

Take Back Lost Revenue by Activating Virtuozzo Storage Today

JUNE, 2017

Introduction

New software-defined storage (SDS) solutions are enabling hosting companies to remove the complexity of legacy storage approaches – including locally attached drives – while opening up big opportunities to keep revenues up and provide better service to customers with high availability and no downtime.

According to the Gartner report, *How to Determine When Hyperconverged Integrated Systems Can Replace Traditional Storage*, “By 2019, more than 50% of the storage capacity installed in data centers will be deployed with SDS or HCIS architectures based on x86 commodity hardware systems, up from 10% today.”¹

Working directly with locally attached storage gets you solid local IOPS, bandwidth, and throughput results, but it means that you are relying on hardware that will inevitably fail at some point. In fact, while the annual rate of disk failure varies across manufacturers, in general, you can expect that 3 out of every 100 drives you depend on for your business will fail next year.

When you lose data from locally attached storage due to hardware failure, your only fallback option is to restore your data from a backup, which can take hours – even days. To compound matters, disasters tend to happen after hours, or on weekends and holidays. Depending on your backup schedule you’ll lose at least 12 hours of data.

From a support perspective, these failures require immediate attention from the IT operations team to ensure minimal customer impact. As the amount of data grows exponentially every year, using locally attached storage will only require a bigger team to manage it and more money to keep it up and running.

Are You Getting the Most of Your Hardware?

At Virtuozzo we found that prior to upgrading to SDS, our study of nearly 1500 servers showed that only 50% of all nodes were handling as much as 90% of IOPS. If you’re not using SDS today, this means that you have an opportunity to improve your IOPS by 2x by pooling all your hardware resources into a single cluster.

You can also overcome storage fragmentation with SDS. For example, when a node has close to 100% CPU and RAM utilization, storage utilization is approximately 40-60%. At the same time, when a node has close to 100% storage utilization, CPU and RAM typically are under-utilized. This happens because workload schedulers are only taking CPU and RAM resources into account when selecting the best node to place a workload on. Storage resources are not taken into account.

Software-defined storage solutions enable you to overcome these challenges by decoupling storage from CPU/RAM and treating storage resources as single, standalone pool with high availability and built-in redundancy.

Lost Revenue Due to Lack of High Availability

Customer churn is one of the biggest challenges facing hosting companies today. Let's assume you have a small cluster of 20 nodes with 15 drives on each node, and you are selling VPS services to your customers. Let's also assume that you have about 500 customers with customer data stored on at least two hard drives. On average, with an annual disk failure rate of 3%, you can expect to replace nine hard drives per year. In addition to the cost and time needed for your team to fix the problem, more importantly, this means that at least 36 customers will face outages that could last hours – putting you at a major risk of customer churn. In fact, these common outages can contribute up to 6% of your annual churn rate, and consequently at least 6% in lost revenue.

Not convinced? Take a look at all the other things that can go wrong in this example below from Google data center operations:

“In each 1800 node cluster's first year, it's typical that 1,000 individual machine failures will occur; thousands of hard drive failures will occur; one power distribution unit will fail, bringing down 500 to 1,000 machines for about 6 hours; 20 racks will fail, each time causing 40 to 80 machines to vanish from the network; 5 racks will “go wonky,” with half their network packets missing in action; and the cluster will have to be rewired once, affecting 5% of the machines at any given moment over a 2-day span.”

Finally, shorter outages that impact more customers can leave you even more exposed to customer churn and lost revenue. Let's assume that only 25% of your nodes are impacted by an outage that is shorter, lasting 2-4 hours. In this case let's also assume that impacted customers have a 50% chance of moving to a different hosting provider. In this scenario, reactive maintenance and system failures can lead to an additional 12% to your churn base, or more.

The absence of highly available storage could contribute 15-20% of additional churn rate in your customer base, resulting in at least 15-20% in lost revenue.

Reactive Maintenance Disrupts Service, Increases Costs

There's no question that the biggest cost for keeping your systems up and running is paying the salaries of the highly skilled professionals that you rely on for your business. Hosting companies and service provider IT staff need to use their time wisely to deliver the most value to the business. Any disk or server failure requires immediate attention to ensure minimal customer disruption – and requires your staff to change priorities on a moment's notice. Since these failures are exceptions, and not normal events, they can happen at any time. As a result, you need more support engineers, covering all possible shifts to provide 24x7 coverage.

The second biggest cost to your business resulting from reactive maintenance tied to system downtime to perform software updates and maintenance. Typically, updates and maintenance require a node reboot that will impact the end customer in the same manner as a short outage. Your only option is to perform a live migration to spare nodes to perform the required maintenance. When using locally attached drives, the live migration of a single virtual machine could take hours, requiring you to copy the whole disk to a different node.

As a result, you're faced with either unhappy customers who may take their business elsewhere due to additional outages, or your spending too much time and money on installing updates in a non-disruptive way with an IT ops team that's constantly reacting to keep systems up and running.

Reactive and unplanned maintenance leads to downtime that impacts the service you deliver and costs your company money with your IT teams spending too much time fixing problems and performing updates.

How SDS Help You to Fight These Issues

SDS solutions treat hardware failure as the new normal. Running highly-available services on top of off-the-shelf x86 servers requires software with high levels of resilience to handle hardware failures and issues. SDS aggregates all existing hard drives within a cluster and represents them as a single highly available, highly redundant storage capacity pool. As a result, all storage resources are shared between all nodes. Due to built-in data redundancy, if one disk or node goes down, data is still available for the rest of the cluster, so workloads will continue running on top of remaining nodes without disruption.

How does this help you as a VPS provider?

1. Customers won't even notice system failures or issues – there is no data loss and no downtime.
2. Backups are only required for more serious recovery scenarios.
3. Support engineers can manage single disk failures or even single node failures and deal with them during standard business hours.
4. Replacement of the failed hardware becomes a routine and planned activity.
5. With data “shared” across the cluster, migrating VMs from a host for maintenance only and installing updates only with minimal impact customers.

How SDS Works

Software-defined storage enables you to create a single pool comprised of a group of physical storage devices that are available across the entire cluster. In the simplest form, replication-based redundancy storage is straightforward. For example, when a client writes a file to SDS, first, it will be represented as a set of fixed size chunks. One piece will be written locally, and then the storage solution will create two more copies of it over the network. Internally, the storage solution keeps three copies of data (see Figure 1).

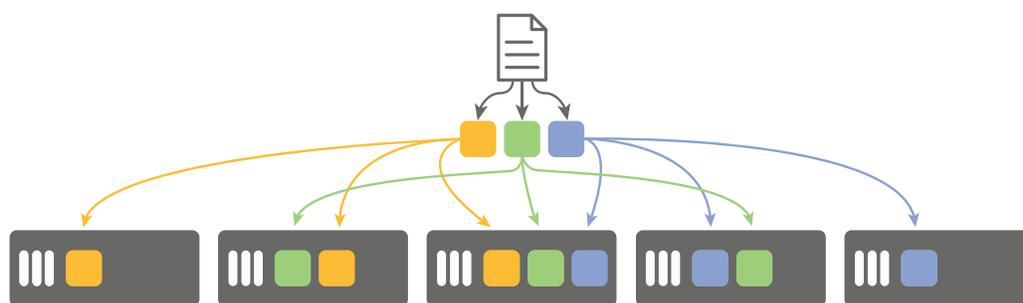


Figure 1

When hardware failure destroys certain data chunk copies, the storage solution will immediately create additional copies of the data in the background (see Figure 2).



Figure 2

Flexible Redundancy

In the case of Virtuozzo Storage, users get flexible redundancy with a replication factor of two or three. If two copies of data are maintained at all times, the cluster will survive the failure of a single node or drive. Three copies of the data provides resilience against two simultaneous failures.

The storage pool spans multiple nodes and expands as the cluster scales. A single pool eliminates fragmented data across individual nodes and disks, and represents storage as a single resource to the clients (see Figure 3).

Also, typically about half of active nodes carry up to 90% of disk operations. The remaining half are underutilized. Due to the distributed nature of SDS solutions, every node in a cluster can contribute to IOPS. As a result, overall cluster performance is not limited by fragmentation.

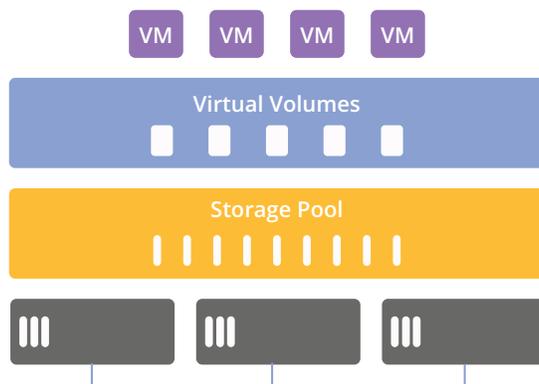


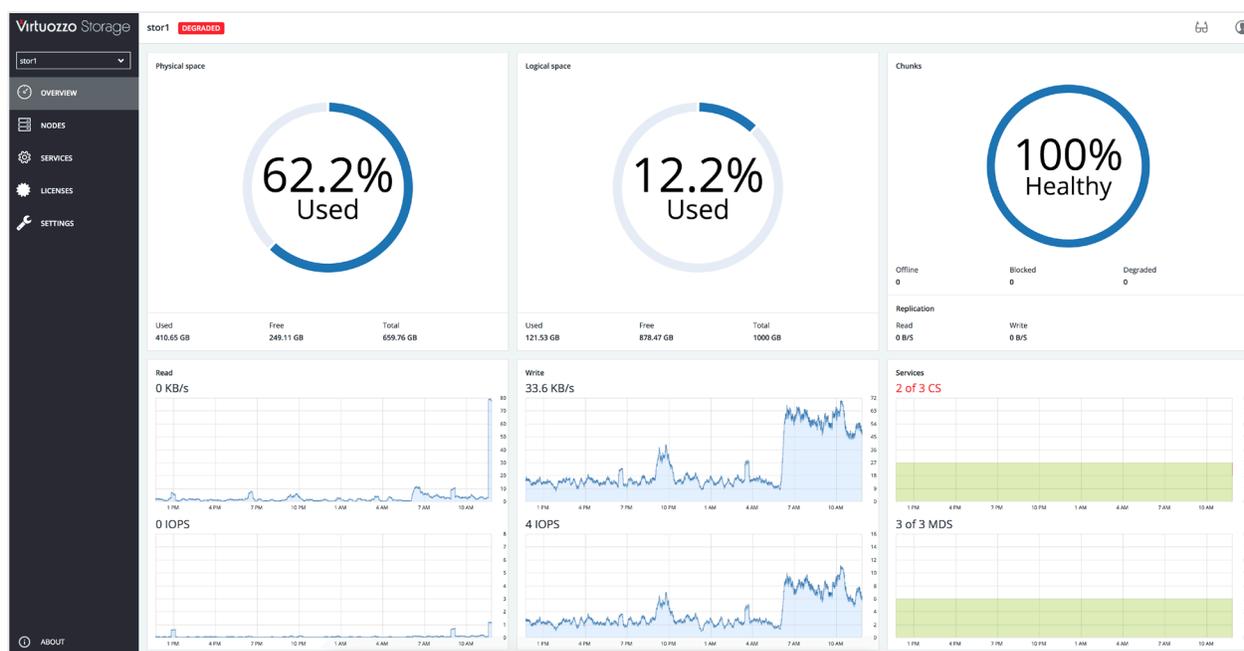
Figure 3

Because SDS decouples compute and storage resources, it enables virtual machines and containers to be instantly migrated to a new server whenever the original server becomes unavailable. Traditional virtualization brings effective compute resource allocation and usage. However, storage resources remain isolated and fragmented much longer. A software-defined storage approach can deliver greater levels of virtualization with a far more efficient use of computing resources, enabling you to take advantage of the performance and efficiency benefits of a hypercovered infrastructure.

Virtuozzo Storage: A Software-Defined Storage Solution Optimized for Hyperconvergence

Virtuozzo is a software defined storage solution integrated and used by the largest hosting and service providers across the globe. Extremely cost-efficient, Virtuozzo Storage is designed to work with commodity off-the-shelf hardware in hyperconverged infrastructure deployments.

- Block, file, NEW S3 object, and Docker container storage in a single solution
- Fault tolerant by design with flexible redundancy schemas: Erasure coding and 2/3 replicas support
- Faster than CEPH – Maximum performance with SSD caching and journaling, data auto-tiering and load balancing
- Efficient hardware utilization: low-cost scalability for up to multiple petabytes in a single cluster
- Hot plug for SSD/HDD to increase/decrease capacity or replace a drive without downtime
- Storage tiers to organize storage space and keep different categories of data on different disks
- Web-based UI management and smart automation



Virtuozzo Storage improves overall storage agility and provides a better scaling model for storage. With improved storage resiliency/reliability it is much simpler to deploy and scale out performance while benefiting from dramatic CAPEX/OPEX costs savings.

Activate Virtuozzo Storage Today

If you're a Virtuozzo customer, deploying Virtuozzo Storage provides a fast and low-risk way to improve the service you deliver to customers with high availability to help keep your revenues up, while enabling a hyperconverged infrastructure solution that can drive greater performance and cost efficiency.

Let Us Help You Get Started Today

Deploying Virtuozzo Storage doesn't require any CAPEX investments. It's built to run side by side with environments that use the same hardware running your virtual machines or containers today.

If you are a current Virtuozzo 6 or 7 user, you already have all of the Virtuozzo components installed. If you are running a previous version of Virtuozzo or are a new user, you need to get the latest version of Virtuozzo.

Here are some recommendations to get you started:

1. Install and test storage with even a single node. For production environments, you need at least five nodes; the cluster can survive the failure of two nodes without data loss. The bigger the cluster, the better Virtuozzo Storage performs. It is recommended to create production clusters from at least ten nodes for improved resiliency, performance, and fault tolerance in production scenarios.
2. Use separate networks (and ideally separate network adapters) for internal storage and external client traffic. Doing so will prevent public traffic from affecting cluster I/O performance.
3. Use storage tiers as a way to represent and organize storage space. You can use them to keep different categories of data on different drives. For example, you can use high-speed solid-state drives to store performance-critical data instead of caching cluster operations. Plan your storage tiers according to your capacity and performance requirements.

If you'd like to learn more, contact the Virtuozzo team to assist you in planning and deploying a storage cluster. We'll provide a detailed consultation and help you get a POC setup so you can take advantage of software-defined storage as part of a high performing, highly efficient hyperconverged infrastructure solution today.

[Contact us](#) for more information.