

The Cost of Using the Public Cloud

An Evaluator Group TCO Analysis Comparing AWS and VxRail

By Eric Slack, Sr. Analyst

February 2017



Evaluator Group

Enabling you to make the best technology decisions



For many companies the appeal of the public cloud is very real. For tech startups, the cloud may be their only option, since many don't have the capital or expertise to build and operate the IT systems their businesses need. Existing companies with established data centers are also looking at public clouds, to increase IT agility while limiting risk. The idea of building-out their production capacity while possibly reducing the costs attached to that infrastructure can be attractive. For most companies the cloud isn't an "either-or" decision, but an operating model to be evaluated along with on-site infrastructure. And like most infrastructure decisions the question of cost is certainly a consideration.

In this report we'll explore that question, comparing the cost of an on-site hyperconverged solution with a comparable set up in the cloud. The on-site infrastructure is a Dell EMC VxRail™ hyperconverged appliance cluster and the cloud solution is Amazon Web Services (AWS).

Backed by specific research, Evaluator Group helps IT end user clients with strategic infrastructure decisions regarding technology, economics, etc. Through this work we have developed tools, like TCO models, that can be applied to more general comparisons such as that featured in this report.

TCO Model

Total Cost of Ownership is a concept that's commonly used when making capital purchases. When you buy an iPad, tablet or other consumer product, the "sticker price" is essentially all you need to make a meaningful cost comparison. But figuring the comparable costs for IT infrastructures is more involved. The cost of owning these assets includes the costs of buying and running them, which TCO models are designed to provide.

For traditional infrastructures, Capital Expenses include all one-time costs, such as the acquisition cost, plus all the design, installation and training. Operational Expenses start after installation and include the on-going costs of running the data center and paying the people who operate the equipment, plus recurring maintenance. Hyperconverged Infrastructure (HCI) appliances, like Dell EMC VxRail, have simplified the traditional compute environment to reduce both Capital and Operational Expenses.

HCI Benefits

A big part of the cost associated with capital purchases, like IT systems, is design and implementation. For a moderately sized compute environment this includes choosing servers, storage systems, a storage network and hiring the integration expertise to put the whole thing together and make it work. Hyperconverged Infrastructures greatly simplify this job by essentially reducing the decision to choosing which model of appliance or node to buy, unpacking the boxes and connecting them to the Ethernet. The net effect of this simplification is a lower total cost of ownership, and one that's much easier to calculate.

With that established, let's look at the specifics of this TCO comparison. We started with a cluster of VxRail appliances designed with enough processing power and storage capacity to support a variety of workloads, from "tier 2 or tier 3" applications to some core IT services as well. Then we priced out a comparable configuration using Amazon Web Services (AWS).

The On-site Infrastructure

We configured a 6+1 node cluster of 1U VxRail E Series appliances with an additional node to ensure high availability. This configuration provided two vCPUs, 4 GB of memory, ~100 GB of all-flash storage and 500 IOPS per VM. Storage capacity includes a 10% snapshot allocation. All-flash storage was chosen over a hybrid configuration (disk and flash) because flash provides more reliable, consistent performance and can support any workload, from a storage perspective. And, it allowed us to use deduplication (only available with all-flash) as well as erasure coding, (RAID 5) to gain space efficiency. Also included in this cluster are a pair of 10 Gb Ethernet switches, plus cabling and support.

The cluster supports 650 VMs, or 108 VMs per node, each with the resource configuration listed below:

- Dual E5-2660 v4 CPUs, 28 cores total (14 per CPU)
- 512 GB per node memory
- 800 GB cache SSD
- 10.47 TB capacity SSD storage

For the TCO comparison we calculated a cost per month using a 3-year simple amortization with no interest. The assumption was that companies making this kind of comparison would have their own funding framework and could easily add in the cost of capital from internal sources or through a commercially available lease.

Facilities costs were estimated at \$100 per month for each node of the hyperconverged cluster. For admin time it was assumed that the entire cluster would require 2 hours per week, a conservative estimate given the fact that VxRail systems are tightly integrated with VMware allowing admins to use the tools they're already familiar with. The cost was estimated with the following formula:

Hourly cost of a fully burdened employee (\$150,000 per year or \$75 per hour)

* 2 hours per week * 52 weeks * 3 years

TCO Calculation

Arriving at total cost was actually very simple. We totaled the purchase price for the VxRail hardware, VMware licenses and switches (net of typical discounts), plus implementation services, 3 years of hardware and software maintenance and ongoing overhead. The cost line items were:

• 7 x VxRail E Series appliances (incl. VMware licenses)	\$314,587
• 2 x 10Gb Ethernet switches (including cables)	24,245
• Implementation services	37,310
• Prepaid hardware and software support for 3 years	125,419
• Facilities costs for 3 years (power, cooling, rack space)	25,200
• IT Administrator time	23,400
	<hr/>
	\$550,161

This brings the total monthly cost for this 7-node infrastructure, designed to support 650 VMs, to \$15,282.25 (\$550,161 / 36 mos) or \$23.51 per VM.

The Cloud Infrastructure

The “elastic” economics of the public cloud can be great for workloads that scale in unknown or transient ways. Service providers, like AWS, offer different types of compute and storage to better support dynamic workloads and enable resources to be matched to demand. For web applications, as an example, this creates an opportunity for cost savings, making them a good fit for the cloud. For other workloads, it’s a different story, as the IT Director at a mid-west medical testing company described it:

“For the core services piece, it was difficult to predict what the [cloud] pricing would be, and complicated to decipher. For burstable app delivery capacity, like web services, the cloud can’t be beat. But you need core services (the boring stuff) all the time and at more-or-less the same compute/memory/net consumption levels.”

AWS Elastic Compute Services (ECS) is offered in dozens of different of compute instances, each providing different combinations of resources, different utilization profiles, different storage options, etc., to match variable workloads. But traditional IT applications were designed to be connected to *dedicated* resources, expecting storage capacity, CPU cores and memory to be available all the time. This means a cloud infrastructure that’s comparable to on-site infrastructure needs to be 100% utilized and dedicated. For this comparison, we chose the following ECS configuration:

Compute

“C4.large” ECS instance, 100% utilized, providing 2 vCPUs and 3.7GB of memory

Storage

“EBS General Purpose SSD” was chosen, providing 300 IOPS* per 100GB of allocated storage. The VxRail configuration supplied 500 IOPS per VM, which would require either using AWS “Provisioned IOPS” storage or increasing the amount of General Purpose storage allocated per VM, both of which would have increased costs. It was determined to leave the storage at 100GB of EBS General Purpose SSD storage.

Snapshots

A 10% snapshot allocation per month was added, commensurate with VxRail configuration.

Data Transferred

AWS tracks the amount of data transferred **into** the ECS instance from the internet, **out of** the ECS instance to the internet and **between** ECS instances, and charges for the second two. For this model we estimated that each VM would generate the following amount of data movement each month:

- 10GB of data transferred in
- 30GB of data transferred out
- 60GB of data transferred between ECS instances

Financing

For this TCO comparison, we included the following four financing options from AWS. Each is represented by a different bar in Figure 1.

“On-Demand” refers to buying compute, storage and data transfers as needed. It requires the least commitment but charges the highest cost.

“1 yr, No Upfront” is like On-Demand but with a one-year commitment.

“1 yr All Upfront Rsvd” (Reserved) requires the costs for one year of compute and storage to be pre-paid.

“3 yr All Upfront, Rsvd” extends the 1-year plan to 3 years.

Results and Analysis

Figure 1 below shows the average monthly cost per VM for on-site infrastructure (green bar) that would be expected to support general IT applications and several cloud pricing options at the 650-VM level.

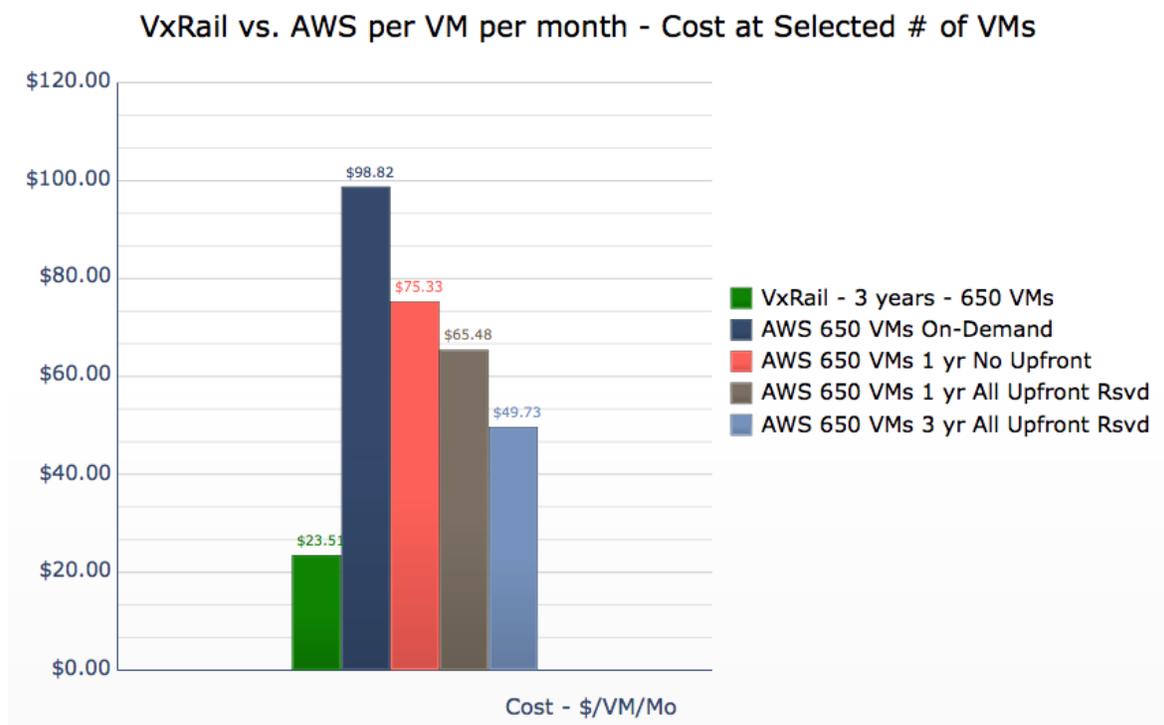


Figure 1: Monthly Cost per VM for On-site and Cloud Infrastructure

Figure 2 shows the same monthly costs extended out for a 3-year term, typical for IT capital purchases.

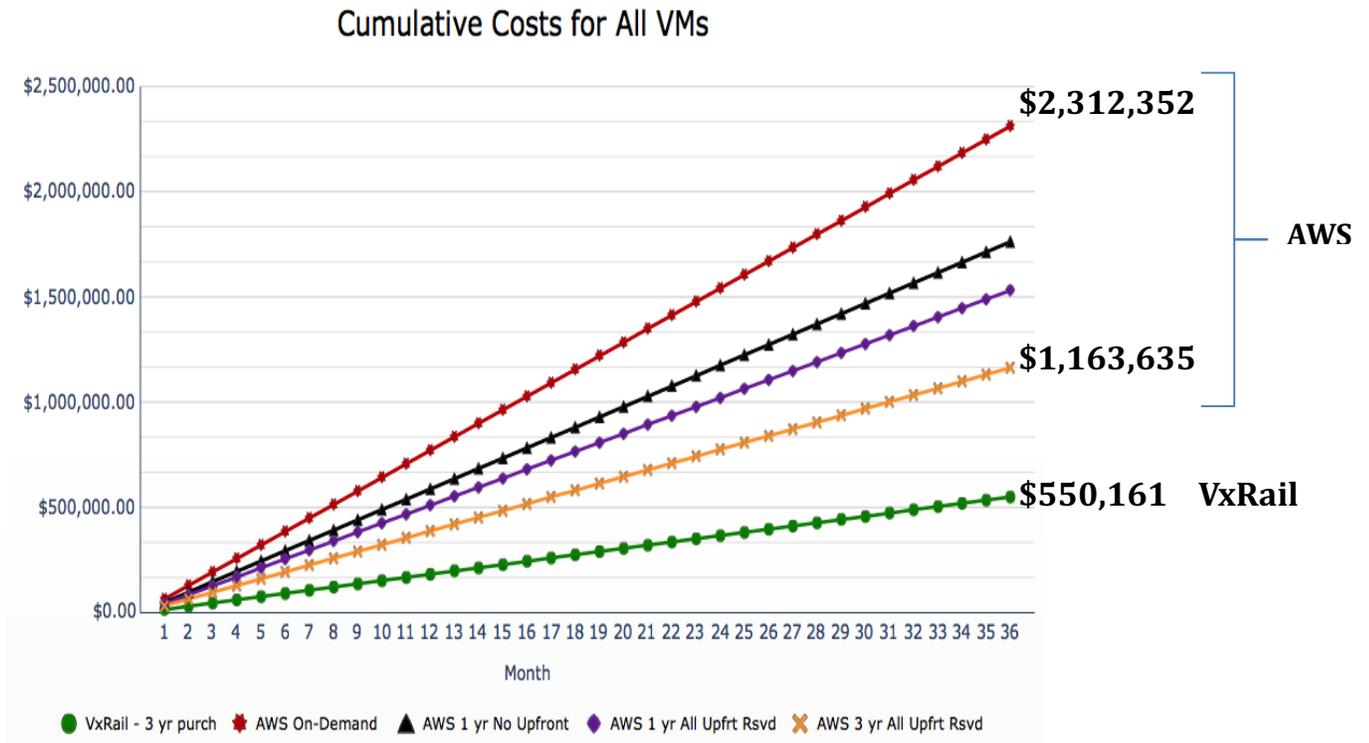


Figure 2: Total 3-Year Cost for On-site and Cloud Infrastructure

The Cloud Costs More – a Lot More

As the results show, running a given number of VMs in the cloud is more expensive than running them on a hyperconverged infrastructure cluster of VxRail appliances. The On-Demand option is over 4x the on-site solution and the next two options are about 3x. In a best-case comparison, the 3 yr All Upfront Rsvd AWS offering is still over twice the cost of a VxRail cluster, including switches, licensing and overhead.

For a three-year period, that difference is more dramatic. It comes to over \$2.3 million for AWS On-Demand. For VxRail the cost for 36 months is \$550,161. So what’s causing this disparity?

Hyperconverged has Evolved

The cost of traditional infrastructure, as well as the administrative overhead can be significant, compared with a hyperconverged solution. As detailed above, HCI technologies greatly simplify the

design and implementation processes compared with a traditional IT infrastructure. They also reduce hardware costs by replacing proprietary storage systems with commodity storage devices in industry-standard server platforms.

The data services that have historically added significant cost to arrays are now implemented in software-defined storage, improving the overall capability of these scale-out systems to deliver fast and efficient operations and further reduce the total cost of ownership. But there's more to the story.

Even compared to a year ago, things have changed with hyperconverged appliances. As the market has matured HCI solutions have gotten more powerful, storage has evolved to include more flash, or all flash, and the cost of hardware keeps going down. Server manufacturers like Dell EMC have the ability to leverage significant cost advantages over others, like smaller start-ups, that must buy hardware from the manufacturer or sell their software through those same OEMs.

As hardware has gotten more powerful and less expensive, we've seen features and functionality go up as well. HCI solutions now routinely include multiple levels of data protection and availability, advanced data reduction and storage efficiency and policy-based management capabilities that make them easier and less expensive to operate.

A Different Cloud Option

Another area that HCI solutions are addressing is the need for a hybrid cloud. In fact, [Dell EMC's Enterprise Hybrid Cloud \(EHC\)](#) platform now supports VxRail appliances, enabling companies to add cloud flexibility to their on-site infrastructure investment and enjoy a "best of both worlds" situation. Hybrid clouds can offer self-service features for users on-site while giving IT the flexibility to connect to a public cloud for disaster recovery, off-site backup or to augment local capacity when demand spikes.

About Dell EMC VxRail

Powered by VMware vSAN and now offered on Dell PowerEdge platforms, Dell EMC VxRail 4.0 Appliances provide the flexibility to add capacity and performance on demand, enabling customers to easily extend use cases across the virtualized environment. VxRail delivers a hyperconverged solution using standard, enterprise-grade building blocks from known, trusted vendors and includes a single point of global 24x7 support for hardware and software. It's quickly and easily incorporated into existing VMware eco-systems, removing IT lifecycle complexity while simplifying deployment, administration and management. Fully integrated with vSphere, VxRail enables IT to expand on what is already known and proven in the environment, creating certainty in IT operations.

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 2017 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.