Modernizing and Optimizing the Data Center

A Hyperconverged Infrastructure Brief for Business Decision Makers

By Scott D. Lowe and David M. Davis
Co-Founders, ActualTech Media

September 2014

HYPERCONVERGED.ORG

BY ACTUAL/TECH MEDIA, LLC

Introduction: The Business and IT Dilemma

Have you ever had a situation in which your Information Technology (IT) department was unable to meet a need due to technology limitations in the company's data center? Perhaps there was too little data storage space or not enough capacity in the system to meet the needs of a new initiative. The result was an unexpected expenditure required in order to patch the deficiency in the current environment.

For too many businesses, IT, particularly the data center, remains an expensive part of organizations' operating and capital budgets. For example, as a business grows and seeks new revenue opportunities, dated IT systems become a hindrance rather than an enabler, forcing the organization to slow down and adapt to the situation rather than accelerating business.

There are myriad technical reasons that data centers have been constructed the way they have. IT architecture planners, mindful that the business will grow, have attempted to build systems that can adapt, but, like so many, these architects can't always predict that future and may run into situations in which their systems do not provide the level of flexibility and expandability that the business needs.

To be fair, building a data center has traditionally been an extremely complex undertaking, where experts were needed to help carefully plan a variety of different vendor and technology resources in order to build a system that meets requirements. At the same time, many businesses are rightfully looking for more consistency and predictability in both IT budgets and operations, making it more difficult for the IT department to react to an unexpected event (even if that event is a positive one, such as a significant new source of revenue).

The Challenges

Here's a look at the high-level serious challenges being that many businesses face today.

Data Center Complexity

As a business decision maker, you may know that the data center, which could even be simply the server room for smaller organizations, is a complex place filled with a rainbow of cables, a plethora of blinking lights, and dozens of fans whirring to keep equipment operating at just the right temperature. Frankly, that's probably all you really want to know about the data center.

However, as the decision maker, you should be aware that there is a lot more going on within all of those wires, lights, and fans.

The underlying equipment is the brain, heart, and circulatory system of your business. Unfortunately, in many cases, these "organs" are so complex to maintain that each requires a team of one or more specialists to maintain it. Every time your business adds a new service in the data center, it's just like adding a new organ to someone's body and hiring additional doctors to maintain it. As you add more and more people to the team and more and more disparate equipment to the data center, the level of complexity increases.

Decreased Flexibility

As things get more complex in the data center, the ability for IT to be flexible to business needs decreases. Data centers are often described, for example, as a "house of cards"—ready to topple if even the slightest change is made. This is one of the primary reasons that IT pros get a reputation for not being able to handle change; more than likely, they're worried that a change will negatively impact the business.

This results in an inflexible data center that eschews change. Unfortunately, as you know, change is inevitable, and it is the only constant in the business world.

Budget Cycles and Optimization

As you consider your organization's data center spending history, it probably includes a lot of peaks and valleys, especially when you consider both operational and capital costs in aggregate. Worse, it's more than likely that you're not managing your data center spending as efficiently as possible.

Here's an analogy: think of your data center as a bus. (In a way, it really is the vehicle for your business' data and success, so this analogy is not too far off.). As it stands right now, your IT department buys separate vehicle components and assembles them all to build something roadworthy. However, they buy an engine (controller) and body big enough to be a bus (this equates to the individual servers that power the business) before knowing how many seats (disks) will be needed for the passengers (applications). After that, they buy seats for the vehicle as the organization grows and needs more places for applications to sit. The problem is, your company may never actually fill all of that space on the bus if growth is not as expected. Or, your company may outgrow that bus and need to add a trailer behind it with more seats for all your data The question then becomes this, "Can the engine that your IT department crafted, handle the weight of the additional passengers (data storage and data)?"

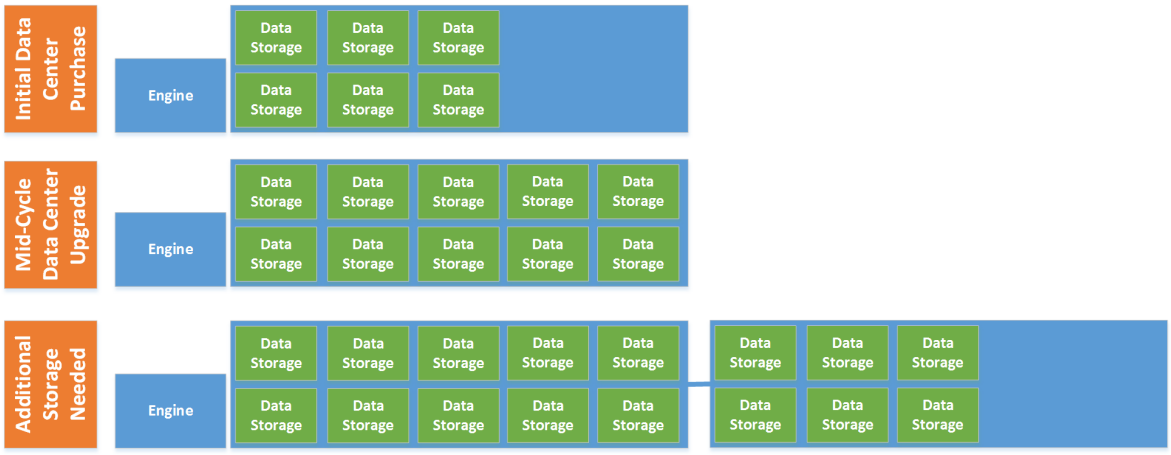

Figure 1: The initial data center purchase (top) has more overall capacity than is needed; during the life of the data center, the organization needed to add storage (middle); finally, additional storage was needed, requiring the deployment of a trailer, but can the engine still pull the whole load?

Bringing this back to the question of money, this situation means that you're spending money up front on capacity that you *hope* to use, and many organizations end up never using the full capacity. Then, if you do reach full capacity, the IT department has to spend more money on the individual components to upgrade the engine of your business "bus" in a very manual way.

Perfectly planning data center systems around budget cycles has traditionally been a tough thing to do.

Increased Cost of Ownership

This legacy data center construct forces expensive IT specialists to spend significant time on planning, building, and maintaining the systems that run the business. Further, the need to make such large up-front purchases in the hopes that the systems will last the duration of the budget cycle, results in inefficient spending as infrastructure just sits there idle awaiting the day when the business grows large enough to use it. Both of these scenarios increase the total cost of ownership for your data center while, at the same time, introduceing an opportunity cyst in the form of inflexibility.

There is a better way.

Hyperconverged Infrastructure and the Modern Data Center

Hyperconverged infrastructure is one of the newer data center architecture options available and it brings organizations some serious solutions to the two problems discussed earlier in this paper. "Hyperconverged" refers to a class of data center infrastructure that comes ready to go. With a hyperconverged infrastructure, the IT staff no longer has to build each of the individual components that comprises the data center. Instead, IT pros buy prebuilt "units of infrastructure" that are then assembled like a building-block set. Everything snaps together to scale seamlessly. And, if you need more, you just buy another block.

Going back to the bus analogy, hyperconverged infrastructure is the furthest thing from a bus. Whereas IT staff had to build the shell of a bus before and then add seats, a hyperconverged infrastructure instead looks like a fleet of passenger cars. The engine in a one of these cars is more than capable of transporting six passengers. When your company needs to add more passengers (in this case, more business software) you simply buy another car with six more seats. This car is then added to the fleet, which is managed by a central dispatcher as one large resource pool versus individual units. In the world of IT, this is what is known as "linear scale out.". Much like the vehicle fleet dispatcher, IT architectures that adopt this scale out methodology are centrally managed from a single administrative interface.

In essence, you add sufficient resources to meet business needs and then, if you need more, you buy only what you need. Think of it as a "just-in-time data center.".

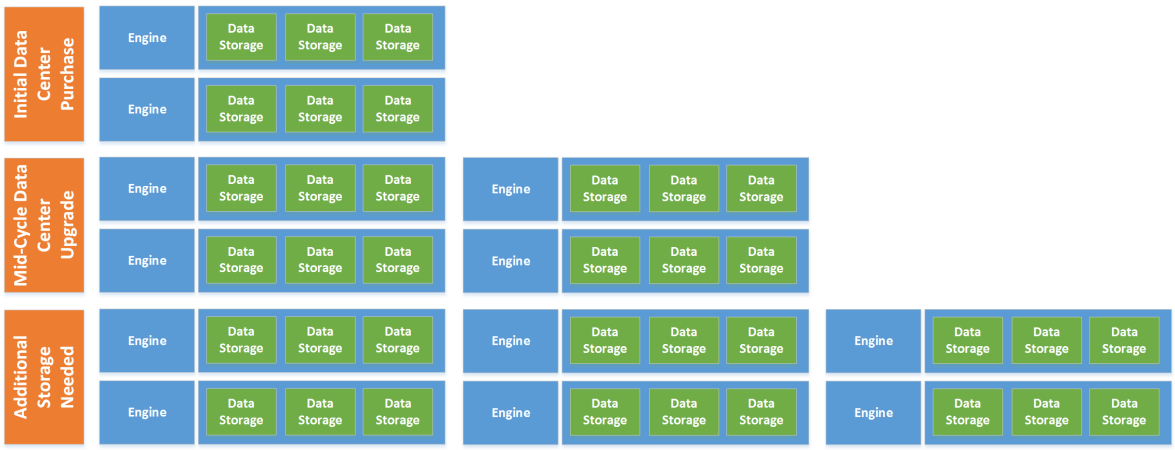

Figure 2: With hyperconverged infrastructure, a small engine is added each time systems are upgraded, meaking sure that workloads have enough horsepower to pull them.

Budget Predictability, Leveling, and Reduced Costs

On the budget front, hyperconverged infrastructure brings stability and simplicity to IT budgets. Because the organization is able to procure IT infrastructure when needed rather than having to buy big up front, there is less invested cost simply awaiting use. Your company buys what you need when you need it. The time-to-value for the incremental unit of investment is very short. You're able to utilize your IT investment very quickly and, if business needs to change, you buy an incremental unit of infrastructure to add capacity. A level IT budget is often the stuff of dreams. Hyperconverged infrastructure can make this dream more of a reality.

Additionally, money can be saved since IT staff can simply add another block to the network without having to re-engineer the existing systems. Of course, this also reduces the risk associated with major upgrades since new blocks of infrastructure simply tack on to existing systems. It no longer takes a team of people to perform an upgrade; a single person can deploy new infrastructure, often in hours rather than days or weeks.

Business Operational Improvements

Although staying in budget and keeping that budget stable are important goals, the real need is to ensure that the data center can fully support the important software needs of the business. The fact is that hyperconverged infrastructure components are, in aggregate, far more capable than many legacy data center implementations. They also come at a lower total cost of ownership. Each node in a fully hyperconverged infrastructure data center is imbued with tremendous processing power as well as with data storage components that provide ample capacity and, just as important, have storage devices that provide the speed that is necessary for many modern business applications. That performance factor is critically important and is often one of the weakest areas of legacy data centers.

Besides being easier to implement, hyperconverged infrastructure components are often far easier to manage than their legacy counterparts. Complexity and difficult management tools plague many IT systems, a situation that increases costs and lead time to new deployments.

Summary

Hyperconverged infrastructure provides organizations with a way to meet business needs with much less complexity, adding unseen levels of flexibility to what used to be rigid systems. As business needs dictate, IT departments that have adopted this agile architecture are far better positioned to help their organizations quickly meet emerging needs in extremely cost effective ways.

About the Authors

Scott D. Lowe is Senior Editor of EnterpriseStorageGuide.com and editor of HyperConverged.org. Scott has been in the IT field for close to twenty years and spent ten of those years in filling the CIO role for various organizations. Scott is also a micro-analyst for Wikibon and an InformationWeek Analytics contributor. In addition, Scott has also written thousands of articles and blog postings and has regularly contributed to such sites as TechRepublic, Wikibon, and virtualizationadmin.com. Because of his unique blend of skills (CIO/strategic & Engineer/tactical) Scott is also a sought-after resource for speaking engagements, and marketing collateral, including white papers and e-books, for a variety of technology firms.

David M. Davis is a well-known virtualization and cloud computing expert, author, speaker, and evangelist. David's library of popular virtualization video training courses can be found at Pluralsight.com. He holds several certifications including VCP5, VCAP, CCIE, and has been awarded the VMware vExpert award 6 years running. His website, covering the latest in virtualization is VirtualizationSoftware.com.