

Panasas ActiveStor Director 200

High-Performance Metadata Director Node

The ActiveStor® Director 200 (ASD-200) is Panasas® metadata director node and the control plane for the ActiveStor high-performance data storage solution. The ASD-200 is built on industry-standard hardware chosen for its carefully balanced hardware architecture that enables independent scaling and to meet the requirements of metadata-intensive applications in manufacturing, life sciences, energy, financial services, academia, and government.

ASD-200 director nodes are powered by PanFS®, the Panasas parallel file system, and control many aspects of the overall storage solution including file system semantics, namespace management, distribution and consistency of user data on storage nodes, system health, failure recovery, and gateway functionality. Together with ActiveStor storage nodes and the DirectFlow® driver on client systems, PanFS delivers the highest performance with unlimited scalability, enterprise reliability, and ease of management. As the system scales, reliability and availability increase while administrative overhead remains low.

ASD-200 Enclosure

The ASD-200 director enclosure is a 2U, 19" rackmount, 4-node chassis, as shown in Figure 1. Each enclosure can house up to four ASD-200 director nodes and includes redundant power supplies, as shown in Figure 2.

ASD-200 Node

The ASD-200 director node is a server node running the PanFS parallel file system. The node has been selected for its form factor and overall quality and reliability. It has been configured and tested with focus on the strength of the CPU, DRAM capacity, and networking bandwidth.

Scalable Metadata Services

ASD-200 director nodes manage system activity and provide clustered metadata services. The nodes orchestrate file system activity and speed data transfers while facilitating scalability and virtualizing data objects across all available storage nodes. This enables the system to be viewed as a single, easily-managed global namespace.



Figure 1. ASD-200 enclosure, front view.



Figure 2. ASD-200 enclosure and nodes, rear view.

PanFS metadata services running on ASD-200 nodes implement all file system semantics and manage sharding of data across the storage nodes. They control distributed file system operations such as file-level and object-level metadata consistency, client cache coherency, recoverability from interruptions to client I/O, storage node operations, and secure multiuser access to files. Storage administrators can easily create volumes within the PanFS global namespace to manage hierarchies of directories and files that share common pools of storage capacity. Per-user capacity quotas can be defined at the volume level. Each volume has a set of management services to govern the quotas and snapshots for that volume. This type of partitioning allows for easy linear scaling of metadata performance.

Superior Manageability

A single point of management for a scale-out file system allows the storage administrator to focus on core business tasks instead of the storage system. Panasas easily addresses capacity and performance planning, mount point management, and data load balancing across multiple pools of storage. ASD-200 director nodes easily integrate into growing heterogeneous environments through the high-performance DirectFlow protocol support for Linux and multi-protocol support for NFS and SMB.

Gateway Services

ASD-200 directors provide scalable access for client systems via NFS or SMB protocol “gateway” services. Director nodes do this without being in the data path. Using these gateway solutions, users can easily manage files created by Windows environments. User authentication is managed via a variety of options including Active Directory and Lightweight Directory Access Protocol (LDAP).

File-level Reconstruction

Data protection in the PanFS operating environment is calculated on a per-file basis rather than per-drive or within a RAID group, as in other architectures. ASD-200 directors also provide an additional layer of data protection called Extended File System Availability (EFSA) for the namespace, directory hierarchy, and file names. In the extremely unlikely event of encountering errors that erasure coding cannot recover from, the system knows which files have been affected and fenced off and which files are known to not have been impacted.

High Availability

All metadata transactions are journaled on a backup director node. All volumes remain online in case of failover, with no required system check. Network failover ensures there is no single point of failure in the system network. All director nodes share the reconstruction workload and enable load balancing during reconstruction, providing fast reconstruction within hours rather than days.

Automatic data rebuilding protects against system-wide failures. Redundant networking data paths automatically fail over. All components are hot swapped for easy field servicing. EFSA takes advantage of erasure coding of directory data to preserve file system integrity and accessibility.

Timely, High-Quality Service and Support

Unlike open-source solutions and even commercial alternatives from broad portfolio vendors, Panasas offers timely, world-class L1–L4 support.

ASD-200 Specifications

ASD-200 Enclosure	
Hardware	19" rackmount chassis with rails
Power Supplies	2x 2200 W titanium-level
Height	3.46 inches (88 mm)—2 rack units
Width	17.24 inches (438 mm)
Depth	29.92 inches (760 mm)
Operating Temp.	0–35°C (32–95°F)
Non-operating Temp.	-40–60°C (-40–140°F)
Operating Humidity	8–90% (non-condensing)
Input Line Voltage	220–240 VAC, 50–60 Hz

ASD-200 Node	
Storage Capacity	U.2 NVMe SSDs: 23–46 TB M.2 NVMe SSD: 3.84 TB
Memory	4x 16 GB DDR4 ECC RDIMMs
NVDIMM	1x 16 GB DDR4 ECC NVDIMM-N
M.2 NVMe	1x M.2 NVMe SSD
U.2 NVMe	6x 2.5 inch Enterprise NVMe SED
NIC	25 GbE Dual SFP28 Network SIOM
Other	Integrated BMC, IPMI, VGA, USB

More Information and Ordering Details

For more information and ASD-200 ordering details, contact your local Panasas representative or visit panasas.com/products/activestor-director.

About Panasas



Panasas builds a portfolio of data solutions that deliver exceptional performance, unlimited scalability, and unparalleled reliability – all at the best total cost of ownership and lowest administrative overhead. The Panasas data engine accelerates AI and high performance applications in manufacturing, life sciences, energy, media, financial services, and government. The company’s flagship PanFS® data engine and ActiveStor® storage solutions uniquely combine extreme performance, scalability, and security with the reliability and simplicity of a self-managed, self-healing architecture. The Panasas data engine solves the world’s most challenging problems: curing diseases, designing the next jetliner, creating mind-blowing visual effects, and using AI to predict new possibilities.

Worldwide Office
1-888-PANASAS
info@panasas.com

Panasas Headquarters
San Jose, CA, USA
Panasas Research & Development
Pittsburgh, PA, USA

Panasas EMEA
Oxford, United Kingdom
emeainfo@panasas.com

Panasas APAC
Sydney, Australia
apacinfo@panasas.com

Panasas China
Shanghai, China
chinainfo@panasas.com