

The Changing Landscape of Storage



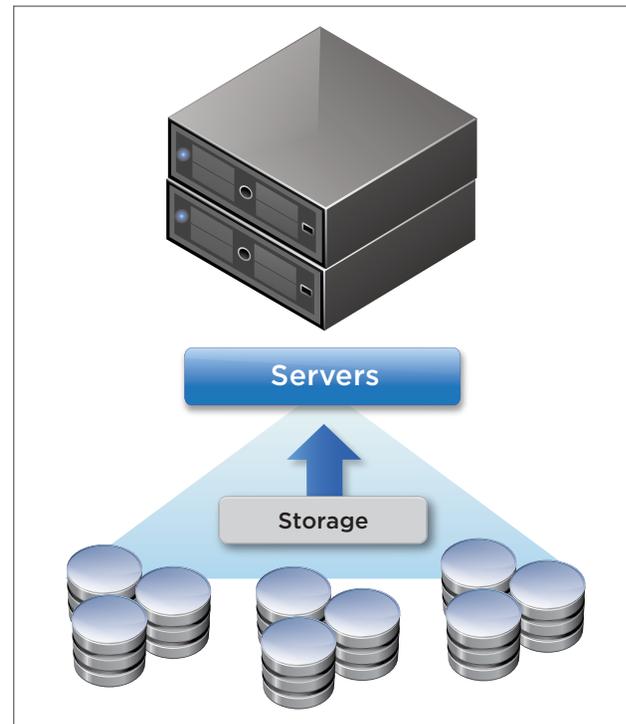
Virtualization Is Reshaping How We Manage Our Data

You may not feel the earth moving, but if you're managing a virtualized infrastructure or serving as a storage administrator be aware: A disruptive, industry-changing shift is in motion.

Much of what we have come to know about data storage is in a state of collapse. Traditional storage structure—characterized by external storage arrays—once loomed as a standalone promontory on the IT landscape. Today that structure is breaking apart and realigning amidst the fabric of servers, software, and hypervisors.

In this brief, we examine this disruption in detail and what it means to you and your business. But in order to do so, we need to consider four components of the realignment:

- Storage technology base
- Storage consumption model
- Storage integration model
- Storage operational model



As storage technology now makes the transition from arrays to industry-standard servers running software, price and performance advantages are emerging. New widgets are available faster, and the benefits are realized with far less upheaval for the installed environment. In fact, some traditional storage vendors are offering array-like packages of storage software running on preconfigured industry-standard servers. The result is lower costs and the simplicity of a singular support model.

The Storage Technology Base

Our familiar storage technology base—traditional external storage arrays built on proprietary hardware and running highly specialized microcode—has held its ground on the IT landscape for the past two decades. In this model, storage array vendors compete for enterprise business based on performance, functionality, support, and price.

But this approach has always posed fundamental pitfalls for IT managers. Because component technology in storage arrays is constantly evolving, vendors face the ongoing, and time-consuming, challenge of qualifying their offerings for the faster, better, or cheaper options. Consequently, customers might find themselves contending with backward compatibility, and whether a new widget can function within the array that's already in place. Or, the potential cost—and delay—required to buy a new array so the business can benefit from the faster-better-cheaper component.

The Storage Consumption Model

Operating under the traditional array model means acquiring storage from one or more vendors. And such storage arrays can be large and comparatively expensive boxes that you need to integrate with your IT environment. It's a process that requires meticulous planning—including vendor negotiations—and bears the risk of a costly aftermath if mistakes are made.

The new model simplifies the acquisition process, and buffers you against the risk described above. Not only can you purchase storage capacity from the server vendor as part of your server acquisition, you can make your acquisition in smaller increments—a few components here, an additional server there. Without the pressure of major storage acquisitions, your IT team is under less strain to plan ahead and enjoys more flexibility to make adjustments as you grow. Furthermore, without the economic overhead associated with large storage arrays—server supply chains can be quite efficient—your acquisition costs should be much less on a per-unit basis.

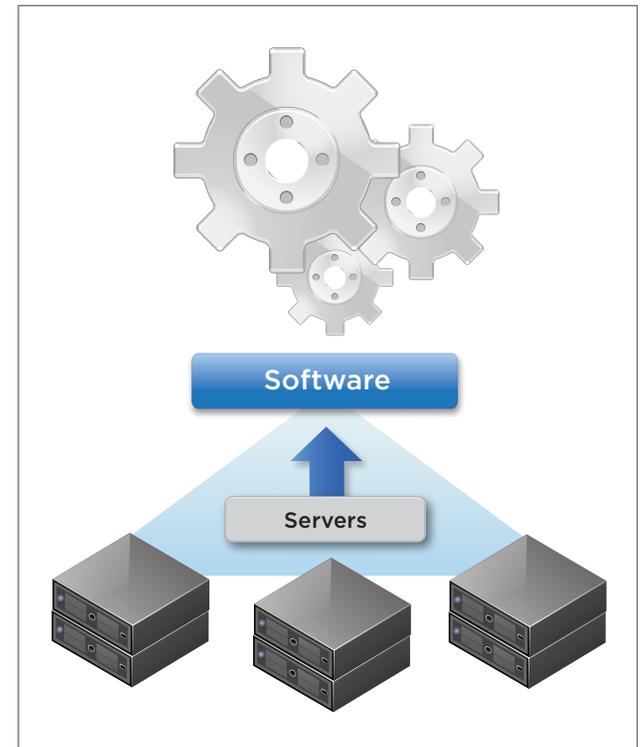
The Storage Integration Model

With external storage arrays standing apart from servers and applications, it's become a necessity to build bridges between them over the 20 years that arrays have held a prominent position. You recognize these bridges in the form of storage networks and protocols, APIs, plug-ins, and management tools. Still, gaps remain. For instance, most arrays don't natively understand applications and application boundaries. As a result, the array won't automatically understand changes in application policy. Put another way, we are operating in a "bottom-up" environment in which storage is acquired as pre-defined buckets, and applications then consume from the buckets as needed.

What would it look like if we flipped the picture? As an application gets provisioned, its requirements are driven via policy down to the supporting storage infrastructure. Resources and service levels aren't carved into buckets until they are actually consumed.

In fact, it's hardly a novel concept. The AS/400 platform that was launched in the late 1980s featured storage concepts integrated throughout. And today storage functionality and data services are on the march back to the servers and applications they serve.

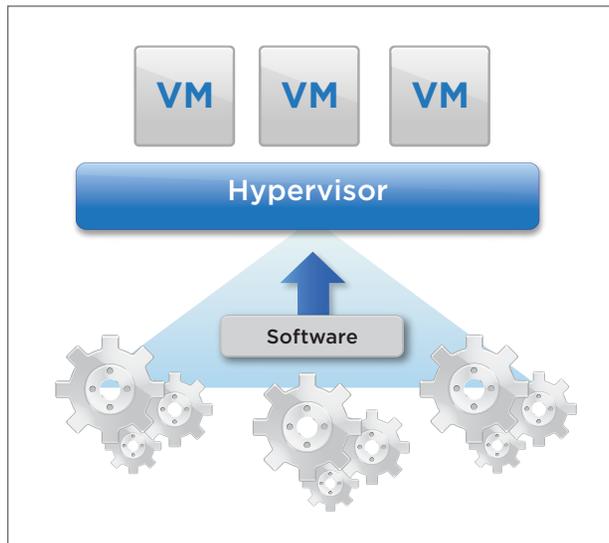
As the march picks up momentum, it's worth noting that the hypervisor is uniquely positioned to absorb this functionality. The hypervisor can see each and every workload, all infrastructure resources, and sits in a convenient point to attach policies that drive application-specific behaviors.



The Storage Operational Model

Operations represents the most significant of the four phases of storage realignment. In most IT settings, storage operation rests in the hands of storage professionals—who adhere to their own workflows, certifications, and perspectives. Similarly, virtualization professionals follow their own path in order to achieve their objectives. Consequently, all the usual IT workflows, such as provisioning, monitoring, troubleshooting, capacity planning, and performance management, move back and forth between the two groups.

But as storage is realigned with servers, software, and ultimately the hypervisor, the need for dedicated storage expertise dissolves, and with it the need for the inefficient workflows between the two groups. In fact, a virtualization management team using VMware vSphere® can fulfill nearly all of its responsibilities without involving or engaging a storage team.



The payoff? Your storage team can cast aside day-to-day demands and provide strategic value to the IT organization by focusing on your most challenging needs. And you experience the simplicity of a storage operational model that's folded into general virtualization—and cloud—workflows, whether manual, scripted, or fully automated. Now you are managing one task from a single control point rather than multiple tasks from multiple control points.

Conclusion

Marc Andreessen once said “software is eating the world.” Within the realm of IT, it seems more salient to say that virtualization is eating infrastructure, from compute to networking, and now storage. In fact, traditional standalone storage is crumbling. Taking its place is a model where storage is collapsing into servers, software, and the hypervisor—becoming part of the service to be consumed. With this shift comes a shift in thinking about data center consolidation and efficiency. Ask yourself if you're in a position to benefit from this changing landscape.

VMware Virtual SAN™, hyper-converged storage for virtual machines, can help you take advantage of the changing landscape of storage. To find more information about VMware Virtual SAN, visit the website at <http://www.vmware.com/products/virtual-san>.

Try VMware Virtual SAN for free today.

Join Us Online



Blog: <http://blogs.vmware.com/virtualblocks/>

Twitter: @VMwareVsan

Facebook: <https://www.facebook.com/vmwarevsan>