Colocation America®

CASE STUDY: HYPERCONVERGENCE VS CLOUD

By: Scott Alan Miller | On Behalf of Colocation America

Dear IT Pro,

Working with Colocation America and their data centers has enabled us at Niagara Technology Group to deploy new systems faster and move our large services clusters out of our own facilities. Running our own facilities was not cost effective nor did it provide us with the service levels that we wanted, as our expertise is not in data center management. Key staff that should have been working on our "wheelhouse" work were instead handling non-expertise data center tasks, which is a problem many businesses—especially small businesses—encounter. Moving our large clusters to Colocation America's data centers freed up our staff so that they could focus on core business operations.

For us, the shift in human capital was the biggest benefit. But we also found benefits in many other areas—the biggest being stability and availability of the environment, which was something we were not able to provide by having our hardware in-house. The move to Colocation America's data centers has given us faster networking with much higher reliability, so while the move was really one about shifting people, we have realized benefits in service levels as well, making it a slam dunk. Below, you'll find a study which states colocation provides better ROI than in-house IT, as well as why colocation is a better option for small businesses than popular public cloud options.

Sincerely,

Scott Alan Miller
@SCOTTALANMILLER on MangoLassi.IT, Spiceworks, Twitter

THE HIGH COST OF ON-PREMISES INFRASTRUCTURE

IT Infrastructure is a challenge for any company, and especially companies that are not large enough to implement their own, full-scale data centers. Like many things in IT, major challenges come in the form of lacking specific, seldom-used expertise, as well as lacking the scale to utilize singular resources effectively.

This lack of scale can come in many forms, with the obvious one being manpower. Managing a physical computing infrastructure uses unique skills that are separate from IT itself and are often desired to be available "around the clock." This can vary from security, to electrical, to cooling the facilities to "data center technician" style staff. Of course, smaller businesses simply do without these roles available to them, but this raises the cost incurred on a "per server" basis to maintain the infrastructure. Large businesses and dedicated data centers leverage an "efficiency of scale" to make the cost of physically housing an IT infrastructure lower—either by actually lowering the cost directly or by raising the quality and reliability of the equipment.

The cost-effectiveness of delivering power, cooling, and data center services is only one aspect of the business cost of IT infrastructure. Where many businesses attack this problem by reducing infrastructure investment and staff, which may counteract some amount of the upfront costs of the infrastructure, it generally does so to the detriment of availability and longevity of equipment. Whether it is a lack of ISP redundancy, an absence of diesel electric generators or shaving a year or two of service off of a server's life, these costs generally add up, often in ways that are difficult to quantify.

We see the effects of low quality infrastructure often come out in the behavior and expectations of smaller businesses. For example, in the enterprise data center an average server lifespan may be ten years or more, but smaller businesses often assume that a server is worn out and unreliable in seven or eight years. This increase in failure rate also leads to more concern about system failure. Smaller businesses often see a higher need to have redundant systems even when lower revenue would normally suggest otherwise. Small businesses are prone to investing heavily in high availability mechanisms, often at great expense, to mitigate a perceived risk of high system fail rates that larger businesses may be less likely to see. These factors can combine to create a high cost through more rapid system replacement and a tendency towards overbuying hardware—sometimes even doubling the otherwise necessary investment to protect against risks created by lower-quality facility's management.

Lower cost gear may run longer and more reliably in a high quality physical environment than more expensive, better engineered equipment will in a lower quality environment.

Of course, the most obvious components of lower reliability come from being unable to maintain redundant generators, independent power rails, adequate fuel supplies, uninterrupted power supply units, steady temperature and humidity, air filtration, and, of course, highly redundant multi-path WAN access. These aspects we think of frequently and are nearly out of reach for all but the largest companies. Even simple things like restricting access to only essential server room staff can be an insurmountable challenge in a small environment.

These challenges create an opportunity to find alternatives for the SME, SMB and SOHO business markets to look for ways to leverage combined scale. While many companies today turn to ideas such as hosted cloud computing, the associated costs to elastically expanding capacity often make this impractical as this same market struggles to have the ability to utilize that type of functionality. Cloud computing can be an answer in some cases, but normally only for the very smallest of companies for whom a single server is too much scale, or for those companies so large that they have a DevOps-style automation infrastructure capable of scaling elastically with load demands and workloads that make sense for this process. But these companies are the exception, not the norm. More often hosted cloud computing makes sense for only a specific subset of public-facing workloads and only in some cases.

For the majority of companies too small to create the scale necessary to build out their own full scale IT infrastructure, the answer is likely going to be found in colocation. It must be noted that there are potentially overarching locational or environmental factors that can make off-premises infrastructures impossible—or at least impractical. Most businesses, however, will not be subject to these limitations.

Colocation tackles the cost challenges of the smaller business market by generating the scale necessary to make high quality, dedicated information infrastructure facilities possible. This includes staff, WAN connectivity, environmental controls, power, and expertise. Cost savings can often come from surprising places

including lower power cost per kilowatt hour, lower cost of cooling and power conditioning, and higher real estate density.

It is often believed that colocation represents a cost premium service for businesses that have needs above and beyond the average, but in reality colocation should be chosen because it represents an opportunity to lower costs while also improving reliability. In most cases, colocation will actually bring a cost savings on a month-by-month basis providing for an impressive ROI potential over time as the initial cost can be equal, or similar to, other investments. The ongoing monthly cost can be lower and, perhaps more importantly, the costs can become far more predictable with fewer risks and unexpected expenditures.

Because the cost of services are potentially very granular it is actually far easier for colocation to lower the overall expenditure than is generally believed. For example, a small business with just one or two servers would still need certain basics such as air conditioning and UPS support plus footprint space and security; all dedicated for only a very small amount of equipment. In a colocation facility these servers may represent less than one percent of the cooling of a large, high efficiency cooling system, or just a small fraction of a large UPS, etc.

Much like I noted at the beginning, colocation also frees IT staff from performing data center functions—at which they are generally untrained and poorly qualified—to focus on the tasks at which they are more valuable and intentioned. Exactly calculating ROI can be challenging because individual cases are very unique and depend heavily on the workloads, use cases, independent needs and environmental factors of an individual business, and the colocation options considered. But it should be approached with a mindset that colocation does not present only an opportunity for improvements in the quality or reliability of IT infrastructure services, not that it can represent a return on investment but that it may, in fact, do both of these things on top of fundamentally lowering costs overall.

SMB, HYPERCONVERGENCE, AND COLOCATION

Given that colocation is a great ROI for small businesses, hyperconvergence and colocation are probably this perfect match for the small and medium business market.

For years, the SMB market has struggled with computational systems that were either rudimentary compared to their big business competitors, or the SMB was stuck paying for systems and tools that were significant overkill for them. Lacking the scale of a large business or enterprise, SMBs have simply been left stuck in the middle with systems that were underutilized, but too complex and costly to use more thoroughly. And just as SMBs struggle with costly and complex computing systems, they also struggle with providing adequate environments to put these computer systems.

The problem is that SMBs, by and large, simply need far too little "capacity" or total size from their systems to warrant the things that make sense in the enterprise—the scale just isn't there. An enterprise can build a data center with all the trappings of a top-tier facility for only a few dollars per server, but a SMB might only need two servers. To do the same could cost a hundred thousand dollars per server!

But today, the needs of SMBs are not that much different than those of enterprise businesses. Of course, every business of every size is unique, but increasingly SMBs also need to maintain systems around the clock, with long data retention, good performance, and access from anywhere in the world, any time. These things are trivial for a big business and staggering for an SMB.

This is where hyperconvergence and colocation come into the picture. Hyperconvergence tackles this from the computational systems side while colocation tackles this from an environmental and physical side. Together, these technologies can allow an SMB to operate far more like enterprises have in the past.

Hyperconvergence changes the landscape by introducing computing, storage, platform management, and high availability as a single, cohesive system, taking many of the most expensive and complex pieces of IT infrastructure and collapsing them into a size that SMBs can consume. Many of the best benefits that enterprises have enjoyed can now be had by nearly any SMB through hyperconvergence by lowering the cost of the physical investment while easing management concerns around complexity and a need for advanced skills. Lower cost, greater agility, more flexibility, and less investment per workload radically changes the SMB IT picture.

Similarly, colocation takes the physical aspects of the environment and makes it available to the SMB market in a consolidated way that allows each SMB to invest in a small amount of real estate at enterprise quality and feature level while paying only for the size that they need, and only for as long as they need it, reducing capital expenditure risks. Colocation gives SMBs the ability to scale

up and down in a reasonable way (not minute to minute).

Because cloud computing is such a bad fit for the general SMB market, the combination of hyperconvergence with colocation can bring the majority of the beneficial ancillary aspects of hosted cloud computing to the SMB without encumbering them with the costs and complexities of elastic workload models.

With hyperconvergence and colocation, companies can reduce, or effectively eliminate, all platform management needing no more than basic capacity projects (normally several weeks ahead) and license management, with no need to understand storage, high availability, virtualization, imaging, and in many cases, even backups. Workloads can be created and destroyed rapidly and with ease. Everything can be centralized and managed through a single pane of glass or platform portal. Unlike cloud that assumes management via scripts and APIs, hyperconvergence assumes snowflake management at the highest level while not barring DevOps management in any way.

While some companies will have technical needs that may keep hyperconvergence from being their solution—or at least, the only solution—the majority of the SMB market should most likely be looking at the marriage of hyperconvergence and colocation as the “go to” solution for their computational infrastructure and support.

The choice for the SMB, then, is whether to invest in a public cloud option or a hosted hyperconvergence solution. *What are we comparing?*

In the first corner: public cloud. Services like Amazon AWS and Vultr. An average mainline Windows server there is about $96/mo and Linux is about $40.

In the second corner: hosted hyperconvergence. Using Colocation America's hosting with Scale Hardware is the easiest and very comparable as it is enterprise, full support, single price, and a Tier IV data center with Amazon-like full time support. Comparing these two is very useful because both are off-premises approaches that overlap in what they provide. Two different approaches to essentially identical needs for SMB customers.

Examples of what would be needed:

- 4x VMs of 8GB, 4vCPU, 80GB each
- 2x VMs of 4GB, 1vCPU, 500GB each
- 1x VM of 32GB, 8vCPU, 160GB
- 3x VM of 1GB, 1vCPU, 20GB each

Working based on Colocation America's $399 for a quarter cab per month, that's 10U which is enough for a firewall (1U), redundant switches (2U), a backup device (2U) and a rather significant cluster (5U of 1100, 1150, 2100 or 2150 nodes). That's potentially a very large cluster if the 1150D style nodes are used. Far beyond the needs of a normal SMB.

If we needed more space, going to a half cabinet at Colocation America is only a jump to $699/mo and gives us enough room for a maxed out Scale HC3 cluster, plus room to add additional storage options.

So a really quick set of numbers....

- Starter Cluster of Three 1100 Nodes ($25K)
- Two Redundant Switches ($6K)
- NAS to use as a backup target ($4K)
- 3x Windows Data center Licenses ($18K alone)
- Enterprise Router, rackmount ($300)

Looking at 11 8GB VMs....

Result: break even on a five year cycle. Almost exactly the same cost. Except for a few things: The hyperconverged solution is high availability while the hosted cloud option is not. And the hyperconverged solution includes backups, the hosted cloud solution does not.

There are lots of reasons that the hyperconverged solution here is the better option:

To add backup to the hosted solution we'd need an additional $5K.

To add high availability, we would add an additional $53K, doubling the base cost!

The hyperconverged solution has already invested in massive amounts of Windows Server licensing for unlimited VMs on the given host. We could move to more, smaller VMs or simply add additional VMs for free after this point. Additional VMs have no licensing overhead again until we add an entire additional node unlike the hosted cloud solution that continues to pay Windows licensing cost for each additional VM used.

The hyperconverged solution has 128GB RAM available (after HA overhead is removed) and we are using only 88GB in this example. Growing workloads by roughly 50 percent is available without additional cost.

The hyperconverged solution has more available storage. The hosted cloud has only 1.3TB storage total to use.

The hyperconverged solution is far more flexible. Instead of locking each VM to exactly 8GB of RAM, 150GB of storage we can tune each VM as needed. A VM that needs only 6.5GB of RAM need not use a full 8GB just because the steps are 4, 8 and 16GB, but can be tuned to exactly what is needed. Same with storage. This makes the potential density of VMs much higher.

To match the high availability and backup of the hyperconverged solution, the hosted cloud would cost $121K compared to the $53K of the hyperconverged cluster. At this size, only five or six VMs of the 8GB size would be necessary, if backups and HA are needed, to equal the cost of a hyperconverged cluster after five years!

But now we need to consider the colocation costs. This difference in cost would only matter if we were getting our colocation for free or "already covered" by other costs. In our example here we want to look at Tier IV (top tier, enterprise) data center colocation which for this is $400/mo. That's an additional $24,000 over five years.

This makes our numbers a little harder to compare. With 11 VMs, if we want backups but not HA for any workloads, the cloud computing comes out cheaper. But it is not apples to apples. If we want HA, the cloud solution is nowhere close.

The final numbers are:
Hosted Cloud Computing: $121K
Hyperconverged in Colocation: $77K

The cloud option has more flexibility if workloads are transient. The hyperconverged option has far more flexibility for growth at essentially no additional cost.

Let's work from the other direction now: What would it cost for a hosted cloud to have the same capacity as a full HC3 cluster from the example above? We already know that the Scale HC3 in colocation with full accouterments would be $77K for five years.

The limiting factor here is really RAM. The Scale HC3 has 128GB which would support 16 VMs of the example size above. That would be just over $161,000 on the hosted cloud with HA and

backups. Pretty much once we have breached the minimum threshold of hyperconvergence benefit, the cost of the cluster continues to improve with each additional workload that we add. Individual nodes with more RAM, more or faster storage, more or faster CPUs allow us to load up more VMs at minimal additional cost. Adding another node is a small investment, about $13K from our example above when we include the Windows data center licensing, gives us the ability to run 50 percent more workloads for only that small additional investment.

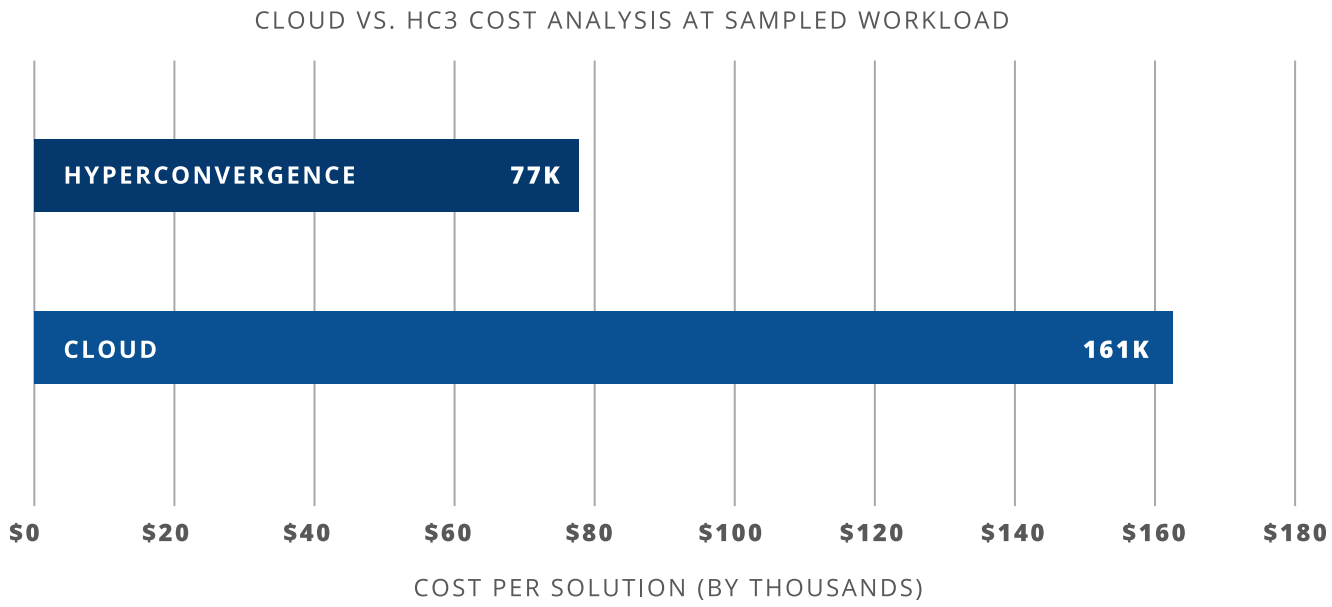

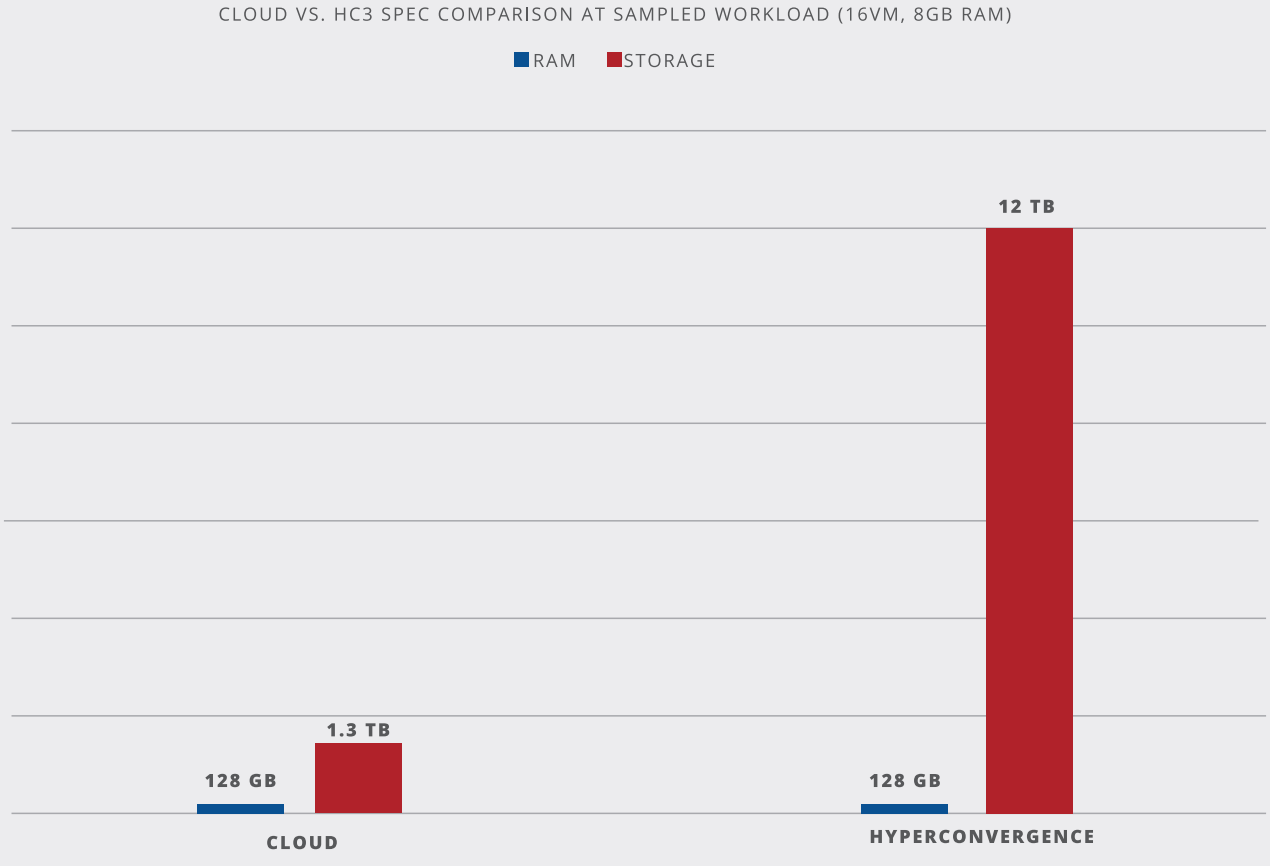

CONCLUSION

By outsourcing our on-premises IT to Colocation America's data centers, we were able to save on infrastructure costs, and allow our staff to return to their core business operations. For a SMB, combining the infrastructure cost benefits of colocation with hyperconvergence allows for enterprise-level IT solutions with availability, scalability, and storage for a greater ROI than popular public cloud offerings over the long term.