

Traditional Data Storage Systems	Resilient Cloud Storage Data Storage System
Performance scalability	
Scale-up systems (always limited by performance of data storage system controllers).	Scale-out system (limited by size of installation).
Capacity scalability	
Limited to several dozens of petabytes (whith the use of slow high-capacity drives).	Scalability up to hundreds of petabytes within single storage system (even all-flash).
Centralized data storage system controllers	
Yes. As capacity or load intensity grows, controllers become "bottlenecks" of any typical data storage system.	No. Each data module functions as controller and drive enclosure at the same time. It provides linear performance (and capacity) growth with addition of data modules to existing installation. So when you feel that you need more performance you simply add data modules. Each storage system client communicates directly with data modules. There are no centralized nodes in data path that store or process data map.
Capacity expansion	
Implemented by adding new drive enclosures to available storage system controllers. The number of drive enclosures per controller is limited. As the limit is reached, purchase of new controller is required. Some storage systems allow to cluster controllers, though the number of controllers in such cluster is quite small (up to eight). Functional limitations are possilble when controllers operate in cluster.	Implemented by adding unified data modules that act as (distributed) data storage system controller. This scales total capacity of a single system (alongside with simultaneous growth of its performance).
Performance expansion (IOPS, Mb/s)	
Implemented by adding drive enclosures up to the limit of controller. When the limit is reached the only way is to buy separate system with more powerful controller.	Implemented by adding unified data modules. This scales total performance of a single system (alongside with simultaneous growth of its capacity).
High-load applications data placement	
In order to achieve high performance for heavily loaded applications, distribution of data among several RAID-groups is usually employed. This is a burdensome and manual operation. More importantly, partial failure of one of RAID-groups leads to significant degradation of performance of access to all heavily loaded application data, and complete failure of a single RAID-group will lead to data loss.	All data is distributed evenly among all disk drives. Through this high parallelism and performance of each application/server are being acheived.
Automated data balancing	
In most cases one or more RAID-groups are organized over newly added enclosedes so it is possible to place new data volumes over them. Manual actions of administrator are required for relocation of part of old data to new space. Possibility of performance degradation.	All data is distributed evenly among all disk drives. When new data modules are added, data is getting automatically rebalanced in background between old and newly added space. It means that the load of existing disk drives is getting decreased hence increasing overall performance.

Traditional Data Storage Systems	Resilient Cloud Storage Data Storage System
Self-healing functionality	
<p>When disk drives are lost and redundancy is being recovered, all drives of RAID-group remain in high load mode for a long time (for several days in a row), while response time for applications with data placed on bad RAID-group increases.</p>	<p>In case of drive/data module or whole rack failure redundancy recovers automatically on free space with minimum performance degradation. Administrator's interference is not required. Urgent hot-swapping of broken-down disk drive/data modules is not required.</p>
System maintenance without service downtimes	
<p>Sometimes it is required to relocate some infrastructure components to other server rooms (or between racks) without service downtimes. For traditional data storage it is often impossible without downtimes of some of the services (volumes) or whole system shutdown.</p>	<p>In Resilient Cloud Storage even the whole rack can be shutdown safely for a while. The equipment then can be relocated and reconnected easily. Consistency will be automatically recovered after restart of relocated modules without substantial performance impact.</p>
Simultaneous usage of SSDs and HDDs in data modules	
<p>Shared write cache is usually used during simultaneous usage of SSDs and HDDs. In order to limit impact of HDDs on SSDs it is required to set write cache size limits on controller for each particular volume/pool (or employ another impact mitigation techniques). This is usually a manual and burdensome operation.</p>	<p>HDDs and SSDs do not compete for write cache because they use separate write caches by design.</p>
Replication speed between storage installations	
<p>Replicas are transferred in single stream. During long distance replication in case of network packet loss, replication speed can significantly drop without further recovery (TCP congestion window shrinks and increases slowly or does not increase at all). As a result - serious data synchronisation delays may occur (up to several days/weeks).</p>	<p>Resilient Cloud Storage employs multistream transfer of replicas, which allows to aggressively utilize available bandwidth and is very tolerant to network loss. Replication bandwidth cap can be defined, data storage system will occupy all the bandwidth provided. This allows to control data synchronization delays.</p>
Number of connection interfaces for storage clients	
<p>Limited to several interfaces per data storage system controller. The number of controllers itself is quite small. Therefore aggregated bandwidth is limited.</p>	<p>Two 10G Ethernet interfaces per data module. Data is transferred directly between clients and data modules. Because there is a significant number of storage modules, aggregated bandwidth is much higher compared to typical data storage systems. Bandwidth increases linearly with additions of new data modules.</p>
Licensing	
<p>Main functions are included in basic license for purchased volume. The rest is usually being licensed separately (for example: replications, file access, metrocluster etc.)</p>	<p>Default license includes all data storage system functionality for purchased volume.</p>