White Paper

# Supercharge AWS with Enterprise Storage

zadara

## Executive Summary

It is no secret that storage performance is one of the main challenges for applications in the cloud. The storage performance problem in the cloud has multiple facets. It is not just low vs. high performance; it's also the fact that performance varies over time, and one customer can impact the performance of others.

This white paper describes a new storage architecture, designed to host multiple tenants with each user receiving dedicated resources to ensure predictable performance and privacy, just like in a private cloud environment.

### Traditional Storage and the Cloud

The main issue is that none of the current storage systems and associated software were designed with the cloud in mind. Even high-end SAN storage arrays were designed for a single tenant, a single manager, the storage administrator who knows what everybody else is doing, and the super guru who knows how to interpret iostat!

Aside from the lack of multi-tenancy isolation, the scalability is not even close to what is needed in the cloud. And, worse, is the inability to accommodate different types of customers with a single system. As a result, cloud-providers start with the lowest common denominator.

All these issues result in a cloud environment in which it is hard to run classic IT applications that need low latency. Instead, the types of databases and applications that are common in the cloud are the ones that can tolerate long latencies, a low IOPS rate, and inconsistencies. But these problems are solvable with a new storage architecture: one that is not based on "scale-out," but is designed to host multiple tenants, each with his own needs.

### A New, Multi-Tenant Storage Architecture

In this new storage architecture, each user — at least, the type of user who needs consistency and performance — has dedicated resources: dedicated drives, dedicated cache, dedicated cores, so that IOPS are guaranteed as if the user had a private SAN array in their own data center. And just like in a private environment, each user has access to performance data (for monitoring) and a management console (for control over and tuning of the environment). This new architecture allows for different amounts of resources to be allocated to different users, depending on their needs. Additionally, this architecture meters these resources in order to charge customers according to their actual usage.

### Virtual Private Storage Arrays

This is the foundation of the Zadara® Storage solution! We took the architecture of cloud computing and applied it to cloud block (SAN) and file (NAS) storage, creating a highly scalable, extremely reliable system, where each tenant controls their own storage. In addition, we provide on-demand resource allocation per customer (i.e., CPU, memory, networking and storage drives), eliminating the noisy neighbor problem. Our virtual private storage array service in the cloud performs and behaves like traditional SAN and NAS arrays in the data center. It has the reliability, performance and features of enterprise storage, with the elasticity and pay-per-use of the cloud.


As mentioned earlier, virtual private storage array performance is highly consistent over time, thanks to dedicated resources. The level of performance depends on three tunable characteristics which the user can change at any time via an online interface: 1) the mix of drives, 2) the RAID level; and 3) the IO Engine.

### Mix of Drives

Users first select the type and quantity of drives. Different drives have different IOPS and throughput, and are optimized for different access patterns. For example, Zadara customers at Amazon Web Services (AWS) can choose SSDs, HDDs, or a mix of both. The SSDs provide the highest random IOPS, especially random reads. In contrast, HDDs excel at sequential IO, and in particular are hard to beat when it comes to sequential writes. And, of course, they cost less per unit of storage capacity.

### RAID Level

Second, users select the RAID level, which protects against drive failures, but also affects performance in two ways. One effect is striping. Striping aggregates multiple drives' capacity and performance such that the result is a group whose performance is the sum of performance of all of its member drives. The total capacity is the sum of drives minus the RAID overhead, which varies by RAID type. The second effect is performance overhead. Certain RAID algorithms (e.g., RAID-6) provide a high level of protection against drive failure, but require more processing power than other RAID levels (e.g., RAID-1), and therefore, incur a certain performance hit versus the simpler algorithms. Each Zadara user gets to choose their optimal RAID level based on performance, capacity utilization and protection.

### IO Engine

The third user selection is the IO Engine, which determines the peak performance available. There are seven sizes of IO Engines: virtual private storage array 200, 400, 600, 800, 1000, 1200 and 1600. Each engine increases performance versus the previous one by adding resources. For example, a virtual private storage array 800 engine, which features 8 CPU cores and 24 GB of RAM, provides roughly four times the performance of the 200 engine, which has 2 cores and 4 GB of RAM. Using our smallest (200) IO Engine, with 4 SSDs and small, fully random IOs (which are the most difficult in terms of performance), we have shown performance at AWS of about 15,000 IOPS. Better yet, one can always upgrade or downgrade among these engines, seamlessly and without disrupting the applications running on Zadara.

### Zadara Container Engine

Each IO Engine is accompanied by a Container Engine. This second engine runs Docker containers inside the storage controllers, offering near-zero latency to persistent, sharable storage and access to the full set of storage features, including snapshots, remote mirroring, encryption, and automatic Backup to Object Storage (B2OS). The Container Engines have their own dedicated CPU cores and RAM, and come in 5 sizes: 01, 02, 04, 06 and 08. Combining them with the IO Engines offers combinations like virtual private storage array 204, 808, 1602, and so forth.

## Performance

Zadara offers features and functions to enhance what you get from AWS directly. And with AWS DirectConnect, your Zadara Cubes appear to be right inside your AWS environment. See the chart below for a summary of the features you get with Zadara.

| | | AWS (EFS, FSx, EBS) | | Zadara |
|---|---|---|---|---|
| NAS & SAN | Universal File Protocol Support | NFSv4 | SMB | NFSv4 + SMB |
| | Active Directory ACLs | ✗ | | ✓ |
| | Large File Volumes | PBs | 64 TB | 250 TB |
| | Large Block Volumes | 16 TB | | 250 TB |
| Privacy | Encryption of Data-at-Rest | ✓ | | ✓ |
| | Customer-Managed Encryption Keys | S3 / EBS Only | | ✓ |
| | Dedicated, On-Demand Hardware | ✗ | | ✓ |
| Perform. | High IOPS per Volume | 20,000 (PIOPS) | | 150,000 |
| | SSD/HDD Selection | EBS Only | | ✓ |
| | Embedded Containers | ✗ | | ✓ |
| Location | On-Premises / Hybrid Cloud Option | ✗ | | ✓ |
| | Multi-Cloud Access | ✗ | | ✓ |
| HA/DR | Clusters/Sharing | ✗ | | ✓ |
| | Instant Snapshots & Clones | ✗ | | ✓ |
| | Remote Replication | ✗ | | ✓ |
| | Availability SLA | 99.95% | | 100% |

## Summary

In summary, Zadara virtual private storage arrays not only solve the storage performance challenges of the cloud, they do so elegantly and flexibly, not to mention affordably.

**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 enteprise storage can help you transform your business. Call or email today.

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