

Resilient Cloud Storage is enterprise-class elastic and distributed data storage system

By virtue of its scale-out architecture the system with several storage modules can be easily transformed into structure with several thousands of storage modules and several hundreds of petabytes of storage volume




«What does it mean to scale?

Essentially, it means how well a particular solution fits a problem as the scope of that problem increases»

«Horizontal scalability is the best (and only true) type of scalability»

Theo Schlossnagle

«Scalable Internet Architectures»



The platform is a perfect fit for such tasks as:

- ✓ CCTV-monitoring
- ✓ archiving and backups
- ✓ fileshares
- ✓ virtualization platforms
- ✓ high-load applications
- ✓ databases (PostgreSQL, Oracle, MS SQL etc.)
- ✓ MS Exchange, MS Sharepoint, IBM Lotus Domino, Bitrix, 1C and any other tasks that require scalability

Advantages of the platform



Support of mainstream virtualization platforms

- VMware, KVM, MS Hyper-V



Capability of read caching using SSDs



Reliability and performance

- triple data reservation
- clients of the system interact directly with data modules
- resource groups reserved for top-priority/high-load applications



Self-healing and automated data balancing

- in case of drive or whole data module failure redundancy recovers automatically on free space with minimum performance degradation
- automated data balancing with addition of data modules



Communication protocols

- RBD, iSCSI, FC, S3, NFS, CIFS



Virtualization of customer's legacy data storage systems via iSCSI



Scalability up to hundreds of petabytes


- linear growth of capacity and performance by addition of standard data modules



Synchronous (metrocluster) and asynchronous replication available



Designed in Russia



Snapshots and clones without performance degradation



No centralized controllers («bottleneck» of typical data storage systems)



Universal data storage system

- provides block-based, object-based and file-based storage



Software updates without downtimes and loss of performance

Usage scenarios

CCTV

- High storage density (1 PB of usable capacity per rack (includes triple replication))
- Video storage for 38 000 days for standard surveillance camera (3 Mbps per camera) per rack (1 PB)
- Video storage for 2 300 days for 4K cameras (50 Mbps per camera) per rack (1 PB)
- Linear growth of capacity with increase of number of cameras and retention period
- Support for S3 and RBD protocols
- High performance even during post-failure recovery of redundancy

Media content

- High storage reliability (triple data replication)
- Replication to backup data center (to similar storage system)
- Availability of acceleration with SSDs
- High performance even during post-failure recovery of redundancy
- Snapshot-based backups
- Integration with legacy data storage systems
- Block-based, object-based and file-based storage of media content
- S3 and RBD protocols support

Servers virtualization and private clouds

- Scale-out architecture allows to expand clouds and server platforms predictably
- High performance even during post-failure recovery of redundancy
- Main virtualization platforms compatibility (OpenStack, VMware, KVM, MS Hyper-V)
- Fast and slow pools support
- Snapshots, thin disks and cloning support
- Virtualization of customer's legacy data storage systems

Virtual desktop infrastructure

- The system is compatible with VMware Horizon Suite and Citrix
- The system is equipped with SSDs for VDI I/O operations acceleration
- Easily handles boot storms during startup of virtual desktops
- Scalability from small-scale to huge-scale installations

HPC tasks (geomodelling)

- Linear growth of performance
- High storage density (1 PB of usable capacity per rack (includes triple replication))
- Supports RBD protocol
- High performance even during post-failure recovery of redundancy

Usage scenarios

Backups

- High storage density (1 PB of usable capacity per rack (includes triple replication))
- High storage reliability (triple data replication)
- Replication to backup data center (to similar storage system)
- Integration with legacy data storage systems
- Block-based, object-based and file-based storage
- Snapshot-based backups

Big data

- Capability to grow from small capacities to hundreds of petabytes
- Ability to grow storage capacities independently from growth of computing power
- Support of Hadoop, Elastic search, NoSQL, Splunk
- High storage density (1 PB of usable capacity per rack (includes triple replication))
- Predictable growth of performance as storage volumes increase
- Self-healing option

Archives

- High storage reliability (triple data replication)
- Replication to backup data center (to similar storage system)
- High storage density (1 PB of usable capacity per rack (includes triple replication))
- Storage capacity and performance of single installation can grow with different generations and types of data modules
- Self-healing option
- Availability of acceleration with SSDs
- High performance even during post-failure recovery of redundancy
- Snapshot-based backups
- Integration with legacy data storage systems
- Block-based, object-based and file-based storage of data
- S3 and RBD protocols support

Unstructured information

- High storage density (1 PB of usable capacity per rack (includes triple replication))
- High storage reliability (triple data replication)
- Replication to backup data center (to similar storage system)
- Availability of acceleration with SSDs
- High performance even during post-failure recovery of redundancy
- Snapshot-based backups
- Block-based, object-based and file-based storage of data
- Support of S3, RBD, NFS, CIFS protocols

Storage of communications providers' billing information

- High storage reliability (triple data replication)
- Replication to backup data center (to similar storage system)
- Snapshot-based backups
- Block-based, object-based and file-based storage of CDR data
- S3 and RBD protocols support