

Enterprise Cloud Storage Technical Overview



Next-generation software-defined enterprise cloud storage, designed to eliminate financial, technical, and operational risk.

Storage Dimensions

100

max. nodes per cloud

3600

max. drives per cloud

21.6PB

max. capacity per cloud

An enterprise data storage solution designed to support cloud deployments must be more sophisticated than a legacy storage solution, as it must perform all the duties of legacy storage, and more. It must support scalable CPU and transport resources, native multi-tenancy, administration and workload isolation, encryption of data at rest and in transit, with different keys for each workload, be API-driven, and be designed with cost control in mind. Back in 2011 we couldn't find an enterprise cloud storage solution that met those requirements. So we built one.

Zadara provides everything you expect in enterprise storage, including multi-protocol support, dual high-availability controllers, a choice of drive types, flash cache, and advanced data management features. But unlike legacy storage, Zadara is a fully-managed, pay-as-you-go service that effectively eliminates your financial, technical, and operational risk.

Use Zadara for any data type (block, file, object), any protocol (iSCSI, iSER, FibreChannel, NFS, CIFS/SMB), and any location (public cloud, private cloud, on premises, hybrid).

Contents

Features

- Scalability

- Low-Cost, Low Latency

- Storage Array Controllers

- Object Storage

Security

Standards Compliance

Fully-Managed

Advanced Data Management Services

- Replicate to Any Location

- Multi-Zone High Availability

- Cloud Hydration

- Microsoft VSS Snapshots

- Docker Containers

- Volume Migration

General Specifications

Zadara Engines

Object Storage Specifications

Scalability

Zadara is delivered as a collection of storage nodes: standard servers running Linux, with storage options that include hard drives, SSDs, and NVMe. Connectivity is built with Mellanox 40Gb Ethernet switches and NICs. Based on standard building blocks (storage nodes), Zadara allows you to start with as few as two storage nodes and scale to as many as 100 storage nodes. This software-defined storage scales in every dimension: 1) CPU resources for the storage controllers; 2) drive resources for raw storage; and 3) transport for interconnectivity.

Low-Cost, Low-Latency Transport

Using 40Gb interconnects and the iSCSI extensions for RDMA (iSER), we built a low-cost, low-latency network fabric with sufficient bandwidth for the most demanding workloads. All Zadara storage solutions use iSER as the protocol for the interconnect. Customers can choose whether or not to implement iSER on their clients. Zadara supports Fibre Channel, iSCSI, iSER, NFS, CIFS, S3 and Swift host connectivity.

Storage Array Controllers

Zadara storage arrays use OpenStack and KVM to run and manage virtual controllers (Zadara Engines), which are virtual machines running on the storage nodes. Each controller is assigned dedicated CPU cores on a storage node plus dedicated disks across several nodes, so there is no over-subscription of these resources. The combination of these dedicated cores and disks is a storage array, and the workloads in each array are completely isolated from each other. Your Zadara storage cloud can run thousands of virtual arrays simultaneously.

Object Storage

In addition to storage arrays for block and file storage, Zadara offers the industry's only 'private' enterprise cloud object storage solution. Zadara object storage utilizes dedicated storage resources, is compatible with public object storage, and supports the same interface (Amazon S3 RESTful API, Swift API), while allowing customers to keep the data on premises or in the cloud. This provides both higher levels of security, but also provides consistent levels of performance.

Security

Security is integral to Zadara, starting with workload isolation. Drives are mapped to a single array at a time and if an array ever vacates a drive it is scrubbed before being made available again. Administrator accounts are never shared between arrays, and the drives in each array can be encrypted with their own, customer-managed key. The password to the key is never stored durably and is never shared outside the array. Arrays support IPsec for encryption of data in transit between the array and clients, and array-array replication traffic is always encrypted. Clients are mapped to individual arrays and a client can access only the LUNs or shares in arrays to which it has been granted access.

Standards Compliance

Zadara conducts ongoing security testing of its clouds and maintains security certifications such as ISO 27001, SOC 2 Type 2, and HIPAA. GDPR compliance is a shared responsibility. We offer a wide set of controls to help you maintain GDPR compliance. For more details please visit the Zadara website. There you will find information about the security measures we have in place.

Fully-Managed, Upgrades Included

Zadara is a fully-managed storage service with 24/7 monitoring and support. And it doesn't end there. Zadara automatically upgrades hardware without any application impact. For Zadara running in the public cloud this happens seamlessly and without the need for any work on your part. When a Zadara storage node running on your premises has reached its end of life, we ship you a new node. The data from the EOL node is copied — online and without a performance impact — to the new storage node. The drives in the EOL node are software shredded and then the node is returned to Zadara. All of these tasks are managed and monitored by the Zadara NOC, so the only thing you need to do is swap the old gear for the new gear in the rack.

Replicate to Any Location

Zadara is identical whether deployed on premises or in a public cloud, so you can replicate between any Zadara storage clouds, wherever they are deployed. For instance, you can replicate to our deployments at dozens of data centers around the world, including Amazon Web Services, Google Cloud Platform and Microsoft Azure.

Multi-Zone High Availability

In certain regions, and in all on-premises deployments, Zadara supports Multi-Zone High Availability. When a Zadara Storage Cloud is deployed in a Multi-Zone High Availability model, the deployment is split into two Protection Zones. Each zone is deployed in a different data center located in the same metropolitan area – generally within 2ms of each other (max. 5ms). Array resources are split across two Protection Zones, effectively providing synchronous replication and controller redundancy for mission critical applications.

Cloud Hydration

Zadara's cloud hydration enables you to adopt cloud computing using a cost-effective and practical way of moving corporate data into the cloud. Cloud Hydration allows migration to the cloud of both online production environments and data that does not need to be continuously online. Cloud Hydration can be targeted to any storage medium from any vendor or cloud storage provider.

Backup to Object Storage

Using a Zadara storage array you can create an automatic, snapshot-based, continuous, incremental backup to low-cost, object storage. This eliminates the need for host-based backup software and provides a simple, easy-to-use, high-performance, block-based backup. And data backed up by Zadara can be restored back to the original volume or anywhere else you choose. You can also migrate data from Amazon S3 to Amazon's Glacier cold storage, creating new tiers of data protection. (Note: data migrated to Glacier cannot be restored.)

Microsoft VSS

Microsoft Volume Shadow Copy Service (VSS) enables online, point-in-time backups of Microsoft SQL Server, Exchange, SharePoint and other Windows-based enterprise applications in Zadara storage arrays. You can replicate these snapshots to any of Zadara's dozens of public storage clouds around the world, including AWS, Azure, Google Cloud Platform and others. VSS for Zadara means that, for the first time, public cloud customers can take online, snapshot-based backups of their Microsoft applications instead of implementing complicated work flows or stopping their mission-critical applications to create backup copies.

Zadara's Snapshots mechanism is very efficient, in terms of capacity and performance. Block volumes and file shares are supported. Snapshot policies define the Snapshots lifecycle via the enforcement of creation and deletion policies. Snapshot Policies are "global" entities and you can apply instances of the policies to one or more volumes. Snapshots can be taken as frequent as every minute. There is no limit to the number of snapshots you may keep per volume.

Docker Containers

Zadara storage arrays incorporate Docker container technology into the Zadara Engines, a pair of virtual controllers. The Zadara Engines have direct, low-latency access to your SSDs and HDDs, and therefore provide high throughput IOPS to the Docker container running within. Like the Zadara Engines, the Container has dedicated CPU and memory, which can be increased and decreased on the fly, non-disruptively. Uniquely, Zadara offers Docker high availability, thanks to Zadara's dual-engine architecture with auto-failover.

Volume Migration

Volume migration allows you to transparently move your Zadara storage array data and snapshots between different pools without affecting servers and workstations mounted to the volume or share. The most common case is to move volumes from a high-performance storage pool to a lower-tier pool when performance requirements are no longer imposed on the data set. Other uses, such as moving between pools with different durability levels, can increase storage efficiency for lower cost per GB or increase performance using the same storage media.

General Specifications

Zadara storage clouds are built upon standardized, cost-effective x86 servers and fueled by Zadara's patented software that creates a virtualizing data storage resource abstraction layer. Storage nodes are interconnected using multiple, redundant 40Gb connections with advanced iSER (RDMA) to minimize latency and dramatically improve performance. Because of virtualization and resource isolation, a Zadara storage cloud can grow to hundreds of nodes while workloads remain isolated from one another. Zadara delivers predictable performance and high Quality of Service.

General Specifications

Cloud

Max. Storage Nodes per cloud	100
Max. Drives per cloud	3500
Max. capacity per cloud (using 10TB Drives)	35 PB

Storage Node

Max. physical drives per Storage Node (SN)	60
Max. drive partitions per SN (counting physical drives & SSD partitions)	256
Max. Memory per SN	512GB
Max. VCs per Storage Node	10
Supported hard drive type	SAS 1TB 10K RPM; SATA 6TB 7200 RPM; SATA 10TB 7200 RPM
Supported SSD type	SATA 800GB; SATA 1.6TB; SATA 3.8TB
Storage Networking Supported	iSCSI; iSER; FC; NFS; CIFS/SMB

Performance

Max. host IO throughput per SN (Dual-port 40Gb NIC per SN)	2.2 GB/s
Max. host IO throughput per cloud	220 GB/s
Max. host IO throughput per array	3 GB/s

Zadara Engines

Zadara Engines intelligently place and protect your data on storage nodes to yield predictable performance, enterprise-class high availability with multi-zone high availability, and powerful data management features. Engines come in a variety of models, designated numerically from 200 to 1600. As the chart below shows, each model is equipped with a specific memory configuration and can handle a specific workload. When you provision a Zadara storage array, you choose the Engine that is appropriate to your requirements.

Engine Model	200	400	600	800	1000	1200	1600	2400
Memory								
Memory	4GB	8GB	16GB	24GB	32GB	48GB	64GB	100GB
Drives								
Max Drives	5	10	20	40	80	80	80	80
Max Raw Capacity	24TB	60TB	100TB	150TB	200TB	240TB	300TB	360TB
Max Usable Capacity (RAID 10)	12TB	30TB	50TB	75TB	100TB	120TB	150TB	180TB
Max Usable Capacity (RAID 6)	12TB	48TB	80TB	120TB	160TB	192TB	240TB	288TB
Base Flash Cache	20GB	20GB	40GB	60GB	80GB	100GB	120GB	180GB
Max Flash Cache	0GB	400GB	800GB	1.2TB	1.6TB	2.4TB	3.2TB	3.2TB
Max RG size: RAID-5	5 (4+p)							
Max RG size: RAID-6	10 (8+p+q)							
RAID Options	1, 5 (\leq 2TB), 6, 10							
Pools and Volumes								
Max Pools	8	16	32	32	32	64	64	64
Max Pool size	12TB	30TB	50TB	75TB	100TB	120TB	150TB	200TB
Max Volume size	12TB	30TB	50TB	75TB	100TB	120TB	150TB	200TB
Max Volumes per array	16	32	64	128	256	512	1024	1024
Max Volumes per host	16	32	64	128	256	256	256	256
Max Snaps per array	Unlimited							
Max Snaps per volume	Unlimited							
Max. SMB File History Snaps per share	64							
Max. SMB File History Snaps per array	512							

Engine Model	200	400	600	800	1000	1200	1600	2400
Hosts								
Max SAN Hosts per array	16	32	64	128	256	256	256	256
Host OS Supported	Windows Server 2008, 2012, 2016, Linux, VMware ESXi, AIX, Solaris							
Mirroring								
Max Mirrors per VPSA	4	8	16	32	64	128	196	256
Max Mirrors per Volume	1	2	4	4	4	4	4	4
Max Mirror Targets	unlimited							
Max Latency between Protection Zones	5ms							
Max Remote Mirror Latency	500ms							
Remote Mirror min RPO	1min							
Users								
Max Users	100							
NAS								
Max Files per Share	100M	200M	400M	600M	800M	800M	800M	800M
Max Files per Directory	100K							
Max File System size	15TB	30TB	50TB	75TB	100TB	120TB	150TB	200TB

Object Storage Specifications

With Zadara's industry-leading 'private' object storage (using dedicated resources), data objects are accessed with user-defined metadata, providing extremely high performance by efficiently storing massive volumes of unstructured data. Zadara object storage provides practically infinite scalability for storing data that is static in nature such as: multimedia, web content, big data, archive and backup files.

Object Storage Specifications

Virtual Controller (VC)

vCPUs	6
RAM	12GB
Max Drives per VC	12

Drives

Min Drives per VPSA Object Storage	4
Max Drives per VPSA Object Storage	3500


Engines

Min Engines per VPSA Object Storage	2
Max Engines per VPSA Object Storage	175

Capacity

Max Raw Capacity (using 10TB drives)	35 PB
Max Usable Capacity	17 PB
Max Objects per VPSA Object Storage	Unlimited
Max Capacity per Container	10PB
Max Objects per Container	Unlimited
Min Object Size	1B
Max Object Size (non segmented)	5GB
Max Dynamic Large Object (DLO) Size	Unlimited

Key Service Providers



Transform your business with zero-risk enterprise storage.

Zadara transforms storage-related costs from a variable mix of equipment and management expenses to a predictable, on-demand, pay-per-use, elastic service that greatly simplifies planning, streamlines budgeting, and improves return on investment (ROI). Find out how zero-risk enterprise storage can help transform your business. Call or email today.

+1 949 251 0360
sales@zadara.com
www.zadara.com

