5 Things Storage Vendors Are Not Telling You

Executive summary

Storage is currently the most expensive and most complex piece in the data center and integral part of any Cloud service - be it public, private or hybrid cloud. Customers usually do not have enough information when buying storage solutions and are often deceived by sales tactics and marketing propaganda. Terms like "software-defined storage", "storage virtualization", "server SAN" and many more claim to be the cure for any storage problem.

This whitepaper is intended for the business leader of any company involved in data storage buying decisions. It gives practical advice on what to look for when buying a storage solution and what caveats to avoid.

This paper also tries to explain in simple terms some of the jargon and practices in the storage industry. It is an attempt to demystify magic marketing statements and help the buyer to make a wise and educated choice.

The 5 Things Storage Vendors Are Not Telling You

1. You might get a huge discount.

Traditional storage vendors do anything to get that initial contract. For a "first time deal", customers might get 30-50% off-the list price of a storage array. If you are a hard negotiator or a strategic account then discounts might go even in the 80% ranges. Have in mind that usually price reductions are made on the hardware only.

Selected models from big-name storage vendors	Total list price	Total street price	Average discount
Including EMC, NetApp, Dell, Nutanix	$ 1 286 400	$ 629 079	-51%

Source: StorPool competitive intelligence.

A legacy storage vendor (think two and three letter vendors) would rarely give you discount on the support or the software packages. And it is the services and software that deliver vital functionality and turn a piece of iron into something that actually delivers fast and reliable storage. There might be a discount on support or storage software, but you can expect it to be much lower. Additionally professional services that are always needed in order to design, evaluate and deploy a storage system are not included and may add a significant amount to the total bill.

Beware. This is a tactic used "to get the foot in the door".

2. ...but there is no "free lunch"

Getting that storage array at a solid discount might initially sound like a good deal. But it rarely is - when renewal time comes or when you need to expand the storage solution... you will pay the price - expect no discount then. This is, of course, a form of a vendor lock in. You will always be dependent to some extent on your vendor or technology of choice. However, some vendors lock you in much more than others. It is much easier to change a storage solution that is software + hardware. You can decouple and change the two pieces independently. And it is much easier to change a piece of standard hardware than a piece of specialized vendor specific hardware.

Tip:
Do not base your business case or ROI calculations on the initial discounts you get from your hardware storage vendor. They might give you a discount the next time, but this is unlikely. And even in the case that they do give you a discount - it will be much lower than the first time.

Companies often turn to StorPool after their 3 year contract with a traditional storage vendor nears expiration - they are shocked by the new price quotes. Another recurring pattern is of a company contacting us when they need expansion of their current storage solution. They are again quoted prices that change the Total Cost of their solution to the point where their ROI vaporizes.

3. The Margin is built in the drive, not the box

Hardware storage vendors use the well known freebie or "razor and blades" pricing model, where the margin is wrapped in the drive price and not the chassis. So the big box comes at near cost but the drives you need are vendor specific and typically cost 3 to 15+ (!) times more than basically the same or alternative drive bought from a standard vendor.

Comparison table: drive list prices from different vendors	Dell	EMC	NetApp	Market price
Hard Disk: 1TB 7.2K RPM SAS	$262	$1 890	$597 *	$90
Hard Disk: 2TB 7.2K RPM SAS	$411	$2 990	$1 151	$140
Hard Disk: 3TB 7.2K RPM SAS	$532	$4 215	$1 856	$190
SSD: 400GB	$4 259	$7 598	$1 718	$509
SSD: 800GB	$2 484	$14 435	$3 251	$739
Average price inflation, times	**4.1**	**19.8**	**6.5**	**1.0**

Sources:
Dell: http://configure.us.dell.com/dellstore/config.aspx?oc=brct132&model_id=powervault-md3860i&c=us&l=en&s=bsd&cs=04 ; **EMC:** http://www.peppm.org/Products/emc/price.pdf & http://www.emc.com/sales/stateoffl/florida-price-list-2014-05.pdf ; **NetApp:** http://www.peppm.org/Products/netapp/price.pdf ; **Market price:** http://www.newegg.com & http://www.intel.com

* no list price for 1TB 7.2k NetApp drive. Estimated as average of $/TB of the 2TB and 3TB drives
Prices as of 12 Nov 2014

Fact:
There are only 3 companies in the world producing hard disks today, regardless of what the logo sticker on top of the drive says! These companies are Seagate, Toshiba and WD. And while the disks might come from 3 original vendors, the drives you buy from your storage vendor are with flashed firmware and cannot work in another array (or a standard server). Also, the storage box from this vendor would work only with the respective vendor's drives, so you cannot buy the standard drive and put it in the storage array either. You have to get these 3-10x more expensive drives.

Besides, the price you pay for a drive when you buy it in the box is different from the price of an additional, upgrade drive. And you guessed it - the latter is much higher. Here is an example for you, taken from EMC. Notice upgrade drives (denoted with "***UPG***") are exactly ***40%*** more expensive:

Item Name	Product Number	Description	Subcategory	MSRP	% increase
V MAX 10K DRIVE	NF4104501C	VMAX 4G 10K450GB DRVSi	Enterprise Tier, Managed	$2 470.00	
V MAX 10K DRIVE UPG	**NF4104501CTU**	**VMAX 4G 10K450GB DRVSi UPG TAA**	**Enterprise Tier, Managed**	**$3 460.00**	**40%**
V MAX 10K DRIVE	NF4106001B	VMAX 4G 10K600GB DRV	Enterprise Tier, Managed	$3 000.00	
V MAX 10K DRIVE UPG	**NF4106001BTU**	**VMAX 4G 10K600GB DRV UPG TAA**	**Enterprise Tier, Managed**	**$4 200.00**	**40%**
V MAX FLASH DRIVE	NF4F24001B	VMAX 4G FLASH 400GBS DRV	Enterprise Tier, Managed	$23 435.00	
V MAX FLASH DRIVE UPG	**NF4F24001BU**	**VMAX 4G FLASH 400GBS DRV UPG**	**Enterprise Tier, Managed**	**$32 810.00**	**40%**

Source: http://www.emc.com/sales/stateoffl/florida-price-list-2014-05.pdf

4. Price does not tell you enough anymore

Storage is already complex and is getting more and more complex. The new software features work magic. One can dramatically change the behavior of a storage system with every new technology and feature added. As with anything in life, there is a trade-off, though. It is getting tougher to estimate the actual behavior of a storage system.

The typical way of buying storage used to be straight forward: "give me a solution that can deliver 20 TB of raw capacity and 20,000 IOPS of performance". Not anymore - with software features like thin provisioning, snapshots & clones, caching, tiering and so on, one can considerably change the output parameters of the storage system. Furthermore, the impact of each feature will depend on the customer's use case - the type of their data, the pattern of their workload. It is getting progressively challenging to predict the impact of a feature on the actual results a storage system will deliver. Now you need to run a PoC or a pre-deployment test just to get a good estimate of what your solution will actually deliver.

Example:
1 TB is not 1 TB any more - intelligent data reduction storage features such as snapshots, clones, thin provisioning, deduplication, in-flight compression will shrink the original 1 TB to… unknown amount of space. This gain is use-case specific - depending on your particular data and its usage.

Another trend we notice is the marketing of big vendors that pushes particular buzz words or a feature as a solution for every problem. Maybe the most advertised word we come across is "deduplication". While "dedup" it is a great feature, it has been over-emphasized mainly by "all-SSD" vendors as The Solution for reducing the physical storage footprint of your data (i.e. raw to usable data). However, there are a number of

other features that achieve the same goal, some of which have much bigger impact - for example thin provisioning, snapshots, compression.

Here is a particular example from a customer in the hosting industry. The use case is a VPS service, running StorPool Storage:

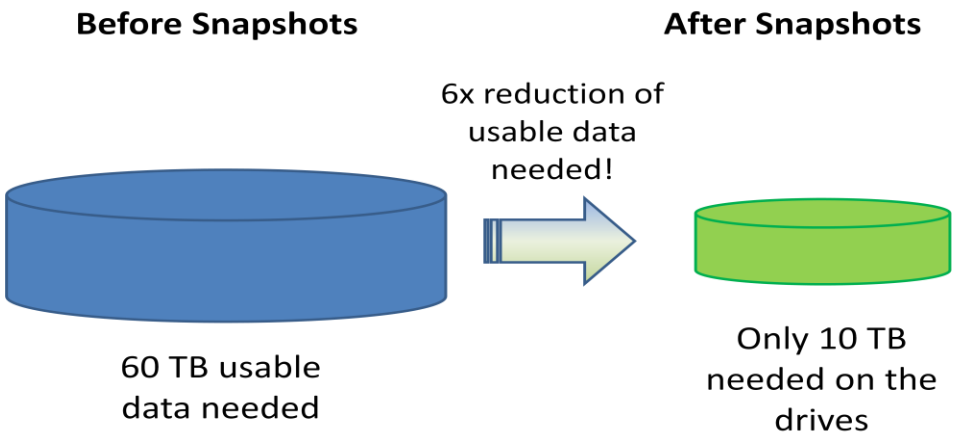

Note: While similar results can also be achieved with deduplication, it uses significant amounts of CPU and RAM, while snapshots do not stress the system.

In resume:
Educated customers should take decisions on the actual benefit/results they would get and they do not insist on a particular technology to deliver these results. As we stated above - the data and the use case have significant impact on the technology that can deliver the most benefits. Customers do not need features they need solutions to real business problems. It usually comes down to increasing performance / capacity or bringing down the total cost of a solution. In many cases - all at the same time!

5. There is nothing special about that storage box

We see many customers who believe that a storage array is something "special". This is not true. A storage box is just a regular server on the inside, with a very special cover design and logo on the outside. And it is not an exceptionally powerful one, as well. Besides the custom box and design, inside the box you find all the components of a standard server - CPU, RAM, disk controllers, network interfaces and of course, drives. Nothing so special.

Example: resume of specifications of EMC VNX 5100 storage array vs. standard server:

Selected Specifications	EMC VNX 5100	Standard server
Protocol	FC (Fibre Channel)	Standard Ethernet network
Storage Type	specialized single purpose storage box - SAN (Storage Area Network)	standard server - use for compute, besides storage
Drive Type	Flash SSD, SAS, NL-SAS	All of EMC's ***plus*** SATA drives
Capacity Optimization	Compression Thin Provisioning	Can have ANY storage feature, provided by intelligent storage software running on the server
Management	EMC Unisphere	simple interface
CPU	Intel Xeon Dual Core 5600	Intel Xeon Processor E5-26xx v2/v3
RAM	8GB	up to 512GB per node
Solution For	Storage Consolidation, Remote Offices, Virtualization	Can cover wide range of cases, depending on the storage software
Product Type	Storage System	Standard server - both storage AND compute

The components in this box are well combined and tested together, but this is not something that is particularly special or hard to do. All software-defined storage vendors provide reference configurations or pre-selected commodity servers that can be used. These configurations can then be used as a standard building block of a scalable storage system.

About StorPool

StorPool is a leading vendor of distributed storage software. The software turns ***standard servers*** into ***high-performance storage system*** and replaces specialized storage arrays (SAN). StorPool aggregates the capacity AND performance of drives distributed in many servers and provides one aggregate shared pool of storage. This technology is far superior to SAN in both technical and business terms (scalability, performance, flexibility, reliability, efficiency, cost / TCO).

Compared to other storage software solutions StorPool delivers very high performance, exceptional efficiency, scalability and simplicity. To achieve this, we have re-engineered the storage stack under the block device from scratch. StorPool can reduce TCO several times and can be licensed on pay-as-you-go (OpEx) or perpetual (CapEx) basis.

Contacts

+1 415 670 9320
info@storpool.com
www.storpool.com

twitter.com/storpool
facebook.com/StorPool
linkedin.com/company/storpool