


RAIDIX 5.X

Software-defined storage
for data intensive tasks
and environments



Key Features

- 5M IOPS in all-flash configuration*
- 22 GBps in hybrid configuration
- Performance penalty in degraded RAID mode less than 10%**
- Fault tolerant HA cluster
- NAS protocols: NFS, SMB, AFP, FTP, NFS-over-RDMA
- SAN protocols: FC, IB-SRP, iSCSI, iSER, SAS, NVMe-over-TCP, NVMe-over-RDMA

*block size 4K, with ERA Engine
**for sequential read

RAIDIX 5.X enables you to deploy all-flash and hybrid storage with high-speed block (SAN) and file (NAS) access. Empowered by innovative RAID technologies RAIDIX storage provides stable performance for data intensive tasks and keeps integrity of your business workflow.

Benefits

Ultra-high bandwidth for the most demanding tasks

Internal logic of RAIDIX storage arrays was originally developed to get maximum benefits for working with massive sequential workloads. It significantly boosts HDDs performance even in low-end storage systems.

To interoperate with solid state drives, RAIDIX storage uses ERA Engine, an innovative software RAID with I/O handling parallelization and lockless datapath. These features raise up storage performance to the level of flash potential bandwidth. Moreover, the system shows high IOPS and low latency even in mixed I/O workloads like VDI and financial transaction processing.

The lowest impact of hardware failures

Driven by innovative RAID technologies, RAIDIX storage demonstrates minimum performance penalty when a drive goes out. It helps to maintain smooth application workflow and data integrity in case of emergency.

RAIDIX has dual controller configuration to eliminate single point of failure and prevent impact from the most common hardware failures. To protect data from natural and man-driven disasters, RAIDIX can create synchronized copies of data in a secondary data center.

Easy integration into existing infrastructure

*file and block protocols cannot be provided at the same time

RAIDIX storage provides access via file (NFS, SMB, AFP, FTP, NFS-over-RDMA) or block (FC, IB, iSCSI, iSER, NVMe-over-TCP, NVMe-over-RDMA) protocols*.

To provide more hardware freedom, RAIDIX data storage works with any drive from any vendor. Software-driven architecture eliminates storage vendor lock-in and minimizes costs of regular hardware modernizations.

Ready for non-standard projects

Due to simple integration with parallel file systems, RAIDIX data storage can be easily adjusted to outsized and non-trivial installations. Depending on the file system features, storage will get additional scalability, ability to work with immense data volumes or gain performance in multi-node projects.

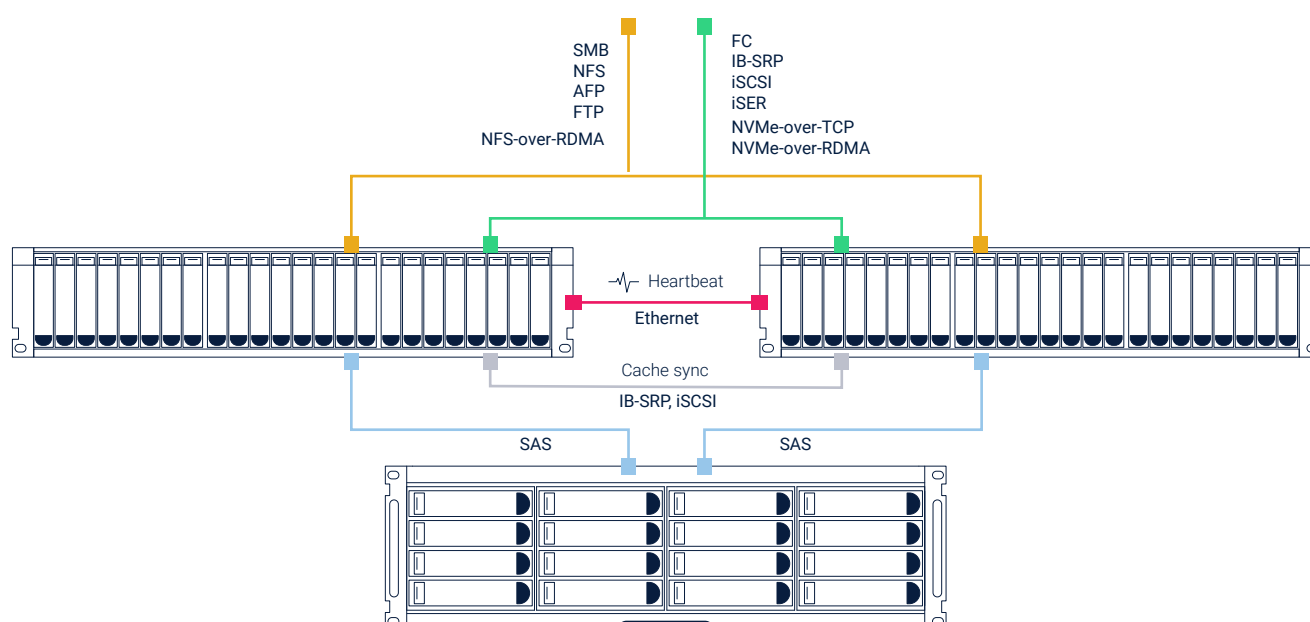
RAIDIX + Lustre. Collaborating with Lustre FS, RAIDIX storage meets the needs of data storage for HPC clusters.

Features

HA cluster mode


RAIDIX data storage can be deployed both as a single and as a dual controller system. To prevent hardware failures RAIDIX uses dual controller configuration as a high-availability (HA) cluster. Doubling hardware platform is the first step to reach an enterprise-grade level of data protection.

The system controllers interact via Ethernet, InfiniBand, iSCSI, in order to synchronize the cache and status of controllers.



The fastest RAID rebuild

The core technology of RAIDIX storage software is the ultra fast RAID engine implemented in the operating system kernel. It is based on unique checksum calculation approach with specific data allocation on CPU vector register.



High checksum calculation speed (25GBps per CPU core) helps RAIDIX to keep stable performance level and decrease system reconstruction time in case of drives failure.

Functionality

Speed Acceleration

Fast RAID for HDD	Innovative RAID algorithms bring maximum benefits for working with massive sequential workloads
ERA Engine	Ultra-fast RAID engine works efficiently with NVMe and SSD devices
QoSMic	Self-learning QoS technology automatically prioritizes workloads from the valuable business services and applications
Advanced Reconstruction	Automatically excludes the slowest drives from reading operations to keep high level of total array performance
SSD cache	Increases system performance due to fast drives used as a buffer for frequently requested data
Adaptive Read-Ahead	Data caching technology analyzes incoming workload, predicts it, and then puts it into cache for overall system acceleration

Data Availability

Fast RAID rebuild	Unique checksum calculation approach ensures record speed of RAID reconstruction
Dual Controller Mode	Solution employs additional hardware controller to eliminate single point of failure. Works in the asynchronous Active-Active mode
Partial Reconstruction	Reconstructs only those particular drive areas where data has been changed. It is possible with the help of RAID logical mapping
RAID 7.3	Performance of RAID 7.3 is similar to RAID 6, but RAID 7.3 has triple parity that significantly increases reliability level
RAID N+M	RAID with the highest rate of data reliability, where users can set the number of disks for checksums allocation
SparePool	Ability to create a drive pool dedicated for hot spare

Data Protection

Remote Volume Replication	Creates synchronized copies of data in secondary data center to lower the impact of threats by natural and man-driven disasters
Silent Data Corruption Protection	System scans and fixes silent errors in the background mode during reading operations from clients with insignificant performance loss

Security

System User Management	User management policy includes two user roles: administrator and operator
Password Policy	Administrator is able to set the strength and validity period of the users' passwords

Software Characteristic

RAID levels	RAID 0/1/5/6/7.3/N+M/10/50/60/70
File protocols	SMB v2/v3, NFS v3/v4, AFP, FTP, NFS-over-RDMA
Block protocols	FC 8Gb/16Gb/32Gb InfiniBand SRP 20Gb/40Gb/56Gb/100Gb iSCSI/iSER 10Gb/25Gb/40Gb/100Gb NVMe-over-TCP, NVMe-over-RDMA
Compatible parallel file system	StorNext, Lustre, Hyper FS, Tiger FS, BeeGFS, GPFS, Xsan
Heterogeneous System Architecture Support	LUNs can be accessed simultaneously through all available block interfaces
Drive interfaces	3.5" SAS / NL-SAS / SATA HDDs 2.5" SAS / NL-SAS / SATA HDDs SATA / SAS SSDs NVMe
Max number of drives in one RAID	64
Max number of drives in the system	600
Max number of LUNs	447
Compatible virtualization platforms	VMware ESX KVM (Kernel-based Virtual Machine) RHEV (Red Hat Enterprise Virtualization) Microsoft Hyper-V Server, XenServer, Proxmox VE