

# Parallels Cloud Storage

---

White Paper

## The Ideal Storage Solution for Hosters

# Table of Contents

- Introduction ..... 3
- The Problem with Direct Attached Storage ..... 3
- The Solution: Parallels Cloud Storage ..... 4
  - Eliminates Downtime Due to Hardware Failures..... 4
  - Reduces Server Maintenance and Support Costs ..... 5
  - Provides a Cost-Effective SAN Alternative ..... 6
  - The Future: Enabling Storage-as-a-Service Solutions ..... 7
- Parallels Cloud Storage: a Component of Parallels Cloud Server ..... 8
- Conclusion ..... 8
- Contact Us ..... 8

---

## Introduction

Most service providers today provision their servers with direct attached storage (DAS). While DAS is fast and inexpensive to procure, it has a high overall cost of ownership and high failure rates that can significantly impact service availability. In order to address the issues with DAS, many service providers have considered storage area networks (SAN) as an alternative; however these are typically too costly for low cost hosting.

This whitepaper describes how Parallels Cloud Storage reduces total operation cost and compares the capabilities and relative cost of Parallels Cloud Storage to DAS and Storage Area Network (SAN) alternatives.

---

## The Problem with Direct Attached Storage

Direct attached storage typically consists of a large machine with multiple disks in a RAID configuration. The advantage of DAS is that it's fast (100MB/s), inexpensive to procure, and scalable – it is easy to add more nodes to increase total storage.

However, the disadvantages of DAS are significant. To begin with, DAS makes the hard disk drive (HDD) a single point of server failure – a major issue for hosters, for whom customer satisfaction depends on high availability. At the same time, DAS prevents you from rapidly migrating your virtual servers to a functional node when a failure does occur. The result is usually a loss of service for end users, resulting in disgruntled customers and higher rates of churn.

A second drawback of DAS is that it's a local solution, limiting available storage to virtual servers attached to that particular node. The total amount of storage available to any particular virtual server on the node is limited to the number of drives and associated storage you can put in the server.

A third problem with DAS is that it doesn't allow you to make efficient use of your available disk space. In fact, surveys we have done of our large hosting provider partners in Europe and North America have shown that they typically are able to use only 36% of their disk space (see Figure 1). This inability to use disk storage efficiently greatly increases the real cost of usable storage.




Figure 1. Parallels research shows that large hosting providers in Europe and North America typically use only about 36% of their available disk space.

For all of these reasons, DAS is an inefficient storage solution even for traditional hosting environments – and it is especially inappropriate for dynamic, highly fluid environments, such as those found in the cloud.

---

## The Solution: Parallels Cloud Storage

Parallels Cloud Storage solves the problems associated with DAS by providing a distributed, shared solution that decouples storage from computation. This approach enables virtual machines and containers to be instantly migrated to an operational physical server if the original physical server becomes unavailable. And with the storage cluster distributed across multiple physical machines, if one data source becomes unavailable, virtual machines can rapidly get their data from another data source.

The following sections discuss the benefits of Parallels Cloud Storage in more detail.

### ELIMINATES DOWNTIME DUE TO HARDWARE FAILURES

HDDs are the single most unreliable component in a server, failing at a typical annual rate of 4% according to recent mean-time-between-failure analyses<sup>1</sup> – a failure rate that has been corroborated in our recent customer surveys. These failures are no trivial matter: by the time you get the drive running again and the data copied to a new drive, each failure can easily result in one to two days of virtual private server (VPS) downtime. These failures are also expensive – not only in terms of the maintenance and support costs involved, but also in the loss of customers due to poor quality of service.

With Parallels Cloud Storage, VPSs will instantly and automatically be switched to a different storage location when a hard drive fails, thereby eliminating downtime and the associated costs of hard drive failures.

Nothing is more important to customers than service quality, especially when it comes to performance and uptime. Poor service quality will result in high churn rates – and churn is expensive to service providers, not only reducing revenue growth, but also increasing costs. That's because the cost of acquiring new customers (typically around \$100 per customer) is usually many times higher than the cost of keeping existing customers.

Table 1 applies the results of the surveys we've done of our hosting provider partners to show how the costs of hard drive failure can add up over a year's time. The table is based on a hypothetical scenario in which the hoster has 500 servers, with four hard drives per server. Given the typical annual failure rate of 4%, in a given year the hoster would experience 80 drive failures (500 X 4 X .04). Of course, not all the failures would have an impact on the quality of service – for example, some might be redundant drives in a RAID configuration. For the purposes of this example, we assume a conservative 20% of the failures as having an impact on service.

Based on industry norms, we also assume each hardware node supports 40 customers (the typical number of virtual machines per node); that each customer has an average lifespan of 52 months (with half of that lifespan remaining at the time of failure); that each customer brings in monthly revenues of \$30; and that the profit margin on those revenues is 30%. Since most customers are not going to stay with a hoster when their site is down for two days, we assume that the customers affected by a disk failure will most likely leave. So, even with our conservative assumptions, 80 hard drive failures per year X 20% (the number affecting service) X 40 customers per node X 26 months (average customer lifespan remaining) X \$30 per month adds up to lost revenues of nearly \$500,000 – and lost profits of nearly \$150,000.

**Table 1: Annual Lost Revenues Due to Churn Resulting from HDD Failures**

<b>Hardware Assumptions</b>	
No. of servers	<b>500</b>
No. of HDDs per server	<b>4</b>
Replacement rate per year	<b>4%</b>
No. of HDD failures per year (500 X 4 X .04)	<b>80</b>
<b>Lost Revenue and Profit Due to Churn</b>	
% of HDD failures that have a service impact (assumption)	<b>20%</b>
Typical no. of customers per hardware node (assuming 40 VMs per node)	<b>40</b>
Average customer life span (per industry norms)	<b>52 months</b>
Average remaining life span at time of failure	<b>50%</b>
Average monthly revenue per customer (per industry norms)	<b>\$30.00</b>
<b>Lost Revenue/Yr (80 X .2 X 40 X 26 X 30)</b>	<b>\$499,200.00</b>
Profit Margin	<b>30%</b>
<b>Lost Profits /Yr (\$499,200 X .30)</b>	<b>\$149,760.00</b>

### **REDUCES SERVER MAINTENANCE AND SUPPORT COSTS**

Lost customer revenues are not the only expense associated with hard drive failures. There are also significant maintenance and support costs, which again can add up over a year's time. By eliminating the impact of HDD failures, Parallels Cloud Storage significantly reduces both maintenance and support costs. These savings are a result of:

- A reduction in maintenance staff, because there is no longer need to keep 24x7 support staff available to carry out fast hard drive replacement
- Fewer support incidents (e.g., no storm of support calls after a hardware node failure)

Table 2 represents another hypothetical example, based on the same 80 failures per year. For maintenance costs, it uses industry norms of 0.25 hours to replace a failed hard drive; an average annual cost of \$120,000 per full-time employee (based on salary plus overhead); and a total of 2,080 hours worked per year (based on 52 40-hour weeks), resulting in an average hourly rate of \$57.69 per employee. Multiplying 80 drive failures X .25 hour per drive X \$57.69 per hour, we see that a year's worth of hard drive failures creates more than \$1,150 of maintenance costs.

Then there are support costs – because the affected customers are going to be calling in to the help desk to find out why their site is down. If we make the same assumptions as in Table 1 (80 hard drive failures, 20% of them affecting service, and 40 customers per node) and use the industry norms of three help desk calls apiece over the two-day outage, at an average cost per call of \$30, we see that over a year's time, hard drive failures can add up to more than \$57,000 in support costs (80 HDD failures X 20% of affecting customers X 40 affected customers per failed node X 3 calls per customer X \$30).

**Table 2: Annual HDD Maintenance and Support Costs**

<b>Maintenance Costs Due to HDD Failures</b>	
No. of HDD failures	<b>80</b>
Time to replace (hrs per HDD)	<b>0.25</b>
Average annual salary of FTE (including overhead)	<b>\$120,000.00</b>
Average hourly rate of FTE (based on 2,080 hrs/yr)	<b>\$57.69</b>
<b>HDD Maintenance Cost (80 X 0.25 X \$57.69)</b>	<b>\$1,153.85</b>
<b>Support Costs Due to HDD Failures</b>	
No. of HDD failures	<b>80</b>
% of service impact (assumption)	<b>20%</b>
No. of customers per hardware node	<b>40</b>
No. of help desk calls per customer	<b>3</b>
Cost per call (industry norm)	<b>\$30.00</b>
<b>Support Cost for HDD Failures (80 X .2 X 40 X 3 X \$30)</b>	<b>\$57,600.00</b>

**PROVIDES A COST-EFFECTIVE SAN ALTERNATIVE**

Hosting customers always will be happiest with a service provider that is able to deliver high performance and high availability on an ongoing basis. In the past, the only way to do this was to add a storage-area network (SAN) – a high-speed, special-purpose network that interconnects multiple storage devices and servers. Typically optimized for enterprise environments, SANS are generally too expensive to be used for low-cost hosted services. Open-source clustered storage solutions, such as Gluster, also exist, but they don't offer sufficient performance at a scale required for hosting.

With Parallels Cloud Storage, however, hosting providers can achieve SAN-like performance and availability without having to purchase expensive dedicated storage arrays or networks. That's because the Parallels solution, which is built on commodity SATA direct-attached storage and a 1 GigE networking infrastructure, leverages providers' existing infrastructure. Instead of using high-cost, dedicated storage arrays and additional networking to connect them all, Parallels Cloud Storage creates a storage pool from a service provider's existing, unused disk space – the unused 64% shown in Figure 1, above. Because this existing infrastructure has already been purchased and is already maintained and supported, the incremental capital and operational cost of leveraging it for cloud storage solutions is minimal to nonexistent.

As Table 3 illustrates, Parallels Cloud Storage provides all of the features and benefits of a SAN, at a much lower cost per TB of data. The table also illustrates the drawbacks and lack of functionality inherent in DAS – the storage option that most service providers are using today.

In comparing costs, it is important to compare not only the procurement cost of storage, but also its true cost in terms of the percentage of available storage actually being used. We know from the study summarized in Figure 1 that with DAS, only about 36% of available storage is actually being used.

Therefore, even though the initial procurement cost is slightly higher for Parallels Cloud Storage than for DAS, the actual cost per TB of usable storage is much lower for Parallels Cloud Storage than for either DAS or SAN. Yet its functionality and benefits are much better than those of DAS, and equivalent to those of SAN.

**Table 3: Cost Comparison between Parallel Cloud Storage and Other Storage Options**

	Capabilities, Features							Costs		
	No VPS down time <sup>1</sup>	Replias <sup>2</sup>	High Availability	Snapshots	Unlimited storage <sup>3</sup>	DAS-like performance	Scalability >1PB	Procurement cost per TB	% of storage utilized	Procurement cost per usable TB <sup>4</sup>
<b>DAS</b>						.		<b>\$50</b>	<b>40%</b>	<b>\$125</b>
<b>iSCSI SAN</b>	.	.	.	.	.	.	.	<b>\$100</b>	<b>80%</b>	<b>\$125</b>
<b>FC SAN</b>	.	.	.	.	.	.	.	<b>\$140</b>	<b>80%</b>	<b>\$175</b>
<b>Clusters (CEPH, Gluster)</b>	.	.	.	.	.			<b>N/A</b>	<b>N/A</b>	<b>N/A</b>
<b>Parallels Cloud Storage</b>	.	.	.	.	.	.	.	<b>\$60<sup>5</sup></b>	<b>80%</b>	<b>\$75<sup>5</sup></b>

<sup>1</sup>Refers to downtime due to hardware failures.

<sup>2</sup>Refers to the ability to configure any number of replicas.

<sup>3</sup>Storage size is able to grow on demand, unlimited by what's physically attached to the server.

<sup>4</sup>Based on percentage of available storage actually used.

<sup>5</sup>Based on license fee per TB.

### THE FUTURE: ENABLING STORAGE-AS-A-SERVICE SOLUTIONS

Future releases of Parallels Cloud Storage will allow its storage to be accessed by third-party applications, enabling service providers to offer their clients storage-as-a-service solutions. This cloud storage can be used for a variety of purposes, including content storage, backup, archiving, and disaster recovery.

---

## Parallels Cloud Storage: a Component of Parallels Cloud Server

Parallels Cloud Storage is delivered as a component of Parallels Cloud Server 6—the only solution available today that permits the coexistence of servers running container and hypervisor technology, as well as migration between these servers. By eliminating virtualization lock-in, Parallels Cloud Server lets you take advantage of the scalability and streamlined operations of operating system virtualization (containers) while also making use of high-performance virtual machines (hypervisors) for isolated applications or simultaneous heterogeneous environments. Parallels Cloud Server 6 comes with two virtualization options:

- **Parallels Containers for Linux**, which sets the hosting industry standard for profitable, secure, and flexible VPS and cloud server offerings. Uniquely suited to cloud server virtualization, it enables near-instant provisioning and on-the-fly modification of hosting and cloud server plans, while delivering up to three times better density and up to 350% better performance than leading hypervisors.
- **Parallels Hypervisor**, which gives you the flexibility to create virtual machines for customers who prefer to deploy and maintain a variety of operating systems simultaneously.

By combining both virtualization technologies into a single solution, Parallels Cloud Server provides you with the freedom to choose the most efficient technology for each virtual server.

---

## Conclusion

Parallels Cloud Storage is the ideal storage solution for service providers. It overcomes the single-point-of-failure and storage limitation drawbacks of DAS, while making much more efficient use of service providers' existing storage. It also overcomes the performance and scalability limitations of open-source clustered solutions. Better yet, it offers all the advantages of SAN storage, at a much lower cost. In short, Parallels Cloud Storage will significantly reduce your overall storage deployment costs while delivering performance equal to or better than that of DAS, without the high failure rate; and equal to that of SAN, without the high cost.

For more information about Parallels Cloud Storage, please visit [www.parallels.com/products/pcs](http://www.parallels.com/products/pcs).

### NOTES

<sup>1</sup> Bianca Schroeder and Garth A. Gibson, "Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to You?" Fast '07: 5th USENIX Conference on File and Storage Technologies. June 13, 2007, San Jose, CA. USENIX Association, 2007, p. 7.

---

## Contact Us

For more information about Parallels hosting and cloud solutions, please contact:

**Parallels, Inc.**  
**500 SW 39th St. Suite 200**  
**Renton, WA 98057**  
**+1 425 282 6448**  
**[www.parallels.com](http://www.parallels.com)**

Copyright© 2012 Parallels IP Holdings GmbH. All rights reserved. Parallels and the Parallels logo are registered trademarks of Parallels IP Holdings GmbH. Other product and company names are the trademarks or registered trademarks of their respective owners.