

Solution Brief

Can Maxta Eliminate the Storage Array?

Date: November 2013 **Authors:** Mark Peters, Senior Analyst; and Monya Keane, Research Analyst

Abstract: As IT organizations deploy virtual servers and implement software-defined data centers, they are encountering storage-related challenges. To solve those challenges and support the new data center architectures, a startup software vendor called Maxta is delivering enterprise storage via a fresh idea and a highly differentiated technology solution.

Introduction: Server and Storage Constructs No Longer Match

Data centers transformed dramatically over the past decade. But the transformation didn't occur uniformly:

- **Virtualization practically redefined the server side of IT**, enabling asset pooling, hardware resource conservation, and service allocations on demand. Those enhancements, in turn, brought agility, availability, simplicity, and cost savings. The virtual machine was a radically new construct.
- **No similar construct really arose on the storage side**, despite marketing claims—and some progress—to the contrary. Improvements at the storage-media level (i.e., flash/SSD) appeared, but arrays themselves remain based on SAN or NAS architectures optimized for physical, not virtual, environments. And attaching old-fashioned SAN/NAS arrays to virtual servers is causing problems.

In some ways, this contrast in the pace of data center evolution could be described as “Ponzi storage.” In a classic Ponzi scheme, what is disguised as profit is actually just capital investment from newer investors. In Ponzi *storage*, no one does any illegal investing; it's simply that “this year's new storage features” are being continually introduced to sort out the problems that last year's new features supposedly fixed.

In a Virtual Data Center, Storage Complexity Brings Real Pain

Storage challenges in a virtual data center include a lack of simplicity, lack of functionality, and lack of economy. For instance, controlling CapEx/OpEx, maintaining DR preparedness, allocating capacity, and numerous other tasks become more difficult (see Figure 1), and ultimately, managing the whole IT environment grows cumbersome.¹

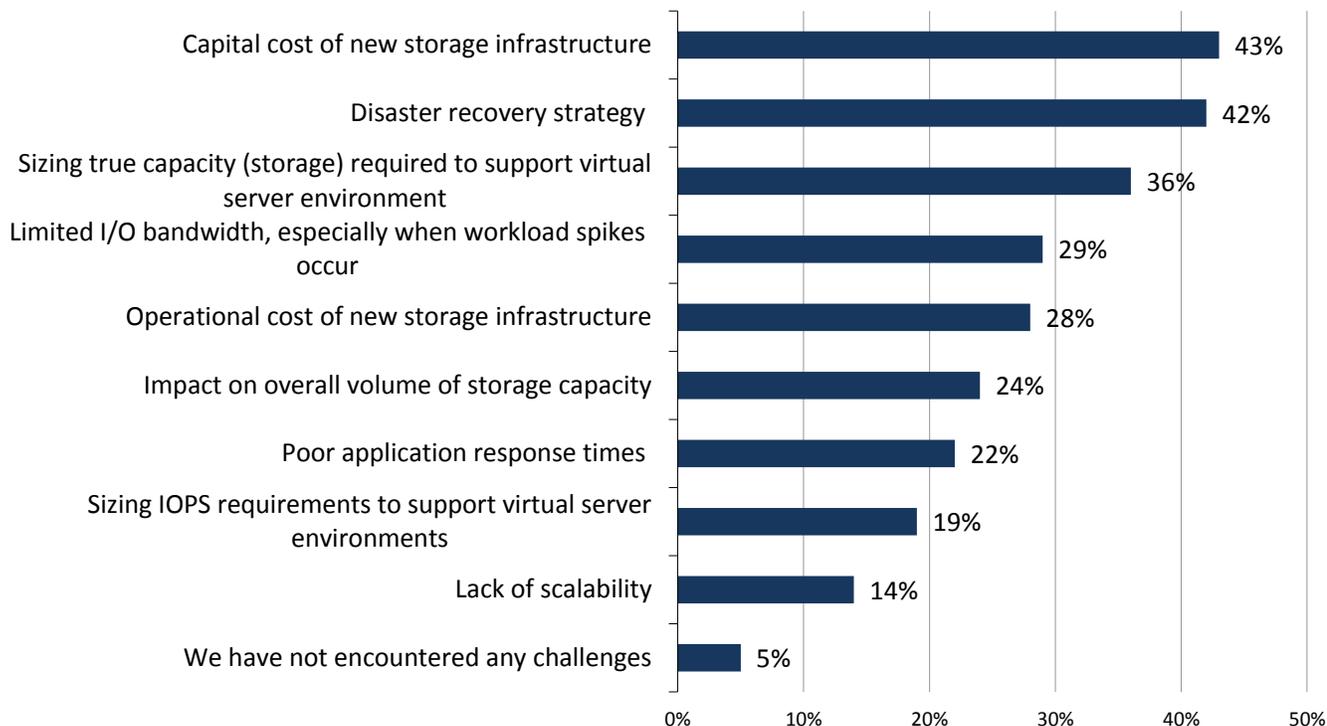
Storage is complex in a virtual data center because storage constructs (volumes, LUNs, file systems) were built to have one-to-one relationships with *other physical constructs* (applications and operating systems on physical servers). Such one-to-one relationships don't exist in the virtual data center because storage doesn't provide a construct that corresponds to a VM. In a data center with virtual servers, one must map multiple VMs to a physical storage volume. This mapping is not trivial because it involves determining the most appropriate volume size and working out which virtual machines go to which volumes—all while taking performance *and* cost into consideration.

It's analogous to visiting a country where you can't speak the language and must refer to a dictionary for every word. If you could only speak fluently, you and everyone you encounter would get more accomplished with much less effort from all sides. But as it stands, a lot of effort is spent creating the message, parts of your message are lost in translation, or your meaning is misunderstood completely. Storage management in a virtual data center is similarly painstaking. Every enterprise storage system has its unique approach for managing storage constructs and its own user interface—its own dictionary for its unique language—leaving IT admins with the complex task of “translating.”

¹ Source: ESG Research Report, [2012 Storage Market Survey](#), November 2012.

Figure 1. Server Virtualization Brings Storage Challenges

From a storage infrastructure perspective, which of the following would you consider to be significant challenges related to your organization's server virtualization usage? (Percent of respondents, N=418, multiple responses accepted)



Source: Enterprise Strategy Group, 2013.

It's Especially Hard to Leverage Advanced Information Management Capabilities

Things become particularly complicated with operations such as deduplication, compression, snapshotting, cloning, and replication. For example, what happens when a VM admin wants to take a snapshot of a virtual machine to protect it? Well, her colleague, the storage admin, can only provide volume-level, LUN-level, or file-level snapshots. So someone will need to match the VM snapshot to a volume snapshot. Storage vendors do try to offer tools to “make it all work,” but those tools are rarely straightforward and invariably require multiple steps to accommodate the gap between virtual servers and physical arrays.

However, enterprise-class storage capabilities are required to get the most from virtual servers:

- If you want the ability to move VMs across servers, then you need the ability to share storage across servers.
- If you want to provision VMs as fast from the storage side as from the server side, then you need efficient VM-level snapshotting and cloning.
- If you want to limit data sprawl, then you need thin provisioning, compression, and deduplication.

To leverage *server* virtualization, you need enterprise-class *storage*. We all know enterprise-class storage arrays are expensive and complex to manage. It raises the question: Does “enterprise-class storage capability” always have to equal “an external storage array”?

Maxta Wants to Do Away With External Storage Subsystems

The gulf between server and storage innovation is at the heart of the problem a “new-ish” software firm called [Maxta](#) is solving by promoting a different way to match the storage side to the compute side in simplicity, availability, agility, and cost.

Founded in 2009, privately held, California-based Maxta was funded by Silicon Valley venture capital firm Andreessen Horowitz. Maxta has been forming partnerships with channel partners to sell the software and package it into IT solutions, enabling VAR and SI partners to sell enterprise-class IT bundles without storage arrays. Maxta's geographic focus is North America, with some paying customers based internationally. Interestingly, all the European deals stemmed from interested parties proactively approaching Maxta after hearing buzz about the solution.

The software itself—Maxta Storage Platform—shipped to select customers in September 2012 and has been used in customer production environments since December 2012. Customers are using the technology in both dev/test and production environments, and in both central and branch locations.

'Going Back to the Drawing Board'—a Technical Overview

Maxta is on a mission to go back to the drawing board completely, "reinventing" storage to make it compatible with a virtualized world. Of course, this is a software company, not a hardware vendor: With this product, capacity allocation, cloning, snapping, replication, dedupe, and compression all happen at the VM level.

The company's executives believe that when you eliminate storage arrays and storage networking, you also eliminate complexity and cost. There is no need to provision storage, decide which VM goes to which volume, or map a volume-level snapshot to a VM-level snapshot. Admins can instead focus on running applications/VMs and on services such as DR and data protection.

Collectively, Maxta describes its approach as enabling true convergence—with applications, server virtualization, and storage running on the same server. But unlike most competitive solutions, this solution runs on any commodity server from any vendor, with any preferred mix of disk and SSD.

Maxta customers report that **snapshot** and **cloning** capabilities are solid. Although ESG has not validated Maxta's technology, it has heard from Maxta that the platform can create thousands of snapshots per VM at sub-second speeds with no performance degradation—an intriguing capability because creating snapshots and clones on distributed systems is usually much more complicated than on traditional centralized storage architectures.

Thin provisioning is the default methodology, and other advanced features include **in-line compression** and **deduplication** to support capacity optimization. In addition to compute/storage convergence, Maxta delivers enterprise-class features (data sharing, high availability, non-disruptive installation) and data services expected from enterprise storage.

Maxta uses replication for **high availability**; two copies of all VM data are always stored on at least two physical servers. If one server goes down, high availability kicks in for all VMs that had been running on that server. Thus, users don't experience any issues if a server goes down unexpectedly or for routine maintenance. Maxta provides data availability inside servers via well-proven mirroring and RAID techniques as well.

Another differentiator is Maxta's **global namespace**. Without it, Maxta would not be able to offer the simplicity it does. A global namespace orchestrated in peer-to-peer fashion and optimizing data placement to get more intra-server bandwidth leverage (rather than inter-server bandwidth) is powerful. Although Maxta uses dispersed storage devices across multiple servers, it all essentially becomes one large pool of storage that any VM on any server can access.

Flash is not mandatory, but, with prices falling, it is highly recommended. Maxta leverages a combination of flash and spinning disk to get "the best of both worlds," leveraging flash's performance and disk's capacity and cost advantages. Maxta is using flash as a read/write back caching mechanism in front of magnetic drives. Moreover, data is never written randomly to the magnetic drives; rather, data written by any VM is done so sequentially to the SSD (to a log). If data isn't being accessed, it is eventually de-staged off SSD and onto high-capacity SATA disk in a sequential manner. Customers *could* deploy an all-flash configuration to get the best possible performance, although Maxta reports no customers yet have opted for that implementation.

Everything runs on commodity servers, but Maxta does offer storage-investment protection by being able to leverage arrays that users already may have on the floor. It is non-disruptive and can be installed gradually over time as applications come online or as data-migration activities occur.

A Maxta Customer's Study

The product appears to be a viable way for companies to move forward with virtualization without having to invest in networked storage. Maxta maximizes CapEx and OpEx savings by leveraging convergence. Only standard servers are needed—not storage arrays or storage networks. And in terms of storage media, Maxta capitalizes on the appealing price/performance attributes of commodity drives and consumer-grade SSDs.

According to Maxta, one customer's three-year TCO study—which examined the expense of purchasing a storage array versus simply buying disk and SSDs, upgrading to a more powerful server, and deploying Maxta software—showed an expected savings of 55% in CapEx alone.

But the bigger savings are in OpEx. Power, cooling, and floor-space costs fall significantly when you don't need arrays. Most important of all, though, is the management-cost reduction. Maxta customers don't need staff with specific understanding of storage, and they save time that was previously needed to do snaps, clones, etc.

In all, the customer cited above anticipates an 80% reduction in TCO over three years.

The Bigger Truth

This is software that manages storage, simplifies IT, and eliminates the need for storage arrays. But make no mistake: It's no "software-defined blah-blah" oversell/stretch. Remember when enterprise-ready flash arrived on the scene? Recall when the cloud "materialized"? Some people dismissed those concepts after hearing IT vendors self-servingly twisting the terms to suit their marketing goals. But now everyone sees the value of flash and cloud.

Likewise, software-defined storage should get to that point, too, as long as the manifest products—such as Maxta—truly do offer value in terms of automation, simplicity, and cost control through the optimized use of inexpensive commodity components.

Maxta's bold yet beguiling idea is to eliminate storage arrays and storage networking—an ambitious quest. The company is all about redefining storage for the virtual data center, and because Maxta believes that upgrading storage arrays and storage networking is obviously turning out *not* to be the way to do it, it decided to start from scratch and eliminate the entire notion of using an external box to run storage independently from the servers running compute tasks and VMs.

Maxta doesn't eliminate hardware entirely, of course; servers and drives are still necessary. But it does essentially eliminate the need for users to deploy or maintain traditional external arrays. As a result, the Maxta approach dramatically simplifies an IT environment—by freeing virtualization administrators to manage virtual machines and applications ... and not storage.