

THE **GORILLA GUIDE** To...[®] EXPRESS EDITION

Top Ten Reasons

IT Leaders Are
Embracing a
New Generation
of Convergence



Inside this book:

- Discover a path to zero-click simplicity
- Find an infrastructure solution with breakthrough economics
- Learn how open convergence enables fully elastic resource scalability

**TAKE A QUICK WALK
THROUGH THE IT JUNGLE!**

Compliments of

Datrium™

THE GORILLA GUIDE TO...®

Top Ten Reasons IT Leaders Are Embracing a New Generation of Convergence

EXPRESS EDITION

James Green

ActualTech Media

Compliments of



© 2017 ActualTech Media, All Rights Reserved.

No portion of this book may be reproduced or used in any manner without the express written permission of the publisher except for the use of brief quotations.

Printed in the United States of America

ActualTech Media

Okatie Village Ste 103-157

Bluffton, SC 29909

www.actualtechmedia.com

Table of Contents

Introduction	4
Converged and Hyperconverged Infrastructure	5
Room for Improvement	6
Zero-Click Simplicity	8
1. Turnkey Purchasing Model	9
2. Configuration-Free Management	10
3. No LUNs to Attend To	12
4. Single UI for VM-Focused Management	14
Breakthrough Economics	15
5. Server-Based Flash is Fast and Affordable	16
6. Inherent Choice	19
7. Cost-optimized Data Node Offers Storage Efficiency	21
Elastic Scalability	22
8. No Controller Bottlenecks	22
9. Scale IOPS and Capacity Independently	25
10. Designed for Massive Scale	27
Summary	29

Introduction

As computing power in the data center has continued to increase at a relatively steady pace, data center infrastructure has become increasingly complex. In fact, in many cases, it has become complex to the point of being detrimental to the goal of running a fluid and nimble IT operation. In the past few years, forward-thinking enterprise IT vendors have begun to address this growing data center complexity problem by offering solutions that streamline the infrastructure acquisition and deployment process as well as the ongoing management of IT systems.

Many companies are seeking to modernize their data center environments. Modernization is a race to *simplicity*, a meta-trend that is impacting the entire IT operation. Customers who do it first and best will win in their own markets, turning IT into a frictionless advantage rather than the complex anchor that it is for many.

Converged and Hyperconverged Infrastructure

With regard to procurement and support, the industry has seen the rise of “converged infrastructure,” which is an integrated hardware stack that has been pre-validated, configured prior to delivery, and is supported by one organization despite leveraging hardware and software from multiple vendors. Using converged infrastructure makes it much easier to procure and deploy systems at rack scale. One drawback to converged infrastructure, however, is that the units are described in racks and rows; the design is not well-suited to small, granular scaling.

An iteration on the philosophy of converged infrastructure is a paradigm known as “hyperconverged infrastructure,” in which the design goes beyond merely validating and pre-building multiple products and actually combines the infrastructure components of CPU, memory, storage, data services, and sometimes networking, into much smaller units of single server units. With this design, the system is intended to scale out

effortlessly, and in only the amounts needed at any given time.

Room for Improvement

Hyperconvergence has shortcomings as well:

- Compute capacity and storage capacity (and related software licenses) inherently scale together in such systems, even when one isn't necessarily needed
- Persistent data is stored on every server and common configurations are limited to 2 or 3 simultaneous server failures before users experience a data outage. As a result, HCI users tend toward smaller clusters in the 8-16 node range, according to Gartner.
- Many organizations who committed to blade servers or quad socket servers, or are simply just a year into a server purchase are out of luck when it comes to hyperconvergence and have had to choose to either ignore the hyperconvergence trend or to buy net-new gear and dispose of their prior investments

- The east-west traffic between hyperconverged nodes as they replicate storage over the network can, and frequently does, require network re-design to support the increased demand and combat noisy neighbor issues that can arise as the cluster grows.
- Medium to large database and data warehousing workloads can underperform in hyperconverged configurations.

To address these concerns and more, Datrium has iterated yet again on the maturing idea of what data center convergence looks like and offers an “Open Convergence” model with the Datrium DVX product line. Following are the top ten reasons why you may be interested in converging compute, primary and secondary storage with Open Convergence! They’re grouped into sections that illustrate some of the higher-level pillars of Open Convergence: it’s easy to deploy and manage; it’s cost-effective; and it’s elastic and massively scalable.










	Converged	Hyperconverged	Open Convergence
VM Processing			
I/O Processing CPU & Flash			
Durable Capacity			

Figure 1. The evolution of converged infrastructures

Zero-Click Simplicity

Many CIOs are currently preoccupied with one agenda above all else: systematically searching out and destroying every contributor to complexity in their data center. In their effort to protect, strengthen and even accelerate bottom line

performance, executives know that eliminating operational inefficiencies can go a long way.

Datrium knows that too, and, to that end, they've designed the Datrium DVX system to be radically simple, starting with the procurement process.

1. Turnkey Purchasing Model

Datrium DVX combines 1-10 DVX Data Nodes for primary and secondary storage capacity and from 1-128 DVX Compute Nodes for CPU, memory, and local flash. The result is a turnkey, purpose-built solution that includes end-to-end support from Datrium. This design and procurement method combines the upsides of converged and hyperconverged infrastructure with the added benefit of having near infinite flexibility over the configuration as it pertains to CPU, memory, local flash, and persistent storage capacity. A system can be configured to be as small as a single Compute Node and a single Data Node and scales up from there.

2. Configuration-Free Management

Because key data services such as erasure coding, compression, and global deduplication are always on, there is never configuration of hot/cold storage tiers. There's also no deliberation about whether to turn any of these services on for a given application or not. They're just always on!

Most traditional storage systems achieve greater bandwidth and availability by leveraging Link Aggregation Control Protocol (LACP) which effectively aggregates multiple network links. Unfortunately, configuring (and later debugging) LACP can be hard. Since Datrium DVX controls the entire infrastructure stack, traffic decisions affecting throughput and availability can be made higher in the stack than at the network link layer where LACP operates. End-to-end adaptive pathing provided by Datrium DVX simplifies networking configuration and takes care of the hard stuff in software. Not only that, but the network monitoring that the software is doing can even make it easier to identify and resolve a network issue when one does occur.

Open Convergence

- Configuration-Free
- One Rackscale Pod
- Zero Knobs

Hyperconvergence

Assessments, Configuration
Per Cluster

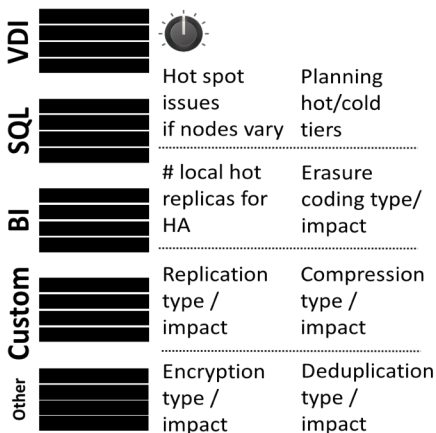


Figure 2. Datrium DVX has no knobs to turn because everything is on and configured by default

3. No LUNs to Attend To

Just about anyone who has managed an enterprise storage system has felt the pain of configuring, managing, and troubleshooting LUNs. Managing LUNs, especially at scale, is a giant pain.

The problem with LUNs is that they're totally abstract and they're an irrelevant construct when it comes to what's actually being stored. In the case of a private cloud environment, an administrator cares about the virtual machines and vDisks, containers, and unstructured data being stored on the system. The administrator should not also be required to figure out which LUN to use to store that data.

Another shortcoming of LUNs in the modern data center is that they're too large a unit for efficient management. It used to be the case that it was acceptable to snapshot, replicate, and recover an entire LUN. But administrators only accepted that reality because they had no better options. How woefully inefficient is that strategy if a single file or virtual machine needs to be recovered? Restoring a 2 TB volume to recover a 150 MB vDisk isn't a

shining example of efficiency, yet that's what LUNs often force administrators to do.

What is a LUN?



A Logical Unit Number (LUN) is technically a number used to identify a logical volume on an array. But colloquially, IT practitioners often use the term "LUN" to refer to the logical volume itself.

For security and performance reasons, LUNs are commonly presented only to very specific hosts. Herein lie some of the frustrations that have come to be associated with LUNs. Remembering and controlling which hosts have (and should have) access to which LUNs is a tall order in the modern data center which is constantly in flux.

For this reason, many modern storage platforms choose to do away with the notion of LUNs in favor of a more suitable unit of management in today's data center: objects, virtual machines, or files.

Datrium's DVX platform was designed to understand and work directly with stored objects as opposed to LUNs. With the built-in secondary storage approach, Datrium DVX enables rich data management at granularities like VMs, virtual disks, and container persistent volumes. In addition, even though there are now 1000s or 10,000s of these objects to manage directly (as opposed to LUNs), the Datrium DVX platform provides simple and scalable ways to manage policies for all of them. You can set policies for services such as snapshots and replication on these objects to map precisely to your business needs, and instantly recover them at will.

4. Single UI for VM-Focused Management

Yet another problem that legacy storage systems presented administrators was the addition of more management interfaces from which to configure and monitor their overall infrastructure. The breadth of possibilities for where to find any given piece of infrastructure information has become a major operational challenge. Administrators may be accustomed to using one console for VM

management, another for storage management, and yet another for replication and backup/restore management, among others. Remember that complexity thing we talked about?

Thankfully, Datrium DVX has accounted for this, and the entire system is configured, managed, and monitored through a single interface. In the same place that administrators manage their virtual machines and containers, they can also manage replication, archiving, restores, and more.

Breakthrough Economics

Storage is notoriously expensive. Purchasing compute and storage together a la Converged and Hyperconverged Infrastructure can be a lot to stomach. Although making major investments can be scary at times, it's significantly easier when a few things are present:

- **Bang for your buck.** Most organizations don't like shopping in the bargain bin if they can help it, but it's important that the purchase feels like a good value. When it seems that the purchase will deliver a lot of

benefits relative to the size of the capital outlay, buyers are happy.

- **Choice.** No one likes being told what to do — especially not when it also involves dropping a serious chunk of change. A big buy feels much better when the buyer has control over exactly what they're getting and how much they're spending.
- **Efficiency.** If the impression at buying time isn't that the solution will save money in the long run, it's bound to be a harder decision. But if a purchase looks to improve future financial prospects, it's an easier decision to make.

As you can imagine, Datrium has considered these items and more when crafting the latest flavor of convergence. The economic implications of Datrium DVX may surprise you — in a good way!

5. Server-Based Flash is Fast and Affordable

If you've ever received a quote for – or purchased – an all-flash array, you know all too well the price tag

that's attached. Yes, a storage array full of the latest generation of flash media is going to blow the doors off of a performance test; but it's going to do something similar to your wallet. There's great news for flash buyers in the convergence model: server-based flash comes at a fraction of the cost of array-based flash. Your money goes further when you're packing servers full of flash, especially when used to store data that has been compressed and deduplicated.

Datrium DVX and Little's Law



Little's Law is a theorem by John Little which states that the long-term average number L of customers in a stationary system is equal to the long-term average effective arrival rate λ multiplied by the average time W that a customer spends in the system.¹ Expressed algebraically the law is:

$$L = \lambda W$$

Little's Law can be applied to storage I/O, and it would work out this way:

$$(\text{Outstanding I/Os}) \div (\text{response time}) = \text{IOPS}$$

This law has a direct impact on storage architecture because as more and more storage operations – virtual machines reading and writing, perhaps – are requested of a system with a fixed number of IOPS, the response time will inevitably increase. Eventually, with enough virtual machines, the response time will grow to an unacceptable level and cause poor application performance or demand that you spend money to increase IOPS capacity.

Because the Datrium DVX model makes use of local flash, I/O requests to the shared storage are limited. Local flash handles all read I/O for the virtual machines and only new write IO, data that is yet to be persisted, is then written to mirrored NVRAM in the back-end Data Node. This model certainly gets the buyer more bang for their buck when it comes to storage!

--

¹ *Little's law*. (2017). *En.wikipedia.org*. Retrieved 3 August 2017, from https://en.wikipedia.org/wiki/Little%27s_law

In the bang for your buck department, the story just gets better. It turns out that due to physics and architectural implications, server-based flash is not only cheaper, but in most cases, it's also faster. That's partially due to its physical proximity to the workload and partially due to architectural implications like the fact that the I/O operation isn't competing for service with other workloads in an I/O queue on an array controller.

6. Inherent Choice

The Datrium folks aren't about telling you what to do. Every buying situation is unique and may require slightly different configurations. Unfortunately, in some buying situations, your options are limited. Some converged and hyperconverged offerings only support certain hardware configurations, and very few support a bring-your-own hardware model. This is often a challenge for organizations with a sizable prior investment in server platforms — especially blade and quad-socket servers.

On the contrary, with Datrium DVX, you can use any of a wide variety of platforms for your Compute

Nodes. This affords organizations the opportunity to continue capitalizing on prior investments as well as the option of leveraging existing supplier relationships to get the best price on their project. It also allows them to stick with what their staff is comfortable with — a very important implication as retraining staff can be costly and time-consuming.

Hardware isn't the only consideration when spec'ing out a new deployment. Organizations are often limited on which hypervisor they can use based on the hypervisor supported by the integrated solution they're hoping to use. In many cases, the solution supports one hypervisor. As VMware vSphere has been the market leader for some time, vSphere is frequently the only supported hypervisor. vSphere has matured over time into a very solid platform, but it isn't for everyone. It's not the cheapest on the market, and some organizations rely heavily on other options from companies such as Microsoft and Red Hat.

Datrium DVX is a choose-your-own-hypervisor solution which currently supports VMware

vSphere, Red Hat KVM, and even a Bare Metal configuration for Docker containers all within a single system. This is a sufficiently broad offering to allow organizations some choice when building their unique Datrium DVX solution.

7. Cost-optimized Data Node Offers Storage Efficiency

Because the Datrium DVX platform separates host-local flash for active data performance from shared storage for persistence, the system can take advantage of the efficiencies of a secondary storage model like a traditional disk-based backup array while still benefitting from the local data services processing and decoupled IO performance scaling that hyperconvergence offers.

Datrium DVX systems are generally sized such that all active VM data (hot and cold) fit in the host-side flash while all copy data (primary and archive) is the only thing regularly stored only on the Data Nodes. That said, however, the Data Node architecture is robust and still performs well thanks to full-feature array-like qualities such as mirrored NVRAM and dual 10Gb network ports on each controller.

Always-on erasure coding for efficient double-fault tolerance, compression, as well as global deduplication across all compute nodes in the system make Datrium DVX both a high-performing primary storage system and at the same time a very cost-effective secondary storage system, capable of economically storing your coldest data. The system currently scales up to ten Data Nodes per DVX system which allows for roughly 1 PB of capacity.

Elastic Scalability

Flexibility in the data center is becoming increasingly important. It's more common than ever before to suddenly need to grow compute capacity, storage performance, or storage capacity. And without a doubt, you'll be expected to do it faster than greased lightning. Thankfully, Datrium has you covered. The Datrium DVX system scales linearly, wisely, and easily.

8. No Controller Bottlenecks

One of the most frustrating storage conundrums of the last decade has been the situation many administrators encounter in which they have plenty of remaining storage capacity, but the controllers in

their standalone storage array have reached their limit. In fact, this has become the norm with the evolution from disk-based media to flash technologies. There are two ways this situation might have been resolved – neither enjoyable.

The first common solution was just to buy a bigger and badder storage system with controllers that could support the load (*and* the load expected over the next couple of years, of course). Once the new system was procured and deployed, the administrator would begin the arduous task of migrating all of the data from the old system to the new one.

Slightly less miserable—but still ugly—there was a second option, but only if you were lucky enough to have a system that allowed this. Some systems allowed for a controller upgrade. This may have involved a complete outage as the storage system was shut down to perform the controller swap. If you were really lucky, perhaps one controller could be upgraded, active connections could be failed over to the new controller, and then the second controller could be upgraded, all without an outage

assuming applications could survive on half the available performance. But the risks involved in undertaking this upgrade are apparent.

DVX: Linear Performance Scaling

- 140K IOPS / compute node
- 2-socket; 22 core/socket, 2.2 GHz
- 4K random reads
- Latency always stays low

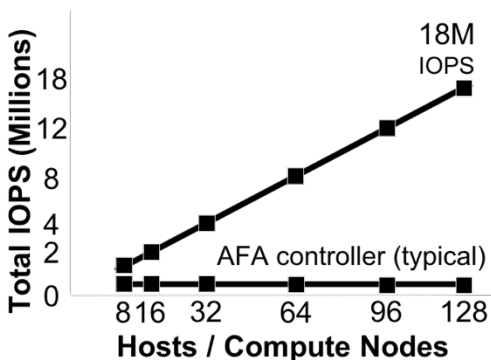


Figure 3. Linear scaling of IO performance capacity in the Datrium DVX system

The Datrium DVX system performs I/O processing on the Compute Node. What that means is that with each additional compute node, the equivalent of “controller capacity” increases. Essentially, IO performance scales linearly all the way to the upper limit of the Datrium DVX configuration maximum, which is currently 128 Compute Nodes. For perspective, at roughly 140,000 4K random read IOPS per compute node, a Datrium DVX cluster at its maximum configuration is capable of 18 million IOPS. And because each node has its very own controller (remember Little’s Law?) the latency stays low.

9. Scale IOPS and Capacity Independently

Now, one of the limited advantages of a traditional shared storage array model is that performance and capacity can be scaled independently of each other. Add bigger controllers in the upgrade described above, or faster disks for more performance; add more or bigger disks for more capacity. With the move to hyperconverged infrastructure, a bit of this flexibility was lost. As was just described, each time a new node is added, IO performance increases.

Unfortunately, in a hyperconverged infrastructure setting, storage capacity is required to expand as well because the architecture demands it.

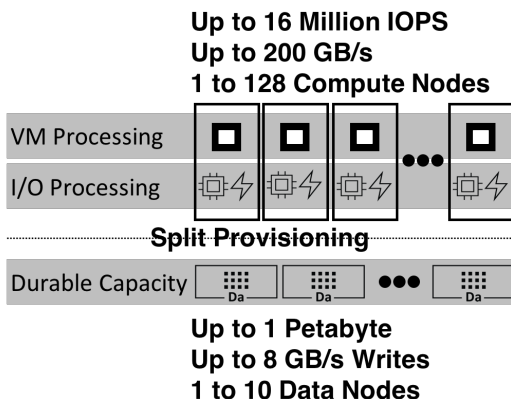


Figure 4. *IO performance and storage capacity scale independently*

The beauty of the decoupling of performance and capacity in the Datrium DVX architecture is that it allows the linear scaling of IO performance that hyperconvergence affords while simultaneously allowing the justified expansion of storage capacity and nothing more. You will not be forced to deploy

more Data Nodes to add a few Compute Nodes. Conversely, if you run out of storage capacity but have plenty of compute remaining, you can simply add another Data Node.

10. Designed for Massive Scale

Because persistent data is stored on the Data Node(s), Compute Nodes are stateless. In the days of boot-from-SAN blade servers, life was good when there was a problem with a server. You'd just yank it out, put in a new one, and get on with your day. Although there were significant benefits gained from hyperconvergence as a model, one of the things that was relinquished was the flexibility that comes with stateless hosts. Because data is now persisted within each compute node, taking a host out of the cluster for maintenance is a carefully executed procedure with some risk involved. Further, losing more than one host in a cluster at one time could absolutely mean that data is lost and certainly unavailable. In a Datrium DVX system, multiple hosts can fail at the same time, and although there will obviously be an impact to running workload capacity, there is no risk to the persistence or accessibility of data.

Because Compute Nodes are isolated from each other (in the sense that they don't replicate data amongst the cluster as hyperconverged nodes do), scaling larger does not come with the performance overhead and network redesign that is seen in many hyperconvergence implementations to support the additional east-west traffic. A network design that has supported a traditional server-SAN architecture will probably serve a Datrium DVX deployment well.

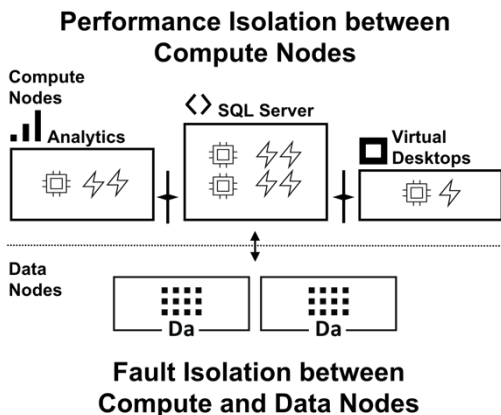


Figure 5. *Compute Nodes are isolated from one another from a data perspective; persistent data is only written to the Data Node(s)*

Summary

You've learned about the various ways that the Open Converged model of Datrium DVX is it's easy to deploy and manage, cost-effective, and massively scalable. To recap, here are the top ten reasons why you should consider Open Convergence:

1. Datrium DVX Rackscale, a turnkey purchasing model, makes it simple to spec and acquire
2. Configuration-free deployment and management removes the nerd knobs so you can focus on what's really important
3. No LUNs! Storage administrators rejoice!
4. Single user interface for VM-focused management, protection and mobility
5. Server-based flash is faster and cheaper than array-based flash
6. The Datrium DVX system gives you all kinds of choices from a bring-your-own

hardware model to a choice of hypervisor (or even bare metal)

7. The cost-optimized Data Node provides secondary storage as well as it does primary storage
8. There are no controller bottlenecks, and I/O performance scales linearly
9. I/O performance and storage capacity scale independent of one another
10. The system is designed for massive scale; a fully populated system can deliver roughly 18 million IOPS and 1 PB of storage capacity

If you remain unconvinced by these ten reasons, you can find even more reasons why Open Convergence could be a great fit for you on the Datrium website, <http://www.datrium.com/>.