

Is Multi-Hypervisor Cost a Myth?

There are lots of exciting things happening in the technology industry. While multi-hypervisor strategies are not pervasive today, I see more and more production environments which are looking to leverage more than one hypervisor.

This raises the question of “is running more than one hypervisor too costly?”

If you are a hypervisor vendor, your obvious answer to this is yes, but let’s stop for a moment and consider the real detail around what defines the “cost” of running a hypervisor, thus the potential increase in cost to run more than one.

I checked the signs of the apocalypse and there wasn’t anything about hypervisors in there:



As a consumer of many hypervisors (Xen, KVM, vSphere, Azure, Hyper-V, and I also count Docker and LXC technically), I hope that it will reduce the risk of someone yelling FUD MONGER!! I can tell you with full certainty that this is a FUD-free post, and my opinions stem from years (far more than I may like to admit) of running production data centers and cloud platforms. I love all my hypervisors equally, and for different reasons ☐

With that, here are some thoughts around operational products and processes that will affect both tangible and intangible costs of operating your IT organization.

Backup

Does your backup product work across multiple platforms? Does it cost more to do so? Most likely, there will be a parity in cost to enable backup of a new guest in a new host. They are either agent-based backups which will grow as you do already, or perhaps they are host-based. Remember, that running multiple hypervisors doesn’t mean running twice as many underlying applications. You may see some duplication of workloads during transition into embracing new hypervisors, but the overlap is minimal.

In my opinion, a key reason to add another hypervisor is to create a scale-out application platform which lends itself to not need traditional node-by-node backups already. Stateless applications change the game already, so it’s not a 1:1 as you grow those environments.

Orchestration

Are you all-in with one orchestration platform that solely supports one hypervisor? Well, that could obviously be a challenge to think that bringing on new underlying hypervisor tools could create some cost of time and potential capital costs also. That said, take a look at your orchestration tool and I would bet that they are doing everything they can to be able to “orchestrate any workload”, right?

If the same people who are telling you that they have best-of-breed integration with their own hypervisor because they want to reduce costs for you, the reality is that they are selling you “stickiness”. By creating tight integration with a single platform, it makes that a sticky environment which ensures future re-up of licensing and support.

If you go outside of the big 2 (VMware and Microsoft) as platform behemoths, you also see the same thing in cloud. Obviously AWS configuration management tools are geared towards operating your AWS environment.

Take a look at the rest of the configuration management and orchestration tools and you will see something quite different. Just looking at CAPS (Chef, Ansible, Puppet, Salt) you see already that they are building towards being the multi-platform capable tool you need. When Microsoft Windows enables native OpenSSH in the platform, the agentless remote management capabilities open up quite a bit. How cool is that?

Conversion of Virtual Machines

This is called lifecycle management if you ask me. Most companies won't be looking to do a massive and sudden conversion from one hypervisor to another. They will run parallel environments and bring new workloads into the platform that makes the most sense. If you do choose to do a mass re-platforming of your production environment, that will be costly both from a capital and operational expense.

Cost of Ownership

Licensing is a capital and operational cost. Free hypervisors are like free kittens, because they may be free to start, but you have to keep feeding them. If you embrace open source alternatives to commercial hypervisors, you will probably engage with a partner or vendor to assist you.

In my mind, I consider this a wash. There are costs on either side, and I've been saying for a long time that VMware should make their hypervisor free (no, not the one that is already free, but the real full version) because they make money off of management tools. There is a reason that they stopped supporting ESXi under OpenStack because it created the ability to use VMware without paying for it and leveraging OpenStack projects for other capabilities. There was nothing innocent about that support and development effort going away. That's their choice, and why stop charging for something that tens of thousands of people are paying for?

Staffing

As a long-time architect and administrator of multi-hypervisor solutions in production and development environments, I actually find it rather offensive that vendors tell me that the learning curve is too steep to operate with confidence on more than one hypervisor. Let's hope they don't talk like that to their kids when they are coaching them through choosing a career path.

Trust me, I put OpenStack into an environment when nobody else on the team had comfort running linux. Guess what happens when you do that? They see the results and then will choose to staff up or train existing staff to bring the new skills in house. Do we really think that we are so locked in to a single product that we can't evolve our skills at all? And how do they think that we learned their platform to begin with? The reality is that we have already made the shift into and out of other platforms, so please stop with using this as the reason why we can't move.

Will it cost to do so? Yes. Will the benefit outweigh the costs? In my opinion, absolutely yes. By enabling your staff to operate and grow your IT platforms in a way that drives your business organization to be the best at what they do, then how can you say that it is "wrong" to bring a new hypervisor or development platform on board.

Why I Think Multi-Hypervisor is a Reality, and is Good

A year or so ago, the conversation around running multiple hypervisors would be one that ended in “there is no good reason to do so because of the difficulty in managing multiple platforms”. We have seen a lot happen in that time that has minimized the truth to that statement.

If you are experimenting with cloud platforms, you are already running multiple hypervisors. I actually had someone say that they didn’t think Xen was workable as a production platform, but they also run multiple applications in Amazon. It was quite joyful to let them know that Xen was powering their applications without them realizing it.

Hypervisors are effectively free, although VMware does cost more than other alternatives. Management of those hypervisors is generally not free, but in many cases can be zero capital cost or close to zero.

If you are migrating between clouds or on-premises platforms, there are tools that achieve this for you. Microsoft has their tool (obviously to migrate inwards), VMware has their tool (obviously to migrate inwards), Amazon has their tool (obviously to migrate inwards), along with other multi-hypervisor friendly tools such as Double-Take by Vision Solutions that can port machines between hypervisors. There are others, but I can speak to these four examples because I have directly experienced how they work. They work great. Almost surprisingly so, but lets face it, these have evolved over years and have been tested well across the years with more consistency and success as hypervisor vendors made inward migration simpler.

Why Run Multi-Hypervisor Environments?

Easy answer in my opinion, which is to enable the best-of-breed platform for your business to run on. If you are experimenting with scale-out application architectures, why wouldn’t you explore doing so on a platform that provides that capability with less cost. Running alternative hypervisors to align with the requirements of your workload.

We no longer switch entirely from vendor to vendor when we add products to the IT portfolio. You probably run more than one orchestration tool today. You probably run more than one development platform today.

If anything, consider what it means to your business to be more agile for you customer. Plus, think of that as you go through your week. You probably shop at more than one chain of stores, and stay at more than one chain of hotels, and you’ve probably owned more than one brand of car. All of those decisions are for one simple reason: choose the best solution for the best result. There are many results we seek, and many solutions that provide them.

The one-stop shop hypervisor is a myth. That’s why we laugh every time you hear “single pane of glass”, because we all acknowledge that best tool of best job is what we really need. Let’s just cool down the fear campaigns that you have chosen your hypervisor once, and you’re not allowed or supposed to look elsewhere. In ten years, you’ll look back and laugh at what you are doing in production today, the same way we look back at physical servers as the “best and only solution” which was how they were sold to us back then too.

Comments are welcome, and I’m sure they could get interesting ☐