



StorPool
DISTRIBUTED STORAGE

Statement on the High Availability of StorPool Storage Deployments

Risk Factors Decreasing the Availability of a Data Storage System and How StorPool Uniquely Increases Availability.....	2
Benefits of Providing a Block Storage Solution as a Fully Managed Service.....	3
Site Reliability Engineering Approach.....	3
Scale-out Architecture with Availability and Continuous Operations in its DNA.....	4
How StorPool Measures Availability.....	5
2024 Measured, actual availability.....	6
Footnotes.....	7
Summary.....	8
About Us.....	8
Contact Us.....	8

Risk Factors Decreasing the Availability of a Data Storage System and How StorPool Uniquely Increases Availability

Achieving an availability of 99.999% in a primary data storage system is challenging and requires eliminating as many common sources of downtime as possible. Below are listed key factors that lower a system's availability *if it is not designed, deployed, and operated correctly*:

1. Storage System Design:

- 1.1. **Interdependency:** the lack of redundancy in storage system components caused by implementing a traditional storage stack with many interdependent roles such as controllers and metadata nodes lowers overall system availability.
- 1.2. **Data Protection:** issues like slow, partial, or misplaced writes during data replication across storage nodes increase the risk of data loss or corruption and, therefore, lower system availability.
- 1.3. **Inefficient Recovery:** a slow recovery process during or after failure (e.g., rebuilding a RAID group or resynchronizing replicas to another node) may lower system availability.

2. Hardware Failures:

- 2.1. **Drives:** hard disk drives (HDDs) and solid state drives (SSDs) may fail, leading to data unavailability.
- 2.2. **Network Equipment:** the failure of switches or network interface cards (NICs) may cause connectivity problems that disrupt write/read operations, leading to I/O errors or even causing data corruption.
- 2.3. **Servers:** the failure of physical servers hosting the storage system may impact availability temporarily or indefinitely.
- 2.4. **Power Supplies:** outages of rack PSUs may temporarily stop the storage system.

3. Software Failures:

- 3.1. **Errors:** software bugs in the component firmware, operating system, or storage software may cause crashes or instability.
- 3.2. **Corruption:** data corruption may be caused by mistakes in handling I/O operations in the storage system or in the I/O stack.
- 3.3. **Updates:** a system not designed for non-disruptive operations of all parts of the system will likely experience disruption of service during patches, updates, and upgrades. In addition, applying poorly-tested software patches or updates can also decrease overall availability.

4. **Operational Issues:** e.g. human errors, lack of sufficient monitoring, or improper maintenance processes.

StorPool Storage is uniquely delivered as a [Fully Managed Service \(FMS\)](#). The integrated FMS process for designing, deploying, fine-tuning, monitoring, maintaining, and performing long-term continuous service improvements addresses all of the above risk factors to achieve an extreme level of long-term, **measured** availability in excess of **99.999%**, with ***maintenance included!***

Benefits of Providing a Block Storage Solution as a Fully Managed Service

- Because of our unique service delivery approach for a software-defined storage vendor, we have full visibility into StorPool Storage's system operational metrics in customers' environments. This is system metadata, not customer data. A storage vendor providing just support services cannot achieve this level of visibility and knowledge about the health and operation of the systems in production.
- Instead of handing over the StorPool Storage solution to be operated by an (often non-specialized) customer team, all StorPool storage systems are managed by StorPool's operations team. The customer has full visibility and control, and they may take over some tasks, but almost without exception, they prefer to let our team manage their StorPool storage systems.
- For the avoidance of doubt, the customer has user interfaces (CLI, GUI, API) and integrations/automations (CloudStack, OpenStack, OpenNebula, Proxmox, Kubernetes, etc.) through which they perform common storage admin tasks such as creating, resizing, and deleting volumes, snapshots, etc.

Site Reliability Engineering Approach

- StorPool applies extremely strict procedures for validating its software with all compatible hardware components, and the full stack of all supported Operating Systems. These validations ensure that all components deliver the required level of performance, performance consistency, and reliability before releasing planned changes or fixes.
- StorPool applies numerous stages of quality control and validation throughout the development and operations processes:
 - Internal fuzzing-style framework for ensuring the robustness of all new changes proposed for release.
 - Each new release goes through several sets of validation runs, including an end-to-end test with environments very close to the ones that we see in many of the production systems using similar hardware components, operating systems and network configurations.

- Each new validated release is staged through all StorPool internal production clusters, all customers' testing/staging clusters, and, finally, customers' production systems
- All components in our hardware compatibility list are battle-tested in many production systems. All new components being added are validated with an extensive procedure to identify the optimal configuration for each of them for use in StorPool storage systems.
- Before installing StorPool Storage, each new StorPool deployment undergoes a series of hardware functional tests, stress tests, and performance tests. After deployment, the operations team performs total solution functional and performance tuning for each storage system based on the use case requirements before handover.
- Production systems undergo continuous software updates (approximately once every 3-4 months). As of 2025-09-26, the oldest version of StorPool running in production is from March 2025, while 95% of clusters are running a version released in June 2025 or later. Software upgrades are non-disruptive by design.
- Production deployments are continuously monitored. Many non-service-affecting alerts are generated and responded to before they escalate to affecting the storage service.
- Instead of expecting zero downtime, our approach recognizes that downtime is, at some point, inevitable. Downtime and service degradation are measured and managed. We do blameless analysis of any issue and don't blame the person who stumbled, we figure out ways to improve the software, our systems and processes so that we eliminate or reduce whole classes of service-affecting issues.
- We write post-mortem documents for all downtime and severe service degradation issues, regardless of whether the cause was within our area of responsibility or outside of it. For example, storage network issues and power outages are common causes of downtime that we track, but they are usually untracked even by companies providing a managed service with an SLA.

Scale-out Architecture with Availability and Continuous Operations in its DNA

- The StorPool storage system is a shared-nothing cluster. A failure of one or two storage nodes does not affect the availability of the whole storage system.
- Data redundancy is across storage nodes (or across racks in rare rack-level redundancy deployments). This means that the system can tolerate 1 or 2 storage nodes being down (or 1 or 2 entire racks, with the rack-level design) for whatever reason without service impact.
- StorPool has a flexible and dynamic data placement scheme, which enables Virtual Spares and Mesh Recovery. When a drive or node fails, mesh recovery means redundancy is restored by writing a new copy of the lost redundant data evenly

across the remaining drives and nodes, thus utilizing the performance of the system and finishing recovery very quickly.

- The StorPool storage system has been designed from the ground up for continuous operation, which is more than simply the ability to continue working without interruption when a piece of hardware fails. It includes non-disruptive software upgrades, non-disruptive capacity additions, or migration to new hardware.
- It isn't as simple as having a scale-out distributed system with 100s of processes (distributed services) working on 10s of storage nodes, as that does not automatically provide higher availability than simpler dual-controller systems. StorPool achieves higher availability because of numerous low-level architecture and implementation features. Just a few examples: network-path-aware service discovery with sub-second convergence times, service liveness detection (non-responsive services are ejected), disk liveness detection (non-responsive disks are ejected), end-to-end data integrity checksums, strong "fail-stop" implementation of all services (when services detect they are running outside of designed/tested state they exit, to contain the failure).
- StorPool has a system of interlocks including "maintenance mode", which prevents operator errors in ways similar to physical interlocks used in safety-critical systems. For example, an operator is prevented from stopping/restarting a service or a node if that would bring the redundancy level down to 0 or if it would directly lead to downtime.

How StorPool Measures Availability

We calculate the availability of the storage solution using the following methodology, data, and rules:

1. We collect post-mortem documents for every degradation and downtime event. We have a list of all post-mortem documents written within the given time period. The list includes post-mortems for each downtime and severe degradation event that happened in production clusters. (Reminder, we have full monitoring and visibility of production deployments of StorPool)
2. Each post-mortem contains details about the amount of affected data in TB, the length of the event, and the main cause of the event.
3. Based on the list of postmortems, we calculate the total number of terabyte-hours affected by downtime. We filter the list for specific severity and cause and add up the affected terabyte-hours.
4. We also calculate the total terabyte-hours of production usage for the period we are measuring by taking the average of weekly usage measurements and multiplying by 365 days and 24 hours.
5. The total un-availability for the year is (TB.h downtime from p.3 above) divided by (TB.h of total production usage from p.4 above).
6. Total availability is 1 minus the total unavailability.

2024 Measured, Actual Availability

Number of downtime events: 43

Total TB-weighted downtime:

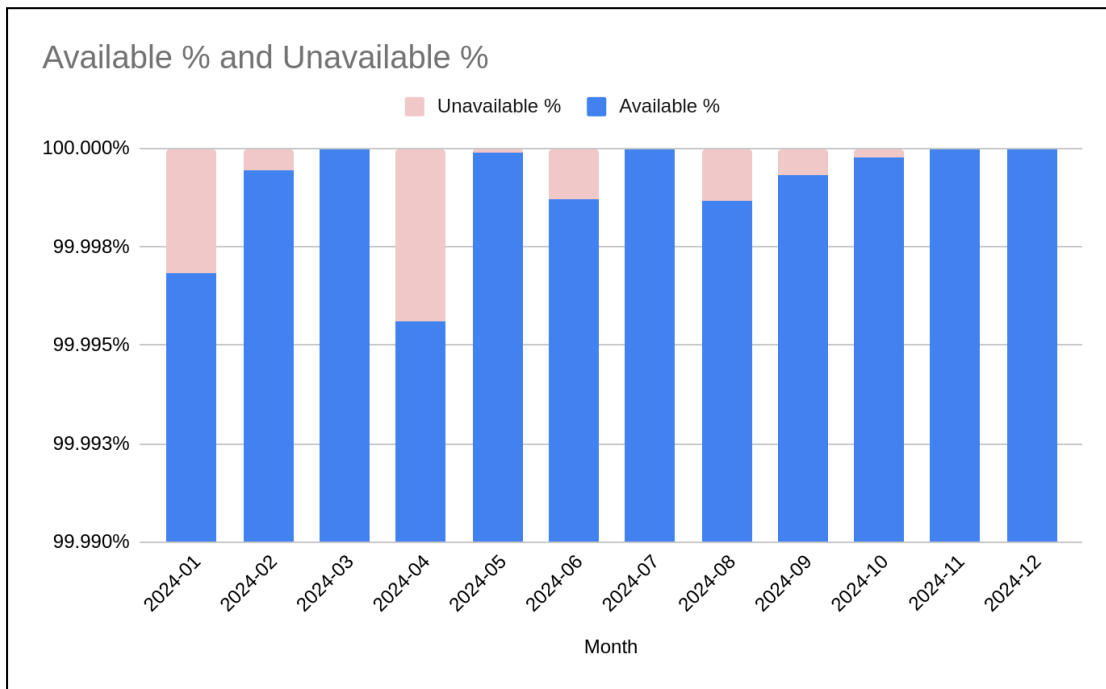
- Categories: Software, Configuration, Operator (StorPool responsibility, 100% post-mortem coverage): 0.0008%
- Categories: Network, Power, Planning, Customer mistake (Customer responsibility, 100% post-mortem coverage): 0.0002%
- Total: 0.001%

On average, you can expect any piece of data stored in a StorPool system to be inaccessible for no more than 5 minutes per year.

Total TB-weighted availability for 2024:

- 99.9990% availability all causes
- 99.9992% availability - only causes which are StorPool's responsibility or under StorPool's control
- **N.B. These numbers are availability measured in production, which is not comparable to theoretical "designed for 99.999% availability" estimations and SLA guarantees of other storage vendors.**
- **N.B.2. These numbers include all causes of downtime, regardless of who was the developer of the software components of the total solution, who changed the configuration, who unplugged the cable, etc. There are no exclusions and exceptions for "maintenance windows", "non-working hours" etc.**

Per-month availability for 2024:



Month	Available %	Unavailable %
2024-01	99.9968%	0.0032%
2024-02	99.9994%	0.0006%
2024-03	100.0000%	0.0000%
2024-04	99.9956%	0.0044%
2024-05	99.9999%	0.0001%
2024-06	99.9987%	0.0013%
2024-07	100.0000%	0.0000%
2024-08	99.9987%	0.0013%
2024-09	99.9993%	0.0007%
2024-10	99.9998%	0.0002%
2024-11	100.0000%	0.0000%
2024-12	100.0000%	0.0000%
Average	99.9990%	0.0010%

Footnotes

Examples of 2024 downtime events to illustrate the wide coverage of tracked and counted downtime events:

1. "Volume/snapshot creation and transactions were not working for 188 minutes" due to "StorPool bug".

Total affected 473 TB.h

In this case, block storage read/write operations were processed, but the API used for creating/deleting/cloning/attaching volumes/snapshots was not accessible. At the moment we don't put API outages in a separate category, so an API outage counts as an outage of the whole cluster.

2. "~94 minutes of downtime for both clusters in total due to missing network" due to "Switch bug".

Total affected 196 TB.h

In this case, the storage network was down. The hardware selection, low level design, configuration and maintenance of the storage network are the customer's responsibility, however our post-mortems include such events, so that we can improve our procedures, systems and recommendations.

3. "Operations were stalled for the cluster for about 24 minutes due to overfilled drive" due to "Planning error".

Total affected: 103 TB.h

In this case, the customer tried to put more data on a StorPool system than its capacity. The system over-filled and stopped working.

4. "VMs were down on a host mistakenly rebooted by us for 13 minutes" due to "Operator error"

Total affected: 13 TB.h

In this case, a StorPool operator rebooted a KVM host by mistake during a planned non-disruptive upgrade.

In Summary

StorPool Storage systems deliver extreme levels of **actual** availability:

- 99.9990% availability - all causes
- 99.9992% availability - only causes which are StorPool's responsibility or under StorPool's control

The above numbers of availability are **statistically measured in production**, which is above and beyond **theoretical "5-nines" availability estimations** and SLA guarantees of other storage vendors. This makes StorPool Storage a preferred storage provider for use cases where both uptime and performance are in focus.

About StorPool

StorPool Storage is an ultra-fast, fully-managed block storage software platform designed for large-scale clouds. It is the easiest way to convert sets of commodity servers into primary storage systems. The StorPool team has experience working with various clients – Managed Service Providers, Hosting Service Providers, Cloud Service Providers, enterprises, and SaaS vendors. StorPool Storage comes as software, plus a fully-managed data storage service that transforms standard hardware into ultra-fast, reliable, and linearly scalable storage systems. [Learn more about StorPool Storage and how we accelerate the world by storing data more productively!](#)

Leading IT organizations like Atos, NASA, CERN, Siemens, Deutsche Bourse, European Space Agency, Samsung, and many others have data running on StorPool.

With a global list of customers in 30+ countries, over 20 major releases, hundreds of all-flash deployments, and solid growth, StorPool Storage is a reliable yet flexible partner for any public or private cloud builder.

Contact Us

To learn more or schedule a free test, contact us at:

StorPool Storage

info@storpool.com

www.storpool.com