

White Paper

Zadara Virtual Private Storage Array Technology



Executive Summary

Zadara offers enterprise storage users an innovative new way to get the best of both worlds: the power of enterprise-class storage, with the convenience of the cloud. Zadara's patented Virtual Private Storage Array technology powers a high-performance, enterprise-grade storage system that can run in any environment: on premises, private or public cloud, or a hybrid arrangement..

Introduction

The IT industry continues to innovate, to drive down cost and increase efficiency. Recently, the industry has capitalized on two key innovations: 1) Cloud-based solutions; and 2) As-a-Service solutions. Zadara® has become a leader in Enterprise Storage-as-a-Service (STaaS) solutions by offering customers a pure Operating Expense (OpEx) enterprise-grade storage solution, available in two delivery models:

- **On-Premises** – Zadara virtual private storage array located on the customer premises
- **Cloud-Based** – Zadara virtual private storage array located in any of 31 locations globally (and counting)

Some Zadara customers use only the on-premises option; some only use our cloud services; and the rest combine the two in a ‘hybrid’ solution.

It is important to know that the stack we build to deliver our public cloud deployments and the stack we deliver for on-premises customers is exactly the same. There is one product, but two very different delivery models.

Regardless of how the customer deploys Zadara, they enjoy the benefits of: 1) a pure OpEx solution, 2) fully elastic and scalable storage, 3) full monitoring and upgrading of the solution by Zadara; and 4) all of the features and functions you would expect from any enterprise storage solution.

A Bit of History

Zadara started offering iSCSI as a service to Amazon Web Services customers in 2011, adding NFS and CIFS soon thereafter. We used commodity servers and our patented technology to build SAN arrays located approximately 1ms away from EC2 in the us-east-1 region. To connect our hardware to EC2 instances we leveraged Direct Connect¹.

As the technology matured, Zadara grew steadily, supporting additional AWS regions and more cloud providers. Today, Zadara supports and services customers globally in AWS, Azure, Google Cloud Platform, Equinix, Omegabyte, Orixcom, Fortrust, Dimension Data, Object Ventures and other cloud and data center providers.

In 2014, Zadara started offering the same hardware and software stack we have been operating in dozens of locations around the world to enterprises and service providers directly. We call the on-premises model On-Premises as-a-Service (OPaaS).

Building an Entirely New Type of Storage Array

A storage array designed from the ground up to support public cloud deployments is going to be fundamentally different from the traditional storage array. This new type of array is going to 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 from day one.

That's exactly what we've built, and we call this patented² array a Zadara Storage Cloud.

¹<http://aws.amazon.com/directconnect/>


² [US Patent 8,819,230](#)

Zadara Storage Cloud

A Zadara Storage Cloud is delivered as a collection of storage nodes – Intel servers running Linux with hard drives and SSDs from companies such as Intel, Toshiba, and Western Digital. Connectivity is built with Mellanox 40Gb Ethernet switches and NICs. A storage cloud can be deployed with as few as two storage nodes and scale to as many as 100 storage nodes.



This software-defined array scales in every dimension: 1) CPU resources for the storage controllers; 2) drive resources for raw storage; and 3) transport for interconnectivity. There is no potential bottleneck.



RAID groups are built from drives spread across the entire pool; this ensures storage node redundancy and efficient use of resources. However, it also means the transport layer has the potential to introduce latency and performance challenges.

Better than TCP

To alleviate any potential issues in the transport layer, Zadara was the first enterprise storage array to implement iSER³. By using 40Gb interconnects and the iSCSI extensions for RDMA⁴ Zadara has built a low cost, low latency network fabric with sufficient bandwidth for the most demanding workloads.

All Zadara Storage Clouds use iSER as the protocol on the back-end network. Customers can choose whether or not to implement iSER on their clients. Zadara supports Fibre Channel, iSCSI, iSER, NFS, CIFS, S3 and Swift connectivity.

Controllers in a Distributed, Scale-Out Storage Array

In a traditional storage array there is a pair of dedicated hardware controllers configured in a failover cluster. Because Zadara arrays are so distributed and can scale to so many nodes, this dedicated approach wouldn't work.

³http://en.wikipedia.org/wiki/iSCSI_Extensions_for_RDMA

⁴http://en.wikipedia.org/wiki/Remote_direct_memory_access

Zadara implements controllers as pairs of virtual machines configured in a failover cluster. We use OpenStack and KVM to run and manage these controller pairs.

Each controller pair is assigned dedicated CPU cores on a pair of storage nodes plus dedicated disks across several nodes; **there is no oversubscription of these resources.** The combination of these dedicated cores and disks is called a Virtual Private Storage Array (VPSA). The workloads in each VPSA are completely isolated from each other. What happens in one VPSA cannot impact the performance of another. A single Zadara Storage Cloud can run thousands of VPSAs simultaneously.

Virtual Private Storage Arrays

A virtual private storage array looks, performs, and behaves like enterprise storage. With multiprotocol support (FC, iSCSI volumes, NFS and CIFS shares), dual high-availability controllers, SSD cache, and advanced data management features including thin provisioning, snapshots, cloning and remote replication – every feature users expect from an enterprise storage array is included.

Because individual VPSAs create strong workload isolation and because they can be created so easily, storage administrators typically utilize a VPSA for each department or per application. VPSAs are sized by the number of CPUs and the amount of SSD cache assigned to each, VPSAs can be scaled up and down with no downtime, and SSD cache can be added and removed, on-demand.

Never Migrate Arrays Again

Storage administrators spend lots of time replacing traditional storage:

1. Negotiate and purchase an array
2. Migrate data onto the array
3. Manage and upgrade the array for 3-5 years
4. Rip-and-replace the old array and start the cycle over

In a large organization with multiple arrays it's possible to always be migrating data from an old array onto a new one.

The distributed nature and remote management of a Zadara Storage Cloud eliminates this cycle by enabling an entirely new approach to replacing old hardware. Zadara will automatically — and for the life of a Zadara Storage Cloud — replace hardware without any application impact. When a Zadara storage node has reached its end of life, we ship the customer 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, the only involvement by the customer is to un-rack the old equipment and rack the new equipment.

Global Ecosystem of Replication Targets

Because the on-premises version of the Zadara Storage Cloud is identical to the cloud deployments, Zadara offers customers a first: **replicate from their on-premises storage into the cloud**. Customers can replicate from their local Zadara deployment to any of the 31 public Zadara deployments around the world including Amazon Web Services, Google Cloud Platform and Microsoft Azure. Customers can also replicate between any public Zadara Clouds.



⁵ Three overwrites

⁶ The password is stored in the memory of the clustered VPSA pair. In the case of dual controller failure (very rare) the administrator MUST provide the password to access the data.

⁷ SSL


Security

Because of our public-cloud heritage security has always been integral to the product. Our first layer of security is workload isolation; drives are mapped to a single VPSA at a time and if a VPSA ever vacates a drive it is scrubbed⁵ before being made available again. Administrator accounts are never shared between VPSAs so the administrator of the “Accounting” VPSA has NO access to the “Human Resources” VPSA. The drives in each VPSA can be encrypted with their own, customer-managed key. The password to the key is never stored durably⁶ and is never shared outside the VPSA. Zadara staff has no access to this password. VPSAs support IPSec for encryption of data in transit between the array and clients, VPSA <-> VPSA replication traffic is always encrypted⁷.

Clients are mapped to individual VPSAs, a client can only access the LUNs or shares in VPSAs to which it has been granted access, not all the shares in the storage cloud.

Geographic Redundancy

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



Protection Zones are Independent Datacenters

Conclusion

When Zadara first introduced iSCSI as a service in AWS, we charged per disk device and VPSA by the hour. As an as-a-service company, we only charge for what you use. When we introduced our on-premises solution, we kept the exact same model, again charging only for actual usage by the hour. Regardless of how much Zadara storage we've physically installed on a customer's premises, that customer only pays for the drives they are using. Whether you are in the cloud or on-premises, you have full elasticity in all directions and only pay for the capacity you are using.

With the flexibility of the cloud business model and the benefits of on-premises storage, technology isn't the only place Zadara innovates. With Zadara, your storage resources are always aligned with your business.

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

