

The Zadara Storage Cloud—A Validation of its Use Cases and Economic Benefits

By John Webster

August 2015



Evaluator Group

Enabling you to make the best technology decisions



Why Build a New Storage Model?

IT organizations are increasingly challenged to operate more efficiently. The cost to deliver a unit of compute to business users, however that unit may be measured, is now under scrutiny from enterprise executives. They have become aware of the comparatively low cost of computing offered by the now well-known cloud services providers like Amazon Web Services (AWS) and Microsoft Azure and are now comparing that to the cost of delivering IT services internally. As a result, IT is challenged to develop a far more cost-efficient computing environment.

In addition, business user groups are demanding a new level of IT agility. While enterprise application users have traditionally been at odds with centralized IT over responding to needs of business groups, cloud service providers have set a new “instant on” standard that IT is now measured against. Enterprise IT must therefore quickly adopt new technologies and application delivery models to remain competitive.

Unfortunately, however, most IT organizations are not getting budget increases that will allow them to keep up with the demand. Paying for the infrastructure that could support the expected services is in many cases simply not possible. As a result, business user groups go outside of centralized IT to get the needed application services, either by going directly to cloud service providers or by creating their own departmental IT environments commonly known as “shadow IT.” Either way, the enterprise executives can wind up spending more on IT services—sometimes much more—without realizing it because they are only looking at the centralized IT budget to determine the overall cost of IT.

IT is challenged once again to transform itself into something else—a process that has taken on a number of different monikers including IT as a Service (ITaaS), IT Agility, and Private Cloud. It is important to understand that transforming the storage environment is essential to answering the challenge. Here, we review a new storage acquisition and deployment model introduced by Zadara Storage. As part of this review, we examine the total cost of ownership (TCO) and reflect the experience of three Zadara users to validate our findings.

Zadara SAN/NAS—In the Cloud and On-Premises as a Service (OPaaS)

Zadara offers a new enterprise storage acquisition and consumption model that combines the features of cloud storage with the ability to locate and manage the storage environment either within the four walls of the enterprise data center or within a public cloud environment such as Amazon Web Services. Zadara storage is actually one storage platform—the Zadara Storage Cloud—with different deployment options. These include:

- On-Premises as a Service (OPaaS) storage arrays

- Cloud-based Storage as a Service (STaaS) from major service provider/colocation facilities worldwide¹
- Hybrid deployments that combine on premises and cloud-based deployments of the Zadara platform

The Zadara Storage Cloud was designed from the beginning to be deployed in a cloud computing, IT as a Service environment. For that reason, it is fundamentally different from a traditional storage array. Attributes built into the platform include:

- Scalable CPU and data transport
- Native multi-tenancy and workload isolation with the individual tenant administration
- Encryption for data in motion and data at rest (*encryption keys are kept by user, not by Zadara*)
- Accessible via REST APIs for openness and portability
- Cost control features accessible by administrators

Zadara Storage Cloud Architecture

The Zadara Storage Cloud software-defined appliance currently consists of interlinked storage nodes—commodity servers with embedded flash and disk storage running Linux and using 40GbE for internode and external server connectivity. A minimum configuration consists of at least two nodes for redundancy and can currently scale up to 600 nodes. Each of the major computing components—CPU, node storage, and connectivity—can be scaled upward and downward as needed. Storage Cloud can be scaled in three dimensions independently:

- CPU and memory resources for storage controllers
- Drive resources for raw storage
- Internal and external connectivity

RAID groups are built from flash and/or spinning disk drives spread across nodes rather than by creating three copies of data for protection as is common among the open source storage platforms. To address the potential for internal I/O latency, Zadara implements iSCSI Extensions for RDMA (iSER) for 40Gb per second connectivity among nodes. iSER connectivity to the server environment is also available as an option along with traditional iSCSI. NFS and CIFS/SMB protocols are also supported. And, rather than using a pair of clustered storage controllers for failure tolerance, Zadara implements controllers as pairs of KVM-based virtual machines that are managed by OpenStack.

¹ Presently these locations include Amazon Web Services (AWS), Microsoft Azure, CloudSigma, Dimension Data, FlexiScale, SerenITaaS, Equinix, CoreSite, TelecityGroup, Eircom, SteelORCA and others.

The VPSA

Initially, Zadara introduced a fully turnkey, SAN or NAS storage solution consisting of Virtual Private Storage Array (VPSA) software pre-loaded onto commodity hardware (Dell and SuperMicro presently supported) and located at cloud datacenters— AWS² and later Microsoft Azure. Now VPSAs can also be located in a customer's data center (referred to by Zadara as On Premises as a Service—OPaaS). Either way, isolation among VPSAs in a Storage Cloud is assured by Zadara.

The VPSA is created by an administrator who assigns a number of dedicated CPU cores, RAM and storage devices (SSD for cache, SSD and/or HDD for capacity), and the associated interconnections to the VPSA. It has the attributes of a traditional enterprise storage array with dual high availability controllers, SSD cache, and the advanced data management and protection features such as snapshots, cloning and remote replication—including the ability to replicate data from on-premises storage into any of the global public Zadara deployments (AWS, Microsoft Azure, etc.) or among these public cloud deployments.

Each VPSA can be assigned, for example, to individual applications or business user groups. And, as a key management and cost control attribute of the VPSA, dedicated physical resources can be non-disruptively added or subtracted in real time. This ability allows the Zadara VPSA user to be charged only for the resources used on a per hour basis.

Security

Security and cloud are two words that, when used in combination, Enterprise IT administrators and storage administrators in particular are very sensitive to. Consequently, the Zadara Storage Cloud offers multi-level security features:

- Workload isolation
- Secure access to individual VPSAs via dedicated VLANs
- Encrypted storage devices per VPSA with customer-managed keys per VPSA
- Client mapping per VPSA—a client can only access the LUNs or shares in a VPSA that are assigned to a VPSA and that client is specifically permitted to access (including support for common domain controllers).
- IPSec encryption of data in transit between the array and clients as well as between VPSAs locally or remotely

² Zadara VPSA cloud services partners now include: Amazon Web Services (AWS), Microsoft Azure, Dimension Data, Equinix, TelecityGroup, Eircom, CoreSite, Steel ORCA

Customer Interviews

In order to determine the viability of Zadara’s new approach in a production IT setting, we interviewed three Zadara users from various vertical industry segments with differing IT requirements.

Public University

The first user we interviewed was at a public university in central California. Most of the storage in the University’s centralized data center was coming off of warranty and IT had to develop a replacement strategy. They initially chose a hyper-converged system with SSDs for critical applications and then Zadara as bulk storage for everything else. This user reported that Zadara allowed them to set up an Amazon S3-like service in-house. Because of budget constraints, they worried at first that unpredictable growth in the new Zadara storage environment would cause a financial problem. However, they see on an hourly basis how much they are using, and are therefore able to make adjustments to capacity usage on a daily or even hourly basis using the management interface. Due to budgeting rules, the University’s IT department does not have the ability to roll over budget money from one year to next. Zadara was the solution for unpredictable capacity needs.

Zadara is now the storage for centralized IT that does not go on the hyper-converged system, which is judged to be more expensive. In addition, they are starting to offer Zadara as a “storage service” to other departmental IT groups and will charge-back each department based on data from the Zadara management application. They are also using Zadara as a backup repository in some cases and are replicating data to Zadara in cloud service provider (CSP) sites for disaster recovery.

In the future, they plan to use Zadara for critical applications and as overflow for the hyper-converged solution, particularly for applications with larger data requirements. They are also looking at deploying Zadara at another University site and replicating between sites.

Overall, this user likes the flexibility and the enhancement to IT services Zadara provides and is particularly impressed with Zadara’s storage isolation in a multi-tenant environment. This allows centralized IT to consolidate departmental storage, achieving economies of scale, but avoiding the “noisy neighbor” problem where one user’s heavy usage will impact other users’ performance. Centralized IT can also give departmental groups their own management portal to self-provision storage on-demand.

Large Vertical Industry Software Solutions Provider

We spoke with a large supplier of application software solutions developed for specific vertical industries. These solutions are delivered as cloud-based applications. The software firm initially delivered their applications from their own on premises IT facility but decided to move their application services delivery engine to a large cloud services provider. Zadara Storage was installed and used within the selected CSP’s data centers and data (approximately 1 PB) is still being migrated.

This user could have employed the storage services offer by the CSP, but chose the Zadara VPSA instead for a number of reasons that included:

- Better security and storage isolation in a multitenant environment
- The ability to control the storage environment and deploy resources in ways that were not supported if they were to use the CSP's storage
- Better storage performance, availability, and reliability.

And, because for this user the cost of cloud services and specifically storage services is a basic factor in determining overall profitability, the Zadara implementation with its lower cost results in greater profitability and provides greater visibility into monthly cost metrics. This user also liked the ability to “spin up and spin down” storage resources as required. However, in reality, they are still growing the storage environment and have yet to take resources off line. Zadara is now their “go to” primary cloud storage environment.

On Line Printing Service

This user hosted their own NAS storage before replacing it with Zadara when they ran out of capacity. Their NAS environment was a combination of units from a commercial vendor and their own software defined storage. With these NAS systems they were quickly out of capacity, forcing them to purchase additional storage units frequently. With their own software defined NAS they were continually juggling the cost of additional servers and drives along with their associated warranty periods. They also found configuration and administration to be a time-consuming process.

Now using Zadara for primary application storage, they eliminated the need to buy servers and drives. They also eliminated other infrastructure needs like data cabinets, networking gear and power. They reported that administration is easy and straightforward but that maintenance and support is where they are saving the most. They also now have better visibility into their storage via the Zadara dashboard. Adding capacity is now “point and click,” and they are able to deliver it to applications immediately without any changes to their application environment. This user also reported that Zadara support and maintenance was excellent.

Overall, this user likes the fact that Zadara storage is very flexible and that capacity can be added immediately in a high growth environment as compared to what they were using previously. They also reported that performance was “way better”. Their exposure to outages has also greatly improved given the storage redundancy and high availability associated with the Zadara architecture. These features have been put to the test a few times already and they have always been reliable.

Interview Summary

In general, we found that all of the users we spoke to were using Zadara VPSAs for primary storage supporting critical applications. One was in the process of migrating a petabyte of data to Zadara to support customer-facing applications. All reported better performance with Zadara versus what they

had previously experienced. All were impressed with Zadara’s level of service, storage expertise and support. The one potential negative perception expressed by two of them was their sense that the management GUI was still somewhat immature but expected Zadara to deliver improvements.

Storage Immortality and TCO Impact

A significant management capability that comes with the Zadara Storage Cloud can be thought of as storage “immortality”. Zadara will automatically, and for the life of the Storage Cloud, replace and upgrade hardware/software without application impact. Zadara manages storage node replacement and data migration from the obsolete node to the new node. The only thing required of the customer is racking of the new equipment and un-racking of the old. Node replacement does not incur additional charges over and above the normal monthly usage fee.

This is a critical cost-saving and management simplification aspect of the Storage Cloud when compared to the way storage array replacement is typically handled. IT administrators no longer have to plan for, acquire, and migrate data from old to new in three-year cycles. This has two major economic benefits that will increase if the number of traditional arrays replaced is increased:

1. Elimination of storage system refresh cycles: Technology and capacity refreshes involving the replacement of an entire storage system are no longer required when performance and capacity is added non-disruptively on a module-by-module basis. This eliminates the costs associated with technology refresh cycles and the need to manage data migrations and capital asset amortization schedules. The ability to non-disruptively add/replace modules as new technology emerges and enters mainstream usage (SSD, higher performance I/O ports, etc.) also helps to ensure the long-term quality of service delivery.
2. Elimination of data migration costs: The significant costs associated with migrating data from an obsolete system to a new system are eliminated as well. These costs include those related to system downtime and a corresponding loss in productivity, as well as an extension of lease terms at premium rates when the replaced system is leased and not returned on schedule.

To demonstrate the potential for cost savings that can result from “immortalizing” the physical storage platform, we have calculated the total cost of ownership (TCO) of deploying traditional arrays over a ten-year time period and compared that to the TCO for equivalent capacity using Zadara Storage Cloud (see Figure 1 below).

zadara-TCO comparison.swf




Figure 1. Ten-year cumulative TCO comparison for a traditional array (purchased and financed) versus Zadara Storage cloud for equivalent capacity.

In Figure 1 above, we started (year 1) with the same initial capacity of 224 TB and applied an annual capacity growth rate of 35% out to year 10. We included in our TCO calculation the cost for storage capacity and controllers, software maintenance, and administrative staff time. For the traditional array, we also added migration costs at three-year intervals. Note that after ten years, the total cumulative cost of deploying a traditional array is twice that of Zadara. We also note that a roughly 2x cost differential actually begins with year 1 and is maintained over the ten-year period.

A similar pattern can be seen when we make the same comparison on a yearly (rather than cumulative) TCO basis (Figure 2 below)

Zadara-TCO comparison.swf

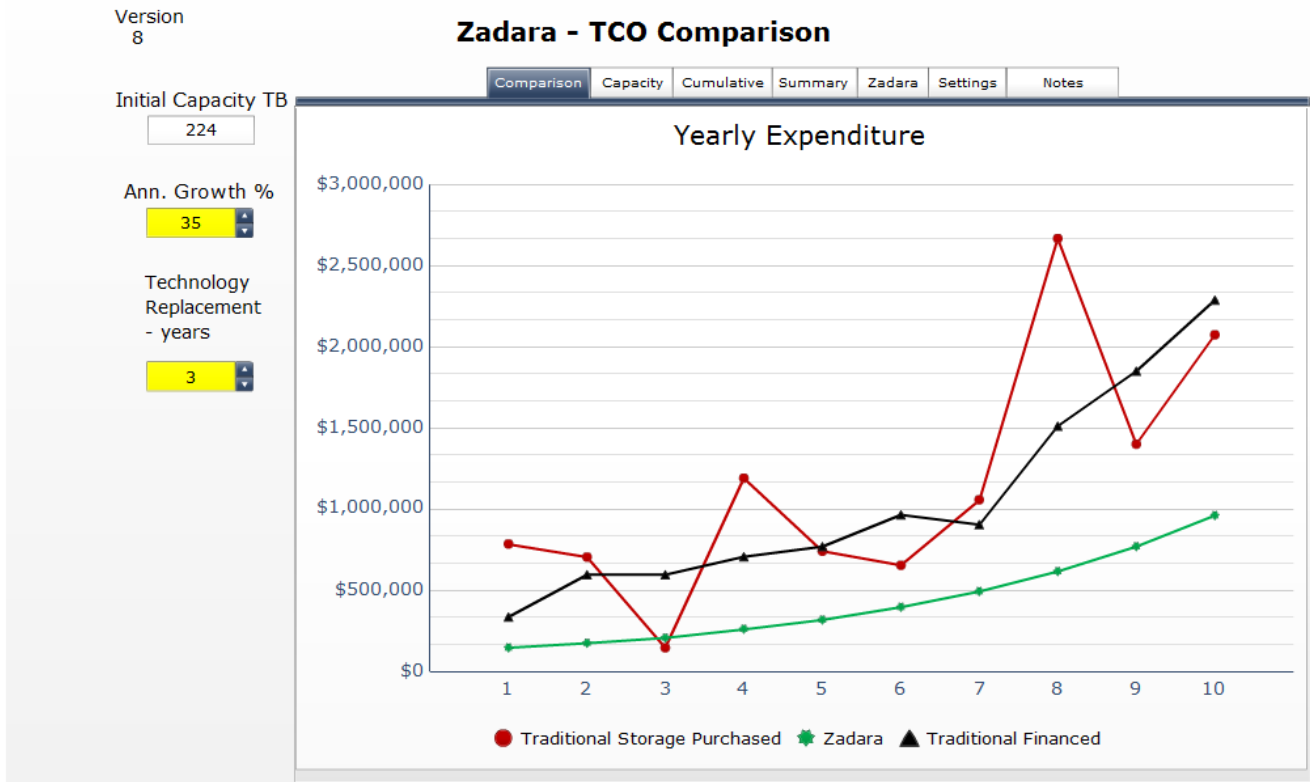


Figure 2. Ten-year yearly TCO comparison for a traditional array (purchased and financed) versus Zadara Storage cloud for equivalent capacity.

We again note that a roughly 2x yearly cost differential begins with year 1 and is maintained over ten years with the exception of the period in year 3 where the TCO of traditional is \$144,800. Because this comparison also assumes that capacity is added in 300 TB increments, we did not need to add capacity in year 3. The capacity growth used to make these TCO comparisons is shown in Figure 3.

Zadara-TCO comparison.swf

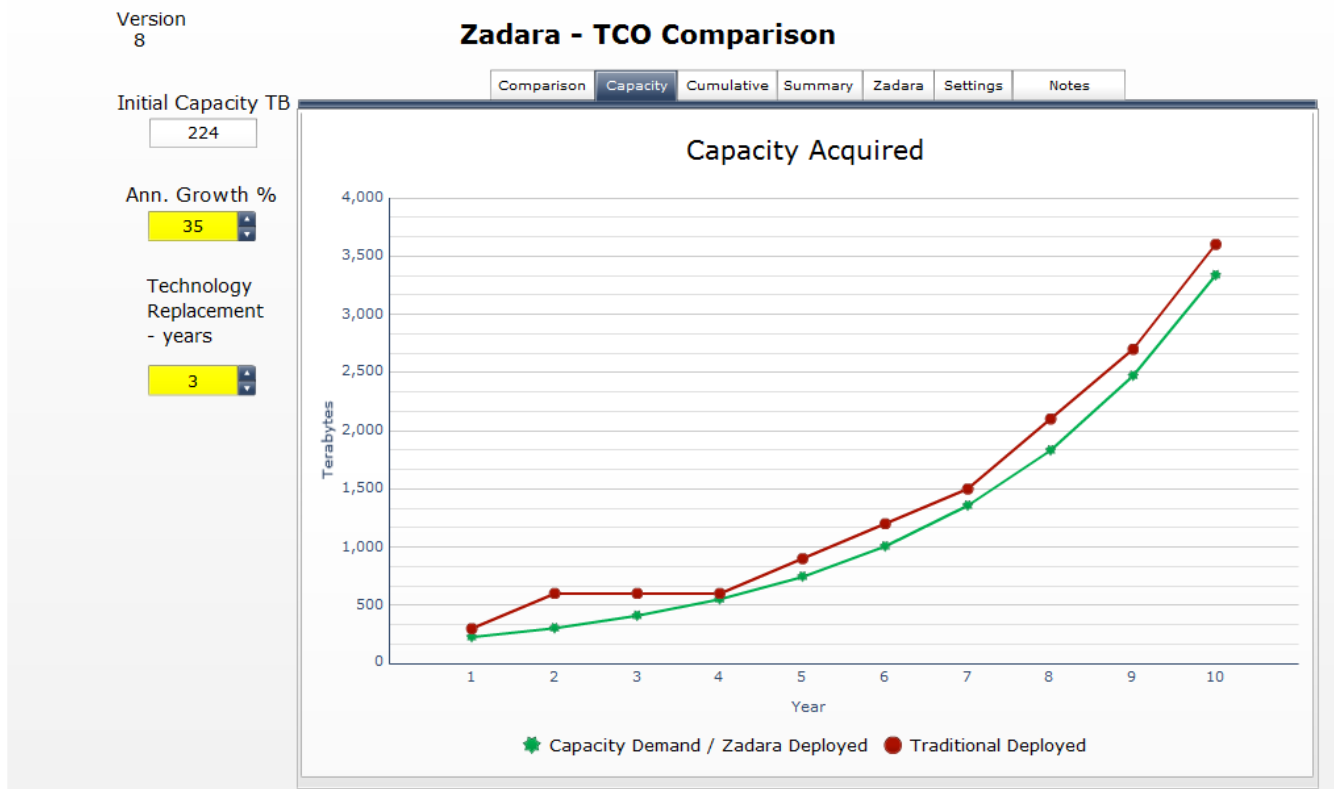


Figure 3. Comparison of capacity growth at a 35% annual growth rate for both the traditional array and Zadara over the ten-year TCO study period.

For more information on the use of this TCO comparison model that was created by Evaluator Group, contact [Zadara Storage](#).

In general, we believe that deploying the Zadara Storage Cloud could cut overall storage costs by roughly 50% yearly and cumulatively over a ten-year period. In addition, capital expenses are avoided because the cost for Zadara storage is entirely an operating expense, freeing up capital budget money for other IT projects.

The Acquisition Model

Zadara delivers VPSAs to users via public cloud partners and/or as an on premises solution. Either way, the monthly cost from Zadara is the same³. Monthly charges are based on the number of drives plus the type of VPSA engine used in one-hour time increments⁴. The type of VPSA engine varies on the basis of I/O performance as follows:

Zadara VPSA Engine Type	Cache Size and Performance
200 (Baby)	2 CPUs, 4GB Memory
400 (Basic)	4 CPUs, 8GB Memory
600 (Boost)	6 CPUs, 16GB Memory
800 (Blast)	8 CPUs, 24GB Memory
1000 (Blazing)	10 CPUs, 32GB Memory

Regardless of how much storage Zadara has installed at a customer's on premise site, a customer is only charged for the drives and engines they are actually using.

Evaluator Group Assessment

Many enterprise IT administrators are currently in the process of transforming their infrastructures to emphasize services delivery. Some may call it a transformation to cloud-based services—private and hybrid—while others may not. But finding ways to call it a transition to the cloud is not really the point. The objective of transformation is to deploy a services delivery engine that can simultaneously support existing applications while delivering new mobile, social, and analytics applications—and to deliver agile application services in an “always-on” way.

Zadara has generated traction both with smaller start-up companies and well-established large enterprises for roughly the same reason—both can have highly variable needs for storage. Small companies can start with entry configurations and grow capacity quickly without a large up-front commitment to overprovisioning storage. If business conditions change, capacity can be scaled back without financial penalty. Likewise, a large data center may have an immediate need for storage to support a new project where capacity usage may be small at first but grow rapidly once the project gets off the ground. Again, a small configuration could be used to get started immediately and scaled upward on demand and scaled back or eliminated entirely at the end of the project.

³ Zadara notes that installing on premises will usually cost less because the costs for cloud data center floor space and environments are avoided.

⁴ Zadara requires a 6-month minimum commitment for on premises (OPaaS) installations.

Based on our TCO analysis, we find that the Zadara Storage cloud is a very economical enterprise solution for storage in both public and private cloud environments. And based on our customer interviews, it is clear that Zadara storage can be deployed as scalable, lower cost primary storage for critical applications versus traditional enterprise arrays or the storage services offered by CSPs.

About Evaluator Group

Evaluator Group Inc. is a technology research and advisory company covering Information Management, Storage and Systems. Executives and IT Managers use us daily to make informed decisions to architect and purchase systems supporting their digital data. We get beyond the technology landscape by defining requirements and knowing the products in-depth along with the intricacies that dictate long-term successful strategies. www.evaluatorgroup.com @evaluator_group

Copyright 2015 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.