

zadara

ENTERPRISE STORAGE CLOUD TECHNICAL OVERVIEW

Revision 2.0

Version 20.01

June 2020

[This document is suitable for publication.]

CONTENTS

CONTENTS	2
PREFACE	4
SYSTEM FEATURES	5
ZADARA STORAGE CLOUD SCALABILITY	9
VPSA STORAGE ARRAY SPECIFICATIONS	10
VPSA ALL-FLASH ARRAY SPECIFICATIONS	12
VPSA OBJECT STORAGE SPECIFICATIONS	14
ZADARA DATA SERVICES SPECIFICATIONS	18

PREFACE

This document describes the features and recommended system scalability specifications of the Zadara Storage Cloud.

System Version: 20.01 (June 2020)

CONTACT INFORMATION

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

©2018-20 Zadara Storage, Inc. All rights reserved. Zadara is a registered trademark of Zadara Storage, Inc. Rev. 0620

SYSTEM FEATURES

Transform Your Business

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).

Scalability

Zadara is delivered as a collection of storage nodes: standard servers running Linux, with storage options that include hard drives, SSDs, and NVMe flash. 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. All storage services and tiers can be delivered from the same cluster of storage nodes.

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 and, optionally, for host connectivity. 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 processor cores, memory, and disks comprises a virtual private storage array. The workloads in each virtual private array are completely isolated from each other and each can deliver both file and block storage based on hard disk drives or solid state drives (SAS or NVMe). Zadara storage clouds 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 fully-managed 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 enterprise storage-as-a-service provides simple, secure, and powerful on-demand data mobility services between virtual private arrays, regardless of location. Use Zadara Remote Mirror for disaster recovery and business continuity applications. Deploy Zadara Remote Clone for automated migrations, offline processing tasks, and data distribution. Either way, you get one-to-one or one-to-many replication service across racks, data centers, clouds, and continents, without extra license or service level expenses.

Multi-Zone High Availability for VPSA Storage Array and Object Storage

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 Protection Zone is deployed in a different location or Availability Zone located in the same metropolitan area – generally within 2ms of each other (max. 5ms and 20ms for Object Storage).

VPSA Storage Array resources are split across two Protection Zones, ensuring data availability is maintained, even in the case of an Availability Zone outage. For VPSA Object Storage, the system will create a complete Data Policy according to the one selected during the creation, ensuring the solution can sustain a complete region loss.

Cloud Hydration Service

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 Service

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.

Antivirus Service

Zadara storage integrated antivirus protection simplifies and scales to provide fast, simple, and highly-available antivirus protection. By integrating the McAfee virus scanning engine directly into the file service, antivirus file scans are performed locally to storage, eliminating network latency, reducing network traffic, and eliminating the requirement for dedicated antivirus servers. In addition, delivering antivirus protection as a service relieves IT organizations of the effort and complexity of managing, maintaining, and supporting antivirus engines and their associated infrastructure.

Microsoft VSS

Microsoft Volume Shadow Copy Service (VSS) enables online, point-in-time snapshots 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. 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.

ZADARA STORAGE CLOUD SCALABILITY

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. With 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.

Cloud	
Max Storage Nodes per Zadara Storage Cloud	100
Max Drives per Zadara Storage Cloud	6400
Max capacity per Zadara Storage Cloud (using 14TB Drives)	90 PB
Storage Node	
Max physical data drives per Storage Node	64
Max drive partitions per SN (counting physical drives & SSD partitions)	256
Max VCs per Storage Node	10
Supported hard drive type	SATA 6TB 7200 RPM SATA 10TB 7200 RPM SATA 14TB 7200RPM
Supported SSD type	SATA 1.6TB SATA 3.8TB NVMe 7.6TB
Supported 3D XPoint Optane SSD type	NVMe P4800 385GB
Storage Networking Supported	iSCSI; iSER; FC; NFS; CIFS/SMB;

VPSA STORAGE ARRAY SPECIFICATIONS

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	3600
vCPUs									
	2	4	6	8	10	12	16	24	36
Memory									
Memory (GB)	6	12	20	28	36	52	68	100	144
Drives									
Max Drives	5	10	20	30	40	60	80	80	80
Max Raw Capacity	24	60	100	150	200	240	300	360	360
Max Usable Capacity (TB) (RAID10)	12	30	50	75	100	120	150	180	180
Max Usable Capacity (TB) (RAID6)	12	48	80	120	160	192	240	288	288
Base Flash Cache (GB)	20	20	40	60	80	100	120	180	240
Max Extended Flash Cache (GB)	0	400	800	1200	1600	2400	3300	3200	3200
Max RG size: RAID-6	10 (8+p+q)								
RAID Options	1, 6, 10								
Pools and Volumes									
Max Pools	8	16	32	32	32	64	64	64	64
Max Pool size (TB)	12	30	50	75	100	120	150	200	200
Max Volume size (TB)	12	30	50	75	100	120	150	200	200

Max Volumes per VPSA	16	32	64	128	256	512	1024	1024	1024
Max Volumes per Host	16	32	64	128	256	256	256	256	256
Max Snaps per VPSA	Unlimited								
Max Snaps per Volume	Unlimited								
Max SMB File History Snaps per Share	64								
Max SMB File History Snaps per VPSA	512								
Hosts									
Max SAN Hosts per VPSA	16	32	64	128	256	512	1024	1024	1024
Mirroring									
Max Mirrors per VPSA	4	8	16	32	64	128	196	256	256
Max Mirrors per Volume	1	2	4	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	1 Minute								
Users									
Max Users (Storage Admins)	100								
NAS									
Max Files per Share (M)	100	200	400	600	800	800	800	800	800
Max Files per Directory	100K								
Max File System size (TB)	15	30	50	75	100	120	150	200	200

VPSA ALL-FLASH ARRAY SPECIFICATIONS

Zadara All-flash VPSAs are ideal for applications and use cases that demand very high sustained performance at a compelling price point. Engines for All Flash VPSAs are configured to deliver uncompromising IOPS and highly-efficient data reduction.

Engine Model	F800	F1200	F2400	F3600	F4800
vCPUs					
	8	12	24	36	48
Memory					
Memory (GB)	48	72	116	176	236
Drives					
Optane Cache Capacity (GB)	40	60	100	140	200
Max Drives	30	60	80	100	120
Max Raw Capacity	120	200	320	400	400
Max RG size: RAID-6	10 (8+p+q)				
RAID Options	1, 6, 10				
Pools and Volumes					
Max Pools	1	1	2	2	2
Max Transactional Pool/Volume size	50				
Max Repository Pool/Volume size	100				
Max Archival Pool/Volume size	200				
Max VPSA Data Capacity ¹ Using Transactional Pools (TB)	20	40	60	80	100
Max VPSA Data Capacity ² Using Repository Pools (TB)	40	60	100	160	200
Max VPSA Data Capacity ³	60	100	160	200	250 ⁴

¹ Capacity as seen by the host

² Capacity as seen by the host

³ Capacity as seen by the host

⁴ Can go up to 250TB if data reduction ratio is at least 1.5:1

Using Archive Pools (TB)					
Max Volumes per VPSA	32	128	256	512	1024
Max Volumes per Host	32	128	256	256	256
Max Snaps per VPSA	Unlimited				
Max Snaps per Volume	Unlimited				
Max SMB File History Snaps per Share	64				
Max SMB File History Snaps per VPSA	512				
Hosts					
Max SAN Hosts per VPSA	32	128	256	512	1024
Mirroring					
Max Mirrors per VPSA	32	64	128	196	256
Max Mirrors per Volume	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 (Storage Admins)	100				
NAS					
Max Files per Share	800M				
Max Files per Directory	100K				
Max File System size (TB)	50	100	120	150	200

VPSA 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.

GENERAL SPECIFICATIONS

Cloud and Storage Nodes	
Max Protection Zones per Zadara Cloud	2
Min Fault Domains per Zadara Cloud (single Protection Zone)	2
Max Fault Domains per Zadara Cloud (single Protection Zone)	4
Min Fault Domains per Zadara Cloud (Multi Zone)	6
Max Fault Domains per Zadara Cloud (Multi Zone)	8
Max Latency between Protection Zones	20ms
Object and Containers	
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

VPSA OBJECT STORAGE PROFILES SPECIFICATIONS

	Standard	Premium	Premium Plus
Virtual Controller (VC) Storage+Proxy			
vCPUs	6	10	10
RAM (GB)	18	52	52
Data Drives per VC	12	12	12
Virtual Controller (VC) - Controller (Ring Master/Ring Slave)			
vCPUs	-	4	4
RAM (GB)	-	16	16
Virtual Controller (VC) - Proxy/Zadara Elastic Load Balancer (ZELB)			
vCPUs	6	4	4
RAM (GB)	16	16	16
ZELB VCs (per ZELB group)	2	2	2
Proxy+Storage VCs (per group)	12	12	12
Capacity & Data Drives			
Max Usable Capacity (TiB)	1,024	4,096	61,440
Min Drives per Object Storage	4	24	48
Minimum Drive Count For Policy Creation			
2 Way Mirror	4	24	48
3 Way Mirror Protection	6	36	72
Erasure Coding 4+2 (3 FD)	-	36	72
Erasure Coding 6+3 (3 FD)	-	36	72
Erasure Coding 6+2 (4 FD)	-	48	96
Erasure Coding 9+3 (4 FD)	-	48	96
Multizone Erasure Coding 4+2 (4 FD)	-	72	144

Multizone Erasure Coding 9+3 (4 FD)	-	96	192
Minimum Drive Count For Policy Expansion			
2 Way Mirror	2	2	24
3 Way Mirror Protection	3	3	36
Erasure Coding 4+2	-	6	36
Erasure Coding 6+3	-	9	36
Erasure Coding 6+2	-	8	48
Erasure Coding 9+3	-	12	48
Multizone Erasure Coding 4+2	-	16	72
Multizone Erasure Coding 9+3	-	24	96
Engines			
Controller Engines	-	2	2
Minimum "Storage+Proxy" Engines	2	2	4
Maximum "Storage+Proxy" Engines	14	56	808
Load Balancer	Internal	Internal	Elastic LB (ZELB)
Data Protection Policies			
Erasure Coding Support	-	✓	✓
2 Way Mirror	✓	✓	✓
3 Way Mirror	✓	✓	✓
Erasure Coding - 4+2 (3 FD)	-	✓	✓
Erasure Coding - 6+3 (3 FD)	-	✓	✓
Erasure Coding - 6+2 (4 FD)	-	✓	✓
Erasure Coding - 9+3 (4 FD)	-	✓	✓
MZ Erasure Coding - 4+2 (3 FD)	-	✓	✓
MZ Erasure Coding - 9+3 (4 FD)	-	✓	✓

ZADARA DATA SERVICES SPECIFICATIONS

Zadara offers data services beyond storage, giving users the ability to do more with their data within Zadara's hybrid cloud storage systems.

ASIGRA BACKUP-AS-A-SERVICE

Data VM Specification	
Engine Name	D200 (dvm.200.vf)
vCPUs	2
RAM (GB)	8
VPSA Model	600
Asigra Management Console Version	14.1 HF3 (14.1.0.3)
Asigra DS-System Software Version	14.1 HF2 (14.1.0.2)