

Is Hyperconverged Cost-Competitive with the Cloud?

An Evaluator Group TCO Analysis Comparing AWS and SimpliVity

By Eric Slack, Sr. Analyst

January 2016



Evaluator Group

Enabling you to make the best technology decisions



The Cloud, once a radical idea in IT, is now mainstream. Whether it's email, backup or file sharing, most consumers probably use a cloud service or two. Similarly, most IT professionals are familiar with cloud service providers such as Amazon, Google and Microsoft Azure, and many companies have moved at least some of their information technology processes into the cloud. In fact, the cloud has become so popular it's easy to assume that running IT applications on-premises is not cost competitive with a cloud-based service. In this report Evaluator Group will test the validity of that assumption with a TCO (Total Cost of Ownership) model analyzing a hyperconverged appliance solution from SimpliVity and a comparable cloud service from Amazon Web Services (AWS).

Total Cost of Ownership

TCO is a calculation designed to expose the true costs associated with a specific system or piece of capital equipment, since the "sticker price" often understates the actual cost of implementation incurred by the business. This is especially true for IT systems, given the level of complexity of traditional infrastructures.

Many of the costs associated with IT equipment can be overlooked, including things like design and evaluation, installation, integration with existing systems and training. In addition to capital expenses (CAPEX), operating expenses (OPEX) need to be included in TCO as well, especially personnel or administrative overhead required to keep the system running and to expand and upgrade it throughout a given lifespan.

TCO models are often used to support capital purchases, but also to help managers make strategic decisions, such as whether to run a certain business process in-house or to out-source it. IT processes are good candidates for this "build or buy" decision since most applications can be run on servers in a public cloud just like they can in the company data center.

The Cloud Assumption

Everyone from CIOs to consumers hears about the benefits of the cloud on a regular basis. Truly ubiquitous, it's easy to assume that the cloud is fundamentally less expensive than running on-site infrastructure. Part of this thinking may be due to the economics of the "web-scale" data centers that cloud providers run, but probably more to the complexity of setting up and running traditional servers, storage and networking. Many of the companies using cloud services have no doubt run TCO calculations of their own and at one point found it to be the right decision. And with the trend towards leaner staffing and lower levels of expertise in IT organizations, the "cloud assumption" would seem to be further supported. But this assumption may no longer be accurate.

The cloud assumption has historically been based on a comparison with legacy infrastructures that consist of separate systems for each IT function, such as compute and storage, even data protection. Given the cost realities, many companies found the idea of running applications in the cloud very appealing and the option to pay for them on a monthly basis certainly simplified the decision process. But what about hyperconverged systems?

Hyperconverged Infrastructures

These solutions combine storage and compute components with a hypervisor and management software to create a comprehensive compute solution. Some vendors go even further to include features such as data protection and built-in data efficiency. Hyperconverged systems appeal to companies for many of the same reasons the cloud does. They are easy to set up and run and easy to use, providing operational agility and greatly simplifying the overall infrastructure, compared with traditional IT solutions.

Hyperconverged infrastructures change the traditional TCO calculations that have been used to justify cloud decisions. Specifically:

- Hyperconvergence greatly simplifies the design process
- “Plug and play” format means very little implementation
- Internal integration with hypervisor management lowers operating complexity
- VM-level focus simplifies storage-related tasks in virtualization environments
- Scale-out architecture is designed for cost-effective, incremental growth

While the trend in IT organizations is toward lower staffing and expertise levels, the emergence of hyperconverged infrastructures fundamentally changes the TCO economics. This begs the question, “Is the cloud assumption still valid?” For this reason, Evaluator Group compiled the following information on the total costs associated with SimpliVity’s hyperconverged infrastructure and a comparable infrastructure running in Amazon Web Services.

Hyperconverged vs. Cloud TCO Model

For this comparison we used an OmniCube CN-3400 hyperconverged appliance and Amazon’s EC2 cloud compute service. In addition to OmniCube SimpliVity offers the “OmniStack” line of products, with

partners such as Cisco and Lenovo, which provide similar features and benefits to the OmniCube systems evaluated in this report. To make these systems highly available, we chose an N+1 configuration for the OmniCube cluster, with “N” being the number of active SimpliVity “nodes” (each supporting 103 VMs), and one redundant node. It should be noted that Amazon offers a 99.9% availability for the EC2 service used in this comparison, which is a significantly lower level of availability than what can be achieved with the OmniCube N+1 configuration.

The baseline configuration in this model was a three-node cluster, supporting 206 VMs (2 data nodes x 103 VMs/node), as shown in Figure 1. Costs for VM counts of 309, 412 and 515 are also calculated in the model corresponding to four-, five- and six-node SimpliVity clusters. However, this could certainly be extended since upgrading these hyperconverged systems with new nodes is easy to do and doesn’t require downtime or data migration. And, as shown below, the economics of these systems get stronger as they are scaled.

We assumed a 3-year period for the infrastructure, a typical IT capital equipment lifespan, and all one-time costs were calculated based on a 36-month simple amortization. Leasing is also an option for many companies, allowing them to turn the capital expenditure into a monthly expense, although it’s not included in this model.

For the baseline comparison with the SimpliVity solution we chose Amazon’s “3-year All Upfront, Reserved” pricing. In the model, the substantial upfront fee of this option is amortized into 36 monthly increments, similar to the hardware costs in the SimpliVity hyperconverged infrastructure. It should be noted that this is the lowest cost option for AWS. Amazon also offers 1-year and “on-demand” options, which are significantly more expensive on a per-VM basis (Figure 1).

Amazon Web Services

AWS offers Elastic Compute EC2 web services on a per-VM basis, so we priced configurations starting at 206 VM instances and then in increments of 103 VMs. The following list shows the specific component charges for the baseline 206-VM AWS infrastructure, based on a configuration created on the AWS website in January 2016.

Amazon EC2 Service (US-East)

- Compute – “M3.medium” instance, 1 vCPU, 3.7GB memory, 100% utilized
- EBS Volumes – For each instance, 100GB of “General Purpose SSD” storage was chosen
- EBS IOPS – for each instance, we chose 300 IOPS for this 100GB of storage
- EBS Snapshots – for each instance, a 10% daily change with 30-day retention was chosen

- AWS Data Transfer In – we estimated 0.5TB of data per month, per 103 VMs
- AWS Data Transfer Out – we estimated 4.5 TB of data per month, per 103 VMs
- AWS Support – we chose Business Support for all AWS services

The total monthly cost for this 206-VM infrastructure was \$12,860.58 or \$62.46 per VM. (Figure 1)

SimpliVity OmniCube

The OmniCube calculations include CAPEX costs for hardware, software and maintenance, plus OPEX costs associated with data center overhead and administration. The following list shows these specific components of the total cost of a three-node cluster.

- OmniCube acquisition price – the list price, less a standard discount, for three SimpliVity CN-3400 nodes required for an N+1 configuration.
- OmniCube maintenance – the 3-year maintenance cost for these nodes
- vSphere licenses – the 3-year cost of a hypervisor license for each 2-CPU node, assuming a 10% discount off list
- vSphere support – the cost of hypervisor support for licenses required for a given configuration
- Power, cooling and rack space – the 3-year data center facility cost assuming \$100 per node, per month

Administrative Support

Perhaps the largest component of total operating expenses is the cost of administering the infrastructure. For this model, we used the following calculation for personnel required to run each OmniCube node:

Hourly cost of fully burdened employee (\$150,000 per year = \$75/hour) * hours required to manage each node (1 hour per week) * 52 weeks * 3 years.

The total monthly cost for this SimpliVity OmniCube cluster is \$12,292.02 or \$59.67 per VM, roughly 4% less than the comparable AWS price. The chart in Figure 1 shows the 3-year costs for the SimpliVity solution and the AWS infrastructure described above. It also shows the costs for AWS pricing based on shorter timeframes which, again, are significantly higher.

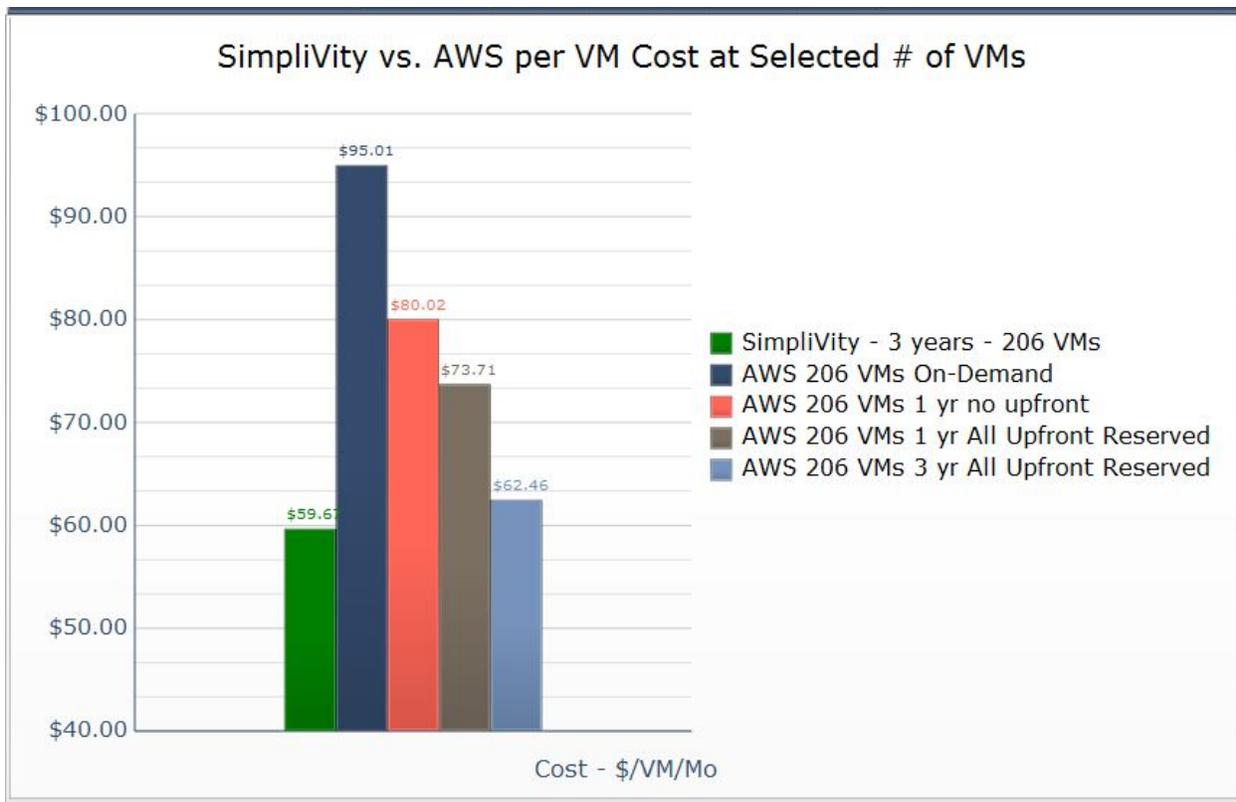


Figure 1

Scaling the Infrastructure

Expansion of this hyperconverged infrastructure involves adding nodes to the cluster. Given the N+1 configuration, each additional node increases the VM count by 103, effectively distributing the cost of the redundant “+1” node across more data nodes. This improves the economics of the system. In comparison, the AWS infrastructure is essentially unchanged on a per-VM basis.

The graph in Figure 2 shows the widening differential between the monthly cost per VM on SimpliVity OmniCube and AWS, as the infrastructure scales from 206 to 515 VMs (two to five nodes). This differential would fluctuate somewhat as the number of redundant nodes increases, but the SimpliVity system is still less expensive than the cloud solution at all levels of scale.



Figure 2

Why Hyperconverged TCOs are Lower

As mentioned above, hyperconverged solutions, like SimpliVity, change the TCO calculations historically used in IT decision making. By simplifying the physical infrastructure, using non-proprietary hardware and reducing or eliminating many of the tasks typically associated with implementing hardware and software they lower total costs, both capital expenditures and operating expenses.

Capital Expenditures

Hyperconverged systems combine storage and server functions, potentially reducing the overall cost by reducing the number of physical systems that must be purchased. They also leverage the economics of industry-standard, “commodity” hardware to further push acquisition costs down compared with proprietary storage and compute systems.

Systems like SimpliVity eliminate the cost and complexity of implementing a complete infrastructure for virtualized workloads using open systems components. The “design” process consists of selecting enough nodes to support the expected VM count and configuring those nodes to fit the expected workloads. There’s no long, involved installation phase; in fact, most hyperconverged appliance

solutions can be deployed in less than an hour with very little need to “integrate” the cluster with existing systems. Expanding the system is simple as well, just add nodes to provide the required resources and the cluster re-balances itself.

As a comprehensive solution, storage services and management software, and in some cases built-in data protection, come pre-installed as part of the package. These services typically include storage efficiency features such as inline deduplication, which when integrated into the base product can provide significant reduction in the amount of storage consumed. Related to efficiency is the way scale-out topologies can expand incrementally, improving capacity utilization and further reducing CAPEX.

Operating Expenses

Management of the hyperconverged cluster is usually integrated into the hypervisor, allowing an administrator to handle more VMs than is possible with a traditional shared storage infrastructure. Hyperconverged solutions, like SimpliVity, are designed with VM-centric management, the ability to provision storage and create management policies for each VM rather than for each storage volume. This further simplifies the day-to-day “care and feeding” of the system making the infrastructure manageable by employees without specific storage skills.

When more resources are required many hyperconverged systems support non-disruptive expansion, eliminating what has typically been a complex process. All these factors add up to an infrastructure that can be managed by fewer people, and those with less IT expertise. This is one of the most significant differences between the TCO of hyperconverged systems and a traditional infrastructure of separate elements for compute, storage and data protection.

Cumulative Costs

As mentioned earlier, TCO models are also used for strategic decisions, such as whether to put new or existing IT infrastructure into the cloud. Part of this strategy discussion involves the long-term, cumulative costs of different options. Sometimes expressing costs incrementally, on a per-VM basis as in this model, can minimize the impact of each alternative. In these instances, a cumulative cost calculation is often useful, such as the graph in Figure 3.

This graph shows the AWS options are \$1.14M to \$1.7M, when the cost of 515 VMs are added for the entire 36-month period. However, the SimpliVity solution is only \$885K, 22% to 49% less when totaled over that same timeframe. What would seem to be a modest per-VM cost differential can extrapolate out to hundreds of thousands of dollars in a few years.

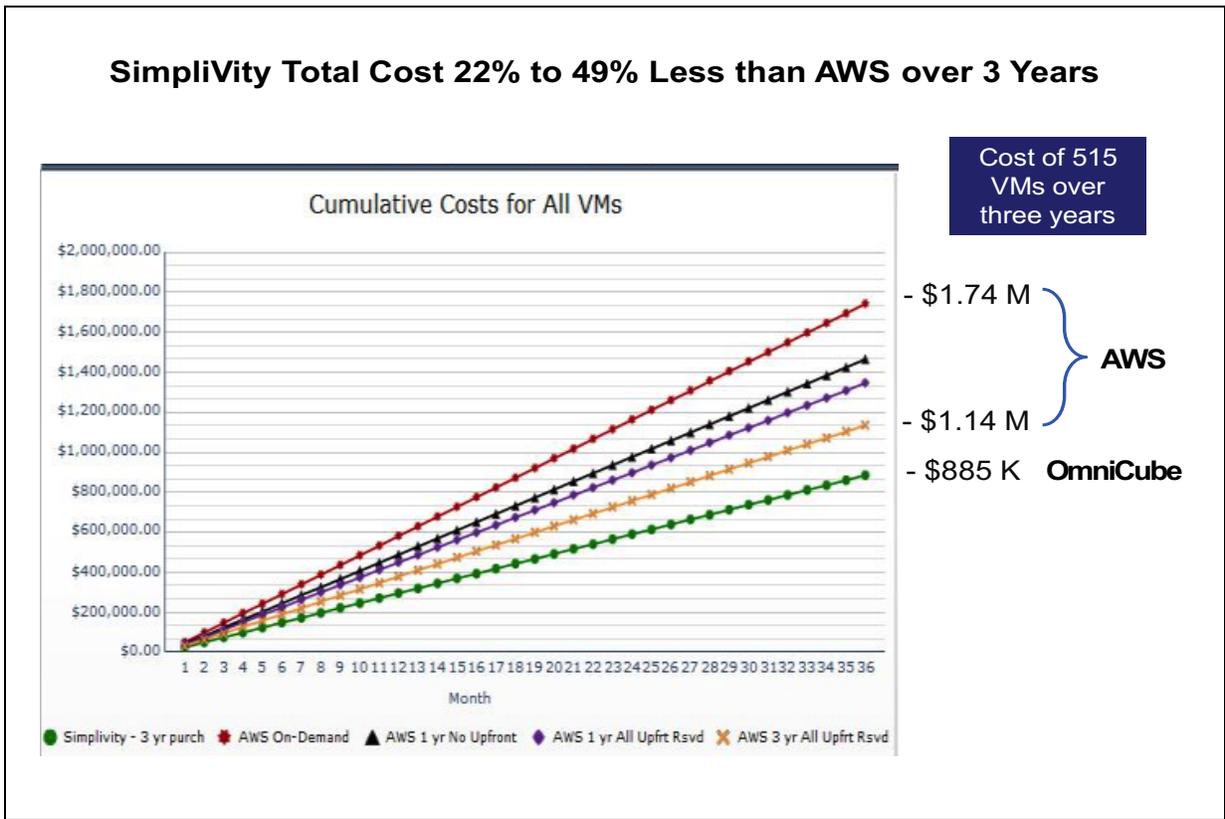


Figure 3

Summary / Conclusion

The Cloud, once a novel idea in IT, is now an accepted, mainstream alternative. This fact has contributed to the assumption that the cloud *must* be less expensive than an in-house infrastructure. For many companies the justification for choosing the cloud has been the high cost of implementing traditional servers and storage components, and operating them over time. Often, a Total Cost of Ownership calculation has been used to make this comparison and support the “cloud assumption”.

In this report Evaluator Group shows that assumption is no longer valid, based on a TCO model that compares SimpliVity hyperconverged infrastructure with a solution from AWS. Due to their comprehensive, low-cost hardware architectures and advanced efficiency and management features, hyperconverged infrastructure can cost significantly less than a comparable cloud solution. And, that cost differential gets better with larger systems and longer timeframes.

About SimpliVity

SimpliVity hyperconverged infrastructure provides a scalable, modular, 2U building block of x86 resources that offers all the functionality of traditional IT infrastructure— in one device. It assimilates storage, compute, hypervisor, real-time deduplication, compression and optimization into a single appliance, along with comprehensive data management, data protection, and disaster recovery capabilities.

Designed to work with any hypervisor or industry-standard x86 server platform, the SimpliVity solution provides a single, shared resource pool across the entire IT stack, eliminating point products and “siloes” IT architectures. The solution is distinguished from other hyperconverged products by three attributes: accelerated data efficiency, built-in data protection and global unified management.

Inline deduplication, compression and optimization are applied to all data at inception, reducing resource consumption (storage and CPU) while increasing application performance. SimpliVity systems include native data protection functionality, while eliminating the need for special-purpose backup and recovery hardware or software. VM-centric management enables system administrators to manage all resources and workloads centrally, using familiar interfaces such as VMware vCenter.

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. www.evaluatorgroup.com Follow us on Twitter @evaluator_group*

Copyright 2016 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.