

2023 Database on VMware Predictions



As we start a new year, it's time for what we all predict *will* happen, *to* happen: sharing predictions for the upcoming year!

I'm going to join the fray with a few of my predictions around the VMware ecosystem and what that means for databases and overall customer direction. Objectively, what I'm about to outline aren't really predictions—they're statements of fact. We see this happening now, and they'll only continue to grow in 2023.

Prediction: The right architecture for the right use case will eliminate trade-offs

Containerization does not mean the death of virtual machines (VMs), just as VMs did not mean the death of bare metal. VMs took market share, of course—in fact, a lot of it from bare metal. Sometimes this was even in the face of it being potentially the wrong technical implementation for the use case. Why? Well, the benefit of the simplicity of management consistency can speak volumes. The potential trade-offs between a less-than-ideal technical solution and the benefits of simplicity and organizational streamlining were worth it.

The shift here is realizing that **we don't have to make these trade-offs anymore**. VMware has embraced this, not only through its Tanzu platform but also to bare metal management with Project Monterey. Now you get consistent management experience and the right architecture support for the right use case: containers, bare metal, and VMs. We saw this play out, and continue to see it play out, even in the public cloud, where you can leverage container services, VM services, and even request bare metal.

Pure Storage® has embraced this too, beyond just simply supporting our storage with all three architectures.

From a VMware perspective, we've been working with the VMware Cloud Native Storage team (the team behind the VMware Kubernetes CSI driver) to make a seamless connection between Kubernetes storage requests via persistent volumes and the underlying [VMware](#) datastore layer. This provides a consistent provisioning mechanism to both containers and VMs. Furthermore, our investment into the vSphere Virtual Volume (vVol) program allows for data mobility between containers, VMs, and even bare metal. This is also why we now refer to vVols as "application-granular," VMware Storage-as-Code. It isn't just about the VM these days, and neither is vVols.

VMware's CNS driver leverages Storage Policy Based Management, which is intrinsic to vVols, making these a tight fit with one another. The vVol architecture around data format openness allows for data sets to be easily moved, or more often, instantly copied and shared between any of the three architectures. The [Pure Storage FlashArray vVol implementation](#) specifically focuses on making this characteristic a priority. The combination of this with [Pure Cloud Block Store™](#) even allows these data sets to be quickly and simply replicated to AWS or Azure to extend this architecture over the hybrid cloud.

We cannot fail to mention [Portworx®](#) here, which directly leverages the VMware CNS CSI driver when running on top of VMware-based Kubernetes deployment to leverage that work too. This is an end-to-end story.

Prediction: More Investment in the Organizational Cloud Operating Model

Cloud is a state, not a place. This is not news. Bringing that state to your organization across all of your operational "places," however, *is* new. How to deliver a [cloud operating model](#) to your on-premises environments in addition to your public cloud deployments is the key to making the best of both places. As-a-service and [infrastructure-as-code](#) are inherent in this. But how do you deliver the resources you have in a service-like model?

VMware has been investing in this realm through offerings like VMware in the cloud, VMware Cloud Foundation, and the litany of orchestration offerings such as the Aria (formerly vRealize) suite. Investments into NSX to provide this to the network, and once again, SPBM with vVols, deliver this to the storage architectures. The importance and value of vVols is as VMware Storage-as-Code. vVols makes the storage object (the virtual disk or VM) an API-driven and managed object.

Pure has been meeting VMware in all of those places as well and has grown this intersection even further with the introduction of [Pure Fusion™, a managed service offering for storage-as-code](#), as well as [Portworx Data Services](#) to provide service-driven databases.

The drive to deliver this cross-cloud operating model will only increase with time and the need to provide consistency of this model across clouds will shed more light on the multi-cloud use case.

Prediction: Ransomware protection will be more than 3-2-1

Ransomware, unfortunately, is still a hot topic. It was a frequent conversation at AWS re:Invent 2022, and of course, anywhere else where IT conversations occurred. Protecting your organization against ransomware is both an offensive and defensive battle. How do you prevent and detect, and also restore? The best offense is a good defense as they say, but you certainly need both!

A key consideration around the restoration capability is *not only can you restore, but how quickly?* Just the investment in backup does not solve that problem. Restoring a 31TB database from backup could cost your business (in terms of time-to-restore-induced business downtime) more than the ransom itself. This situation has played out over and over. Yes, we can restore, but from a business perspective, it's *cheaper* to just pay. How can that situation be de-risked? The 3-2-1 backup adage still applies for all of the normal reasons (three copies, two locations, one off-site) but with this additional axis: Time to restore needs to be considered critical.

Databases are the lifeblood of an organization and often the target of these types of attacks for that very reason. How to provide protection, protection that cannot be turned off, and protection that can be restored quickly is key.

Solving this is the essence of [FlashArray™ SafeMode™](#). These snapshots cannot be manually deleted and protection policies cannot be reduced in protection frequency. Using the resulting snapshots, you can restore a VM or VM disk instantly, regardless of size. Even your FlashArray administrators cannot change these configurations. If they could, so could an attacker. *Zero trust indeed.*

The focus on those rules, plus the granularity of vVols, makes critical database restores fast, simple, and effective. Whether you script it, click it, or use our vSphere Plugin that integrates instant VM or disk restore directly into vSphere itself, flexibility in the choice of management is wide.

Addressing these issues head-on is a topic we expect our customers to continue to invest in, and we're working to make that easy on VMware.

Prediction: A Farewell to SCSI

Lastly, adoption of new fabric and data paths will really take hold in 2023. More specifically, a focus on mass adoption of [NVMe-oF](#). NVMe-oF/FC, RoCEv2, and TCP provide the fabric flexibility of choice, while ditching the decades-old SCSI architectures that were simply not built for modern applications—VMs, bare metal, or containers alike.

The roadmap for NVMe-oF in 2023 is quite exciting—the past has been about the fundamentals, which mostly led to *tactical* adoption for the most demanding of application use cases. Now that we're past that, it's time for the true value-add features to be delivered for full ecosystem integration. This is where it gets fun, and that fun is a major focus for VMware and Pure Storage this year. 2023 will be the year of NVMe-oF, no doubt.

This is shaping up to be a fun year!

