

The Six Factors that Determine Your Storage TCO

When you're in a buying cycle, calculating the cost of different storage options can be overwhelming. And so, more often than not, it's easiest to default to a well-understood metric that can be applied across solutions—such as cost-per-gigabyte.

But that represents just a fraction of the total cost of ownership. Consider that over the past two decades total capital expenditure has grown 3x, and in the same time total operating expenditure has grown 8x¹. Why? Hypervisors have driven new levels of application density, but also introduced far greater management complexity.

So, how much will your storage really cost you? Here are the six factors you need to consider and suggestions for how to think through—and even measure—each one.



Factor 1: Capital expense

Capital expense is the easiest factor to measure, compare and communicate—and the most commonly used metric is cost-per-gigabyte.

Part of the reason is the attention that vendors place on cost-per-gigabyte; storage marketing too often touts this factor in the absence of real differentiation. For example, all-flash vendors are quick to highlight that their cost-per-gigabyte is rapidly approaching the cost of disk—encouraging comparisons on this metric, not on which storage platform is best suited to the environment.

But if you're going to boil capital expense down to a single number, is cost-per-gigabyte still relevant? Consider that:

- Performance often runs out before capacity—forcing organizations to over-provision storage to maintain a performance buffer
- Continuous hardware and software upgrades that optimize for this metric can be highly disruptive to business operation (and thus impose high cost)
- Now that nearly 80% of applications have been virtualized, the premium isn't effective capacity, it's the density with which you can fit VMs into available space and performance resources

Recommended metric: Cost-per-VM. Storage does not create business value—applications do. Since virtual machines are the value driver, they should be the measuring stick for capital expenses. Now, some vendors are wildly efficient while others are not—and that makes for a far more telling comparison.

Tip: If a vendor can't tell you how many VMs can fit on their storage system, that should send up a warning flare right away—probably an indication that they are ill-suited to support your virtualized applications, period.

¹ IDC, Trends and Technologies Impacting Today's Storage Infrastructure Market: September 2015.



Factor 2: People

As mentioned above, operating expenses are growing at a far faster rate than capital expenses and people costs are the number one driver. When it comes to storage, there are two dimensions: skillset and management effort.

Most storage systems require deep storage expertise to manage, and storage PhDs are highly sought after and well compensated. Since the cost of catastrophic failure is so high, having storage expertise on your team is often a necessity.

But then occupying their time with low-value tweaking, tuning and troubleshooting is a waste of resources. If you can afford storage expertise, you want their skills focused on larger strategic issues.

To tame TCO you need storage that anyone on your data center staff can manage, and do so in a fraction of their time. That's rarely possible with conventional architectures built on LUNs and volumes—the RAID, striping, queue depths, etc. that they require are not a shared language across the data center. Outside the storage admin, the rest of the team is thinking and acting in virtual machines and applications. So, be careful to understand how your storage is designed—the building blocks of its operating system are a key indicator of its ease of use.

Recommended metric: Management effort. This can be hard to quantify, so break it down into the hours required to install & configure, manage daily, troubleshoot and more. Be sure to ask for references when investigating storage vendors, and ask those peers about the management effort they invest in storage.



Factor 3: Ecosystem Integration

Some storage solutions arrive highly integrated with the ecosystem that matters to your business, some will require substantial development investment on your part. Here's where you can anticipate and avoid future costs.

One example is hypervisor support. Is the storage built specifically for VMware, or does it work equally well with Microsoft, Citrix, Red Hat and OpenStack? Two-in-three companies run multiple hypervisors—if you are running multiple, concurrent hypervisors (or plan to) will you need separate nodes specific to each hypervisor

Beyond hypervisors, there is the next degree of integration—that includes (in the VMware ecosystem) VAAI, VCAI, VVOL, SRM and more, or with Microsoft SCVMM, SMB, etc. To return to the earlier point about architecture, be forewarned that LUN and volume-based storage do not have operating systems that integrate easily with the virtualization ecosystem.

Recommended metric: Development Cost. Cost this out up front—what integrations will you require? Are they available out of the box, or will they require custom work? If it's the latter, factor in the costs to your timeline and budget.

Tip: Look to certifications from major hypervisors as a proxy for integration breadth and depth. For example, VMware vSphere certification, Microsoft Gold level data center competency and/or OpenStack ecosystem certification.



Factor 4: Scale

Now it starts to get harder, but it's worth it. Every vendor under the sun touts their ability to scale up, out and sideways—even floating terms like web-scale. How do you quantify this?

Consider these two dimensions:

1) Will additional workloads require additional nodes?

2) Do components scale independently or all together?

The first question presumes the number of workloads you need to store over time is growing (a pretty safe assumption). If you've started with virtualized servers, but will soon be virtualizing desktops, can you roll those out on the same system, or will you need to provision more storage? Look for proof points (via customer references and POCs) that a vendor can support multiple workloads on one system and ensure that each gets its fair share of capacity and performance resources.

The second question addresses whether you need a standalone or (hyper)converged solution. (Hyper)converged has a value prop clearly built around simplicity, and they're quick to label their technology as web-scale. But if you go that route, it's because looking ahead you want to scale compute and storage at the same rate. History suggests these two components rarely scale in linear fashion, and so standalone offers greater flexibility and is likely to be more cost effective.

Recommended metric: Historical rate of scale. This is something you need to investigate internally. Have compute and storage needs tended to scale at the same rate or not? That should help you determine the right path forward. You can model it out if time allows to see whether (hyper)converged or standalone offer greater economies.



Factor 5: Business Agility

You should weight this factor if your storage footprint will cover DevOps / Test and Development. Those teams are under pressure to reduce cycle times—to design, build and test products and services faster. Can your storage be an enabler?

In our experience, there are three ways that storage supports business agility. First, it can be so easy to manage that individual teams can manage their own footprint rather than take dependencies on storage admins. They should be able to spin up and tear down VMs at their pace, and/or troubleshoot issues without getting stuck in a queue of trouble tickets.

Speaking of getting stuck in a queue... second, if the storage is shared, it has to be able to isolate the performance of DevOps applications from other virtual machines; that's because DevOps may be running time sensitive exercises that cannot be queued behind IO intensive database or analytics workloads.

Finally, at any point in time the DevOps team will have many developers working on "child" VMs linked to a common "parent" VM. It is critical to keep those "child" VMs up to date as changes are made to the "parent". Either your storage makes it dead simple to push out updates and ensure every developer is working with the most current version, or it will require many cumbersome steps.

Recommended metric: Cycle times. If the DevOps team is a stakeholder, make sure they have a seat at the table in the purchase decision. They should think about the three points above, and poke at how storage will help them reduce cycle times. Storage is often seen as a cost center, but in this case it can be a critical business (and revenue) driver.



Factor 6: Opportunity Cost

Last but not least is opportunity cost—in many ways, a summary of the above. For example, if you are spending more on capital expenses for storage that is low cost-per-gigabyte but high cost-per-VM as your virtual footprint grows... how else could that money be used?

Or in the case of talent, what is the opportunity cost of tasking your highly skilled storage administrator with shuffling VMs between LUNs and addressing trouble tickets day after day? That individual could be spending their time on projects with major strategic impact, like planning for private cloud or establishing ITaaS policies.

And consider business agility. What is the opportunity cost of the DevOps team losing days, weeks or even months to storage bottlenecks? All these elements of opportunity cost are hard to quantify, but they need to be part of the discussion with stakeholders.

Recommended metric: Gut check. Will your storage simplify your environment? Can it offer predictable performance in an environment that's easy for any member of the IT team to manage? Looking ahead, choose the storage that will help you realize your data center vision, whether that's virtualization or private cloud.

Tip: The storage purchase decision is a chance to be a change agent. Only one person can be first to introduce a new technology. Just like TCO is a calculation, so is the cost-benefit of betting on a highly differentiated platform.

Hopefully the above six factors will play into your next storage decision. What you'll find is that conventional storage vendors (and that includes newcomers that continue to rely on LUN and volume-based architectures) will be incrementally different across these six dimensions; a few percentage points here and there. Vendors with an entirely different architecture—one designed to be VM-aware—can offer more radically differentiated capabilities and value. For example, a third party study of Tintri VM-aware storage customers showed that on average those organizations sped performance 6x, shrank footprint 4x and managed storage in 98% less time as compared to conventional storage. And if those numbers seem unbelievable to you, make us prove them. We'd be happy to work through Tintri's impact across the above six factors and help you see storage differently.



Global HQ

303 Ravendale Dr.
Mountain View, CA 94043
United States
+1 650-810-8200
info@tintri.com

EMEA HQ

Fountain House 10th Fl
130 Fenchurch Street
London EC3M 5DJ
+44 (0) 203 053 0853
emea@tintri.com

APAC HQ

9 Temasek Boulevard
Suntec Tower 2, #09-01
Singapore 038989
+65 6407 1359
apac@tintri.com

Japan HQ

Level 6, Kishimoto Building
2-2-1 Marunouchi, Chiyoda-ku,
Tokyo 100-0005 Japan
+81 (3) 6213-5400
info.japan@tintri.com