

DATRIUM'S OPTIMIZED PLATFORM FOR VIRTUALIZED IT "Open Convergence" Challenges HyperConvergence

NOVEMBER 2016



The storage market is truly changing for the better with new storage architectures finally breaking the rusty chains long imposed on IT by traditional monolithic arrays. Vast increases in CPU power found in newer generations of servers (and supported by ever faster networks) have now freed key storage functionality to run wherever it can best serve applications. This freedom has led to the rise of all software-defined storage (SDS) solutions that power modular HyperConverged infrastructure (HCI). At the same time, increasingly affordable flash resources have enabled all-flash array options that promise both OPEX simplification and inherent performance gains. Now, we see a further evolution of storage that intelligently converges performance-oriented storage functions on each server while avoiding major problems with HyperConverged "single appliance" adoption.

Given the market demand for better, more efficient storage solutions, especially those capable of large scale, low latency and mixed use, we are seeing a new generation of vendors like Datrium emerge. Datrium studied the key benefits that hyperconvergence previously brought to market including the leverage of server-side flash for cost-effective IO performance, but wanted to avoid the all-in transition and the risky "monoculture" that can result from vendor-specific HCI. Their resulting design runs compute-intensive IO tasks scaled-out on each local application server (similar to parts of SDS), but persists and fully protects data on cost-efficient, persistent shared storage capacity. We have come to refer to this optimizing tiered design approach as "Server Powered Storage" (SPS), indicating that it can take advantage of the best of both shared and server-side resources.

Ultimately this results in an "Open Convergence" approach that helps virtualized IT environments transition off of aging storage arrays in an easier, flexible and more natural adoption path than with a fork-lift HyperConvergence migration. In this report we will briefly review the challenges and benefits of traditional convergence with SANs, the rise of SDS and HCI appliances, and now this newer "open convergence" SPS approach as pioneered by Datrium DVX. In particular, we'll review how Datrium offers benefits ranging from elastic performance, greater efficiency (with independent scaling of performance vs. capacity), VM-centric management, enterprise scalability and mixed workload support while still delivering on enterprise requirements for data resiliency and availability.

DATA CHALLENGES IN VIRTUALIZED ENVIRONMENTS

Virtualized environments present a number of unique challenges for user data. In physical server environments, islands of storage were mapped uniquely to server hosts. While at scale that becomes expensive, isolating resources and requiring a lot of configuration management (all reasons to virtualize servers), this at least provided directly mapped relationships to follow when troubleshooting, scaling capacity, handling IO growth or addressing performance.

However, in the virtual server environment, the layers of virtual abstraction that help pool and share real resources also obfuscate and "mix up" where IO actually originates or flows, making it difficult to understand who is doing what. Worse, the hypervisor platform aggregates IO from different workloads hindering optimization and preventing prioritization. Hypervisors also tend to dynamically

move virtual machines around a cluster to load balance servers. Fundamentally, server virtualization makes it hard to meet application storage requirements with traditional storage approaches.

Current Virtualization Data Management Landscape

Let's briefly review the three current trends in virtualization infrastructure used to ramp up data services to serve demanding and increasingly larger scale clusters:

1. Converged Infrastructure - with hybrid/All-Flash Arrays (AFA)
2. HyperConverged Infrastructure - with Software Defined Storage (SDS)
3. Open Converged Infrastructure - with Server Powered Storage (SPS)

CONVERGED INFRASTRUCTURE - HYBRID AND ALL-FLASH STORAGE ARRAYS (AFA)

We first note that converged infrastructure solutions simply pre-package and rack traditional arrays with traditional virtualization cluster hosts. The traditional SAN provides well-proven and trusted enterprise storage. The primary added value of converged solutions is in a faster time-to-deploy for a new cluster or application. However, ongoing storage challenges and pain points remain the same as in un-converged clusters (despite claims of converged management as these tend to just aggregate dashboards into a single view).

The traditional array provides shared storage from which virtual machines draw for both images and data, either across Fibre Channel or IP network (NAS or iSCSI). While many SAN's in the hands of an experienced storage admin can be highly configurable, they do require specific expertise to administer. Almost every traditional array has by now become effectively hybrid, capable of hosting various amounts of flash, but if the array isn't fully engineered for flash it is not going to be an optimal choice for an expensive flash investment. Hybrid arrays can offer good performance for the portion of IO that receives flash acceleration, but network latencies are far larger than most gains. Worse, it is impossible for a remote SAN to know which IO coming from a virtualized host should be cached/prioritized (in flash)- it all looks the same and is blended together by the time it hits the array.

Some organizations deploy even more costly all-flash arrays, which can guarantee array-side performance for all IO and promise to simplify administration overhead. For a single key workload, a dedicated AFA can deliver great performance. However, we note that virtual clusters mostly host mixed workloads, many of which don't or won't benefit from the expensive cost of persisting all data on all flash array storage. Bottomline - from a financial perspective, SAN flash is always more expensive than server-side flash. And by placing flash remote across a network in the SAN, there is always a relatively large network latency which denigrates the benefit of that array side flash investment.

HYPERCONVERGED INFRASTRUCTURES - SOFTWARE DEFINED STORAGE (SDS)

As faster resources like flash, especially added to servers directly, came down in price, so-called Software Defined Storage (SDS) options proliferated. Because CPU power has continuously grown faster and denser over the years, many traditional arrays came to be actually built on plain servers running custom storage operating systems. The resulting storage "software" often now is packaged as a more cost-effective "software-defined" solution that can be run or converged directly on servers (although we note most IT shops prefer buying ready-to-run solutions, not software requiring on-site integration).

In most cases software-defined storage runs within virtual machines or containers such that storage services can be hosted on the same servers as compute workloads (e.g. VMware VSAN). An IO hungry application accessing local storage services can get excellent IO service (i.e. no network latency), but capacity planning and performance tuning in these co-hosted infrastructures can be exceedingly difficult. Acceptable solutions must provide tremendous insight or complex QoS facilities that can dynam-

ically shift IO acceleration with workloads as they might move across a cluster (eg. to keep data access local). Additionally, there is often a huge increase in East-West traffic between servers.

Software Defined Storage enabled a new kind of HyperConverged Infrastructure (HCI). Hyperconvergence vendors produce modular appliances in which a hypervisor (or container management), networking and (software-defined) storage all are pre-integrated to run within the same server. Because of vendor-specific storage, network, and compute integration, HCI solutions can offer uniquely optimized IO paths with plug-and-play scalability for certain types of workloads (e.g. VDI).

For highly virtualized IT shops, HCI simplifies many infrastructure admin responsibilities. But HCI presents new challenges too, not least among them is that migration to HCI requires a complete forklift turnover of all infrastructure. Converting all of your IT infrastructure to a unique vendor appliance creates a “full stack” single vendor lock-in issue (and increased risk due to lowered infrastructure “diversity”).

As server-side flash is cheaper than other flash deployment options, and servers themselves are commodity resources, HCI does help optimize the total return on infrastructure CAPEX – especially as compared to traditional silo’d server and SAN architectures. But because of the locked-down vendor appliance modularity, it can be difficult to scale storage independently from compute when needed (or even just storage performance from storage capacity). Obviously, pre-configured HCI vendor SKU’s also preclude using existing hardware or taking advantage of blade-type solutions.

With HCI, every node is also a storage node which at scale can have big impacts on software licensing (e.g. if you need to add nodes just for capacity, you will also pay for compute licenses), overbearing “East-West” network traffic, and in some cases unacceptable data availability risks (e.g. when servers lock/crash/reboot for any reason, an HCI replication/rebuild can be a highly vulnerable window).

OPEN CONVERGED INFRASTRUCTURE - SERVER POWERED STORAGE (SPS)

When it comes to performance, IO still may need to transit a network incurring a latency penalty. To help, there are several third party vendors of IO caching that can be layered in the IO path – integrated with the server or hypervisor driver stack or even placed in the network. These caching solutions take advantage of server memory or flash to help accelerate IO. However, layering in yet another vendor and product into the IO path incurs additional cost, and also complicates the end-to-end IO visibility. Multiple layers of caches (vm, hypervisor, server, network, storage) can disguise a multitude of ultimately degrading performance issues.

Ideally, end-to-end IO, from within each local server to shared capacity, should all fall into a single converged storage solution – one that is focused on providing the best IO service by distributing and coordinating storage functionality where it best serves the IO consuming applications. It should also optimize IT’s governance, cost, and data protection requirements. Some HCI solutions might claim this in total, but only by converging everything into a single vendor appliance. But what if you want a easier solution capable of simply replace aging arrays in your existing virtualized environments – especially enabling scalability in multiple di-

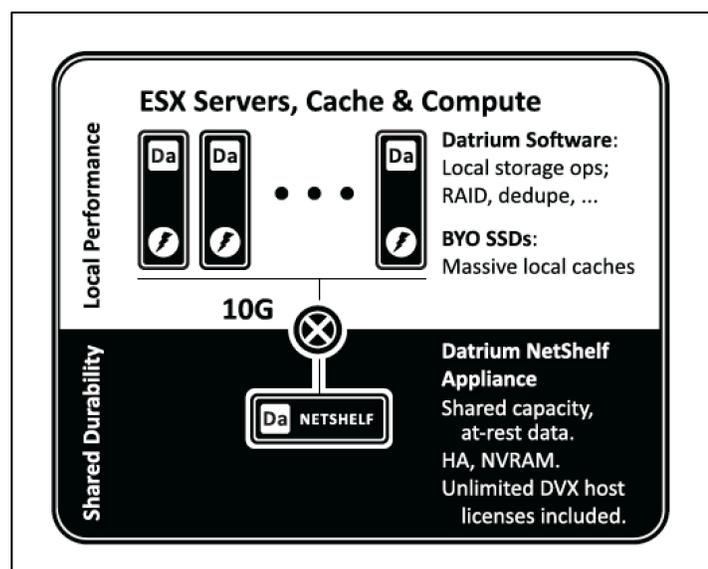


Figure 1 - Datrium DVX Architecture

reactions at different times and delivering extremely low latency while still supporting a complex mix of diverse workloads?

This is where we'd look to a Server Powered Storage (SPS) design. For example, Datrium DVX still protects data with cost-efficient shared data servers on the back-end for enterprise quality data protection, yet all the compute-intensive, performance-impacting functionality is "pushed" up into each server to provide local, accelerated IO. As [Datrium's design](#) leverages each application server instead of requiring dedicated storage controllers, the cost of Datrium compared to traditional arrays is quite favorable, and the performance is even better than (and as scalable as) a 3rd party cache layered over a remote SAN.

In the resulting Datrium "open converged" infrastructure stack, all IO is deduped and compressed (and locally served) server-side to optimize storage resources and IO performance, while management of storage is fully VM-centric (no LUN's to manage). In this distributed, open and unlocked architecture, performance scales with each server added to naturally scale storage performance with application growth.

Datrium DVX makes great leverage for a given flash investment by using any "bring-your-own" SSDs, far cheaper to add than array-side flash (and can be added to specific servers as needed/desired). In fact, most vm's and workloads won't ever read from the shared capacity on the network – it is write-optimized persistent data protection and can be filled with cost-effective high-capacity drives.

TANEJA GROUP OPINION

As just one of IT's major concerns, all data bits must be persisted and fully managed and protected somewhere at the end of the day. Traditional arrays, converged or not, just don't perform well in highly virtualized environments, and using SDS (powering HCI solutions) to farm all that critical data across fungible compute servers invokes some serious data protection challenges. It just makes sense to look for a solution that leverages the best aspects of both enterprise arrays (for data protection) and software/hyperconverged solutions (that localize data services for performance).

At the big picture level, Server Powered Storage can be seen as similar (although more cost-effective and performant) to a multi-vendor solution in which IT layers server-side IO acceleration functionality from one vendor over legacy or existing SANs from another vendor. But now we are seeing a convergence (yes, this is an overused word these days, but accurate here) of those IO path layers into a single vendor product. Of course, a single vendor solution that fully integrates distributed capabilities in one deployable solution will perform better and be naturally easier to manage and support (and likely cheaper).

There is no point in writing storage RFP's today that get tangled up in terms like SDS or HCI. Ultimately the right answer for any scenario is to do what is best for applications and application owners while meeting IT responsibilities. For existing virtualization environments, new approaches like Server Powered Storage and Open Convergence offer considerable benefit in terms of performance and cost (both OPEX and CAPEX). We highly recommend that before one invests in expensive all-flash arrays, or takes on a full migration to HCI, that an Open Convergence option like [Datrium DVX](#) be considered as a potentially simpler, more cost-effective, and immediately rewarding solution.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned here-in are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.