fully understand the operating status of the storage system. The NexaVM N5000 storage array Web interface also records rich log information, including system logs and operation logs. These logs can be notified to third-party monitoring platforms in real time through email, SNMP, etc., ensuring timely management and monitoring of the system.

The NexaVM N5000 storage array supports configuring and managing the basic functions of the storage system through serial port and Telnet command lines, meeting users’ basic business configuration needs.

The NexaVM N5000 storage array also supports SNMPv2 and SNMPv3 protocols. Ports, LUNs, storage pools, and other device information can be viewed through SNMP, as well as monitoring system log information.

In addition, the NexaVM N5000 storage array provides interfaces including WebService and Restful, enabling the storage system to be centrally managed by third-party management platforms.

4.1.19 Out-of-Band Management

Out-of-band management refers to monitoring and managing storage devices through a management channel independent of the main system operation, not relying on the operating state of the operating system or applications. Out-of-band management enables system administrators to remotely access and manage storage arrays even when the main system fails or the operating system cannot respond. This management method provides higher reliability and flexibility, especially when dealing with hardware failures or emergency events. The advantages of out-of-band management include:

- ▶ High availability: Even if the operating system crashes or is inaccessible, administrators can still perform management operations.
- ▶ Remote management: Administrators can remotely access and control devices through the network without being on-site.
- ▶ Fault diagnosis and repair: BMC provides hardware health monitoring, logging, and emergency management functions, facilitating timely measures when system failures occur.
- ▶ Intelligent adjustment: Based on monitoring results, BMC can automatically adjust device working status, reducing manual intervention.

The NexaVM N5000 storage array uses BMC (Baseboard Management Controller) as the core component of out-of-band management. BMC is an independent hardware management system that does not depend on the device's CPU, memory, and other hardware, nor does it depend on BIOS or the storage operating system. When the NexaVM N5000 storage array fails, BMC can restart the storage system through reset operations, ensuring device availability. The BMC of the NexaVM N5000 storage array can monitor multiple key hardware parameters such as device temperature, voltage, fans, and power supplies. Based on these monitoring data, users can intelligently adjust system working status to ensure healthy device operation. Users can also export hardware information and log records through BMC. This information is crucial for subsequent problem location and troubleshooting. The BMC of the NexaVM N5000 storage array uses the IPMI (Intelligent Platform Management Interface) protocol as the management interface. The autonomous characteristics of IPMI enable BMC to continue performing management operations such as power on/off and information extraction even when the operating system is unresponsive or not loaded, overcoming the limitations of traditional in-band methods. During operation, BMC receives and processes system event messages based on the IPMI specification and maintains device sensor records, describing the status of various hardware in the system. When remote access to the system is needed, BMC supports LUN-based Serial Over LAN (SOL), which can change the local serial port transmission direction during IPMI sessions, providing administrators with remote console access. This enables administrators to remotely view the system startup process, operating system loader, or enter the emergency event management console, thereby diagnosing and repairing

device-related problems and configuring system boot phase components. Therefore, the BMC of the NexaVM N5000 storage array, as an out-of-band management tool, provides comprehensive remote management support for the storage system, ensuring system stability, maintainability, and high availability.

4.1.20 Power Failure Protection

Power failure protection refers to protecting data in the storage array from loss and ensuring data consistency and security through specific hardware and software mechanisms in the event of power failure or system power loss. It has the following benefits:

- ▶ Data consistency assurance

Avoids cache data loss or corruption caused by power interruption, ensuring data consistency after power failure.

- ▶ Improve system reliability

Even if a power failure occurs, the storage array can still safely save data, reducing the complexity of fault recovery.

- ▶ Enhance enterprise disaster recovery capability

Power failure protection provides enterprises with higher reliability and disaster recovery capability, especially in critical business systems, ensuring that important data will not be lost due to power interruption.

- ▶ Reduce data recovery costs

Because data is protected after power failure, the recovery workload and cost caused by data loss or corruption are reduced.

The NexaVM N5000 storage array's power failure protection is implemented based on the BBU battery backup unit, which can provide temporary power support for the storage array's cache and controller when the main power fails. When power is interrupted, the BBU provides power to maintain cache data in the array, ensuring that all unfinished write operations can be saved to persistent storage media (such as hard drives or SSDs). Once power is restored, the BBU can write the data in the cache to disk, ensuring data integrity.

Figure 4-15 BBU Working Principle Schematic Diagram

Text in figure: 220V AC Input 1; 220V AC Input 2; Power Module; Power Status Report; Mainboard CPLD; Power Monitoring; BBU; Charge/Discharge Circuit; Charge/Discharge Control; DC Output; 12V; 5V; 3.3V; 5VSTB.

The NexaVM N5000 storage array supports a built-in BBU module (backup power). The BBU module is located on the controller, meaning that if the controller is accidentally pulled out, the BBU module will also be pulled out along with the controller. After the software module on the controller detects that the controller has been pulled out, it can use the BBU module (backup power) on the controller to back up user cache data to the cache, ensuring data integrity. The system's power-off cache flush process is completed by the underlying system, not relying on upper-layer software units, ensuring that the cache flush process is not affected by business, further improving the reliability of user data.

4.1.21 One-Click Destruction

One-Click Destruction is a data security function designed to quickly and thoroughly destroy all data in the storage array. This function is typically used in the following situations:

- ▶ Data privacy protection: Before equipment is scrapped, sold, or transferred, ensure that sensitive data in storage devices is thoroughly deleted to prevent data leakage.
- ▶ End of equipment lifecycle: When storage devices reach the end of their service life or are no longer needed, all data can be safely deleted through the one-click destruction function.
- ▶ Compliance requirements: Meet enterprise or industry compliance requirements for data destruction, ensuring that the data destruction process meets specified standards.

The NexaVM N5000 storage array has a one-click destruction function, which can destroy all configuration information and data information of the storage array with one click. The one-click destruction of the NexaVM N5000 storage array includes fast destruction and secure destruction functions. Among them, fast destruction deletes configuration information and hard drive management data; secure destruction not only deletes configuration information and hard drive management data but also writes zeros to all hard drives, ensuring data cannot be recovered.

4.1.22 Quality of Service Control

Quality of Service control (QoS) is a technology used to optimize storage resource allocation and management, aimed at ensuring that different applications and workloads receive priority processing and reasonable allocation in the storage system. The main functions and characteristics of QoS are:

- ▶ Performance assurance

QoS ensures that the performance requirements of specific applications or workloads are met by setting rules and policies. For example, some critical applications may require lower latency or higher throughput. QoS can prioritize the performance of these applications, avoiding the impact of other application traffic.

- ▶ Bandwidth control

By setting bandwidth limits for different LUNs or storage volumes, ensure that high-priority tasks are not affected by low-priority tasks, ensuring the performance of important businesses.

- ▶ Latency control

By controlling the data access latency of different applications, optimize the response time of the storage array, reducing the impact of network congestion or storage device performance bottlenecks on applications.

- ▶ IOPS limitation

According to actual needs, set maximum IOPS for different workloads, preventing a task from occupying too many IO resources and affecting the execution of other tasks.

The NexaVM N5000 storage array QoS mechanism guarantees the needs of high-performance applications through I/O traffic control policies and I/O priority scheduling policies.

- ▶ I/O Traffic Control Policy

The I/O traffic control policy is based on the traditional token bucket mechanism. By limiting the IOPS or bandwidth of a certain LUN or file system, it limits the upper limit of non-critical business traffic, avoiding excessive burst traffic from these applications, which would cause critical business performance requirements to be unmet. According to the direction of I/O, it is divided into read I/O, write I/O, and read/write I/O control. According to the controlled resources, it is divided into IOPS control and bandwidth control. The main implementation mechanism of the control policy is to increase I/O delay time, thereby reducing non-critical business traffic.

- ▶ I/O Priority Scheduling Policy

The I/O priority scheduling policy is based on the priority of LUNs and file systems to divide queues. Critical businesses have higher priority, and non-critical businesses have lower priority. When resources are tight, high-priority businesses are guaranteed first.

Users can configure LUNs and file systems as high, medium, and low levels. The system includes three I/O queues: high, medium, and low. When accessing a LUN or file system with high priority, I/O enters the high-priority I/O queue. When accessing a LUN or file system with low priority, I/O enters the low-priority queue. When system resources are tight, I/O in the high-priority queue is processed first, thereby ensuring high-priority businesses as much as possible.

Therefore, through QoS control policies, the NexaVM N5000 storage array can efficiently manage storage resources, ensuring the performance requirements of critical tasks while optimizing bandwidth and IOPS allocation. The QoS mechanism not only ensures the efficiency and reliability of the storage system but also prioritizes the operation of important businesses when resources are tight, providing stable storage support for enterprise applications.

4.1.23 Heterogeneous Data Migration

Heterogeneous data migration of storage arrays refers to the process of migrating data between different types of storage devices or storage arrays. “Heterogeneous” refers to systems involving different vendors, different hardware architectures, different storage protocols, or different storage technologies. For example, migrating data from an old hard disk array to a new flash storage array, or migrating data from one storage vendor’s device to another vendor’s device. The main characteristics of heterogeneous data migration:

- ▶ Cross-device migration

Data can be migrated from different types of storage devices (such as traditional hard drives, solid-state drives, cloud storage, etc.) to other types of storage devices, whether the same hardware type (such as SAS hard drive to SAS hard drive) or cross-hardware (such as SATA hard drive to SSD).

- ▶ Cross-vendor migration

Data migration can span different vendors’ storage arrays or systems, involving different storage protocols (such as iSCSI, Fibre Channel, NFS, etc.), different RAID architectures or different operating systems, and even different file systems.

- ▶ Seamless migration

Heterogeneous data migration usually needs to be seamless, i.e., not interrupting normal business operations during the migration process, ensuring data integrity and availability.

The NexaVM N5000 storage array supports intelligent heterogeneous data migration, which can fully migrate data from the source LUN to the target LUN without interrupting the original business, achieving data migration without business perception. Data migration not only supports local data migration but also supports data migration between heterogeneous systems.

Data migration achieves complete copying of source LUN data to the target LUN and uses the target LUN to replace the source LUN after replication is complete. Data migration includes two steps: business data synchronization and inheriting source LUN identification information.

- ▶ Business data synchronization

All data in the source LUN is fully copied to the target LUN. During the migration process, if source LUN data changes, the modified data is also copied to the target LUN.

- ▶ Inherit source LUN identification information

After data migration is complete, the target LUN will inherit the source LUN’s WWN and other information, ensuring that server business is not perceived.

Therefore, the heterogeneous data migration technology of the NexaVM N5000 storage array not only optimizes storage resource configuration but also provides enterprises with a more flexible, reliable, and seamless data migration method. By achieving business-unaware migration, enterprises can reduce maintenance costs and improve storage system performance and efficiency while ensuring data consistency and high system availability.

4.2 Software Specifications

The main software specifications of the NexaVM N5000 storage array are shown in Table 4-6.

Table 4-6 NexaVM N5000 Storage Array Key Software Features

Specification Name	Detailed Parameters
RAID Levels	RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60, Linear RAID
System Compatibility	UNIX, Linux, Windows, etc.
Virtualization Support	Supports VMWare, Hyper-V, NexaSphere, KVM
OpenStack Support	Supports docking with OpenStack cloud computing platform, with OpenStack uniformly managing storage resources
Number of Supported LUNs	4096
Management Software	Has a Web-based graphical management interface, supports unified management, supports operation logs, fault statistics, email alarms, performance monitoring, system configuration export and import, system diagnosis
Efficiency	Storage pool reconstruction can intelligently schedule reconstruction rate according to business busyness
Resource Efficiency Improvement	Intelligent RAID verification, RAID automatic reconstruction, RAID online expansion, media scanning, data integrity detection, disk positioning
Active-Active Controller Architecture	High-speed channel data exchange between dual controllers, seamless request switching, greatly reduced latency
Intelligent Hard Drive Management	Supports intelligent disk hibernation
Cache Acceleration	Supports SSD as cache, configurable SSD hard drive ratio
Data Protection	Data replication, file migration, data mirroring, automatic thin provisioning, snapshot, SSD acceleration, compression, deduplication, automatic tiering, local logical volume mirroring

Hardware	...
----------	-----

5

CHAPTER

Abbreviations and Terms

Reference glossary of acronyms and technical terms.

No.	Abbreviation/Term	Explanation
1	ASIC	Application-Specific Integrated Circuit, an integrated circuit designed specifically for a particular application. Unlike general-purpose integrated circuits (such as CPU, GPU, etc.), ASIC is a circuit optimized for performing specific tasks or solving specific problems. Its design is usually highly specialized, so its efficiency in that task far exceeds general-purpose processors.
2	BBU	Battery Backup Unit, typically used in storage devices to provide short-term battery power support for the system to prevent data loss caused by sudden power failure.
3	BMC	Baseboard Management Controller, an embedded hardware used to monitor server hardware health status, temperature, fan speed, etc., supporting remote management.
4	CIFS	Common Internet File System, a network file sharing protocol that allows applications to access resources such as files and printers through the network.
5	CPU	Central Processing Unit, one of the core components of a computer, responsible for executing instructions in computer programs, performing arithmetic and logical operations, and input/output processing of data.
6	FC	Fibre Channel, a high-speed data transmission protocol used in storage area networks (SAN), designed specifically for high-speed communication between data storage and servers.
7	FTP	File Transfer Protocol, a standard protocol for transferring files

		between computer networks, widely used for uploading and downloading files.
8	HTTPS	Hypertext Transfer Protocol Secure, the encrypted version of HTTP, encrypting data through SSL/TLS protocols to ensure confidentiality and integrity of data transmission.
9	ID	Identifier, used to uniquely identify objects, devices, users, and other entities.
10	IO	Input/Output, refers to the data exchange and processing process between computers or other devices, which can be inputting data from external devices or outputting data to external devices.
11	IOPS	Input/Output Operations Per Second, used to measure the number of read/write operations a storage device (such as a hard drive or SSD) can process per unit time, reflecting the performance of the storage device.
12	IP	Internet Protocol, the basic protocol used for packet switching in the Internet, ensuring data can be transmitted from source devices to target devices.
13	IPMI	Intelligent Platform Management Interface, a server management interface standard that allows administrators to remotely monitor, manage, and repair server hardware.
14	iSCSI	Internet Small Computer Systems Interface, a storage area network (SAN) protocol based on IP networks, used to connect storage devices with servers through Ethernet for data access.
15	LUN	Logical Unit Number, an identifier for logical partitions in storage devices, typically used in SAN

		(Storage Area Network) to represent a virtual disk device.
16	NAS	Network-Attached Storage, a solution for sharing file-level storage through the network. Storage devices connect to clients (usually computers or servers) through standard networks (such as Ethernet). Unlike SAN, NAS provides file-level access through network file protocols (such as NFS or SMB/CIFS).
17	NFS	Network File System, a network protocol that allows computers with different operating systems to share files and directories through the network.
18	NVMe	Non-Volatile Memory Express, a high-performance protocol for flash (SSD) devices connected through PCIe interfaces, providing low latency and high throughput.
19	OLAP	On-Line Analytical Processing, a data analysis technology designed to support complex queries and multi-dimensional analysis, widely used in data warehouses and decision support systems (DSS).
20	OLTP	On-Line Transaction Processing, a system type used for processing and managing large amounts of transaction data, typically used in business environments requiring real-time data processing, such as banking, retail, e-commerce, and online payment systems.
21	PCIE	Peripheral Component Interconnect Express, a high-speed data transmission bus standard used to connect computer internal components (such as graphics cards, network cards, storage devices, etc.).
22	QoS	Quality of Service, a network management technology that

		optimizes network performance through data traffic control, ensuring high-priority traffic receives more bandwidth, low latency, and other resources.
23	RAID	Redundant Array of Independent Disks, a technology that combines multiple hard drives, using different data distribution and redundancy strategies to improve storage performance and data protection.
24	SAN	Storage Area Network, a high-speed network specifically used to connect storage devices (such as disk arrays) with servers. SAN provides block-level storage, meaning storage devices directly provide storage volumes (LUNs) through the network for servers to access.
25	SAS	Serial Attached SCSI, a high-performance serial interface standard mainly used to connect computer storage devices such as hard disk drives and solid-state drives.
26	SATA	Serial Advanced Technology Attachment, a computer bus interface standard for connecting hard disk drives and solid-state drives (SSDs), more widely used than SCSI.
27	SNMP	Simple Network Management Protocol, a protocol used for network management, allowing administrators to monitor, configure, and manage network devices.
28	SOL	Serial Over LAN, serial communication over LAN, used for remote server management, accessing and managing server serial interfaces through network connections.

29	SSD	Solid State Drive, a hard drive using flash memory storage, with faster data read/write speeds and better shock resistance than traditional mechanical hard drives (HDD).
30	TCP	Transmission Control Protocol, a connection-oriented protocol that ensures reliable, sequential data transmission in networks, widely used in HTTP, FTP, and other protocols.
31	TOE	TCP/IP Offload Engine, a network adapter technology used to offload TCP/IP protocol stack processing tasks from the host computer's CPU to the network card, thereby improving data transmission efficiency.
32	UDP	User Datagram Protocol, a connectionless transport layer protocol that provides simple, fast data transmission but does not guarantee data reliability.
33	Web	World Wide Web, a hypertext-based Internet application that allows access to various content such as web pages, videos, and images through browsers.
34	WWN	World Wide Name, a unique identifier used to identify Fibre Channel (FC) devices or LUNs, similar to a MAC address but used for storage devices.

Appendix A. Hardware Specifications (Reference Table)

Detailed technical specifications of the N5000 storage array, presented as a reference for hardware evaluation and procurement.

Specification	Description
Processor	Multi-core processor
Number of Controllers	Dual controllers, expandable up to 16 controllers
Height/Bays	2U/ 12 bays
System Cache	Dual-controller 32GB
Storage Methods	Supports NAS/FC SAN/IP SAN storage, supports unified block storage and file storage
Supported Storage Protocols	FC、iSCSI、NFS、CIFS、FTP、HTTPS
Interface Types	2*management ports + 2*serial ports + 2*GE; provides 4 interface card slots, optional interface cards: 1/10/40GbE, 8/16/32Gb FC, 10/40Gb FCoE, including storage-specific chips;
Disks	SAS、NLSAS、SATA、SSD
Maximum disks	2048
Hot backup disk	Support local/global hot backup disk and dedicated hot backup disk
Disk status indicator	Disk errors (dual LED)
Power and fan module	2 for each
PCIe expansion slots	4 (dual controller)
dimension	447(L)*657.5(W)*88(H)mm
humidity	10% ~ 85%
Temperature	0°C ~ 45°C



Thank you for reading.

For more information, contact your NexaVM account team.

© 2026 NexaVM. All rights reserved.