

# Pure Storage: #1 Most Deployed Platform for vVols



For many years, organizations across industries have counted on Pure Storage® FlashArray™ to power and protect virtual machines most commonly managed by VMware vSphere. This has been—and will continue to be—an important part of our customers' infrastructures.

Why do they choose Pure for [VMware](#)?

- The Evergreen Storage™ model's elimination of forklift upgrades
- Predictable and consistent performance
- Simplicity of Pure vs. traditional complex storage
- True "as-a-service" consumption
- Six-nines availability in FlashArray
- Pure's wide and deep VMware integration
- Affinity for the color orange ☐

I'm pleased to announce that Pure is now the number one provider of vVols. We have our telemetry data to show that you've actually surpassed every other storage company in adoption. And that's actually a huge tribute to the innovation you've helped drive the new standards that we have

Let me dive into a few of these within the context of VMware...

Before Pure came along, block storage in a VMware world typically used the Virtual Machine File System (VMFS). These VMFS datastores map 1:1 with a volume on an array. Arrays provide storage features on the volume level. In other words, the array features operate against an entire datastore. When you took an array snapshot, every VM on that datastore would be in that snapshot.

Consequently, you'd have to configure these datastores with a "best-guess" configuration to meet the needs of the VMs that were likely to be on it. This was... problematic. Take VDI, for instance, which has a very heavy workload in the morning, a lighter workload all day, and then gets heavy again to recompose or update. That's a fairly common VDI workload description: different needs during different parts of the day.

Volume configuration on an array was tied to features, which often translated into performance capabilities. For example, configure a volume like so, get efficiency *X*, with performance *Y*.

Datastores tied to spinning-disk arrays were either under-configured (users suffered during boot storms) or over-configured (wasted resources when the VDI requirements slowed down for the day). Configurations were usually some flavor of:

- **Lots of disks:** This gave the required IOPS but then sat idly for a lot of the day.
- **Lots of cache:** But by the time the cache warmed up, it didn't matter anymore, which puts us back to the first situation.
- **Tiers:** But tiering couldn't act fast enough to keep up with VDI demands.

This created complex, costly, and/or inefficient solutions.

## Then Pure Arrived

Pure has built a robust, resilient product on a foundational idea that many so-called "traditional features" from traditional storage vendors were actually limitations of design or tradeoffs disguised as features.

### **FlashArray delivers the following advantages:**

- Data reduction is always enabled and global to the entire array, so you don't have to decide whether it should be enabled or think about how to lay out your data to make the best use of it.
- RAID is automatically and intelligently configured and maintained.
- Encryption at REST is always on.
- Application type or block size is irrelevant since we manage back-end storage at a 512-byte scale.
- There's no thick or thin provisioning. FlashArray volumes aren't associated with cache, CPU, memory, or a specific number of flash modules. They're tied to the array. Provisioning one volume or a few volumes won't give you better performance from a given array. As a result, volume size, count, or configuration doesn't affect performance or efficiency.

We put a lot of thought and effort into these characteristics to make things as effortless and easy as possible, whether that means things tune automatically or just that you don't have to think about them. It's the job of our engineers to solve these problems, not you.

For example, when creating a volume on the FlashArray device, we just need the size. That's it.

This simplicity of storage configuration, along with superior data reduction, make VDI a fantastic use case

for FlashArray. You can use one VMFS datastore for it all. VDI gets the FlashArray performance it needs in the morning. Then, once users have started their days—and the morning boot/login storm has ended—performance capacity frees up.

Customers would ask us, “Why not put some databases on here too?” So our engineers did. Often, it’s when VDI slowed down that the databases would start to crank up, allowing consolidation into single arrays.

Over time, as customers moved more complex applications from traditional storage to Pure, they also needed more sophisticated features like snapshots, quality of service (QoS), replication, and more.

From a performance perspective, you could throw those databases on the same VDI VMFS datastore, but many users chose not to. Why? They wanted to use our array features—features that VDI often didn’t need at all or not in the same way. Snapshots are a prime example. Array snapshots can provide significant benefits to databases such as for restore and dev/test, but often you only need to restore or refresh one database from a snapshot. As mentioned earlier, with VMFS, an array snapshot includes the entire datastore—requiring a bit more work to orchestrate and not offering the granularity that is preferred.

So, now you have these great array features, but you can’t quite use them as intended in VMware environments unless you sacrifice something. RDMs or in-guest iSCSI? Choose wisely.

How can we solve this trade-off? Do we build integrations to abstract the disparity in granularity between a virtual machine disk and a FlashArray volume? No. A solution for this already existed: Virtual Volumes (vVols).

*I’m pleased to announce that Pure is now the number one platform of vVols. We have our telemetry data to show that you’ve actually surpassed every other storage platform in adoption. And that’s actually a huge tribute to the innovation you’ve helped drive the new standards that we have. Lee Caswell, VP of Marketing, Cloud Platform BU, VMware*

## Next, Virtual Volumes Came Along

Virtual Volumes are a natural choice to solve this trade-off because this is exactly what they’re intended to do.

We looked at vVols and the issues with VMFS. Are they inherent and fundamental? Or are they opportunities? At the time, there were also issues around vVols, including the lack of replication as well as a lack of adoption and customer interest. Were those inherent and fundamental as well? Or were they opportunities?

I would define a problem that is inherent and fundamental to a solution as a core characteristic of the solution. That is, to really fix it requires either tremendous amounts of work (basically making it unrecognizable with its previous incarnation) or tremendous amounts of effort to work around it.

Problems that are opportunities are solvable without having to reinvent the wheel. They don’t require a redesign or workaround—they’re just an addition. That doesn’t mean it’s less work, but the solution is often more effective and long-lasting.

We saw the problems with VMFS as inherent and fundamental. We saw the problems with vVols as opportunities with much more potential for value and differentiation. Once we laid the foundation by

adding vVols support, it would be an easier path to innovate.

We focused on building a great vVols platform and staying true to our engineering values. We worked directly with VMware to solve design challenges and spread the value proposition that vVols offers. We also worked with third-party vendors to support vVols in general.

At the very end of 2017, in Purity 5.0, FlashArray introduced support for vVols. Over the past three years, adoption has grown steadily. As of November 2020—according to VMware’s latest call-home data—**FlashArray is now the #1 most common product for vVols deployments.**

This is a big deal for us! Over the last few years, we’ve continually invested and committed to vVols not just in our VMware integration team (which has quadrupled in size) but also as a company. At the same time, we’ve focused on partnering with VMware to make vVols a success in all of its products. We leveraged vVols not only as an opportunity to improve our VMware integration, but also to improve our engineering collaboration with VMware. Most important of all, we believe that vVols provides a better joint solution to our customers and a way to make Pure’s world-class data plane more accessible to more people.

## What Lies Ahead

Our investment and commitment is paying off in spades, providing vVols design partnerships, like VMware Cloud Foundation, Cloud Native Storage, and Site Recovery Manager. We’re particularly excited about how vVols can add value to Kubernetes on top of VMware and how we can deliver even more value with Portworx®, the market-leading Kubernetes Data Services Platform that’s now [part of the Pure family](#).

And the momentum is growing. There were monumental shifts in VMware’s investments into vVols this year as vVols become a foundational part of many of VMware’s solutions and products. And we’re honored to be design partners in many of them.

vVols provide an excellent opportunity for storage vendors to add value and differentiate within VMware environments. We’d like to thank VMware for its continued partnership and investment, and most of all, our customers who have adopted vVols on top of Pure Storage!

We look forward to our continued collaboration and success in the world of vVols.