

# Hyper-Converged Infrastructures are not Storage Arrays



From the title of this post you may be surprised to find that I'm a huge fan of hyper-converged infrastructures (HCI). Frankly, they fill a market need that for years customers have either been struggling to address with storage arrays or worse ignoring all together. HCI is a perfect solution for retail locations, bank branches, militarized and mobile emergency response vehicles, etc; all of which require data center services on a scale that is too small to justify the cost of a storage array.

The HCI market has the potential to be huge; however, will it replace enterprise storage arrays? Simply put, yes and no.

I ask you to read on before you or another reader posts the comment stating I'm a "hater" or the "Law of Instrument" (aka - if you only have a hammer...).

## **Hyper-Converged Infrastructure Arrays 101**

I believe HCI should be viewed as a form of converged infrastructure, one that is comprised solely of server hardware (and top of rack network switches). This software-defined model leverages a virtualization layer to provide shared data services comprised of local storage, a distributed file system, data mirroring mechanisms and Ethernet networking.

All persistent storage, this includes disk and flash, inherently have media errors. These imperfections require the storage system to provide data protection in the form of RAID, mirroring, erasure encoding, etc, to ensure data availability in the of a component failure. The HCI architecture implements data mirroring on a per-virtual disk basis. HCI vendors do not use uptime like storage vendors do (i.e. five nine's availability), instead they speak in terms of Failures To Tolerate (FTT). I think it is fair to compare the setting of FTT=1 to RAID 5 on an array (four nine's of availability) and FTT=2 to RAID 6 (five nine's of availability).

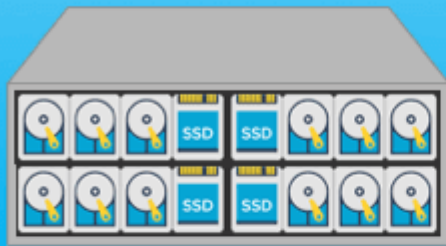
### **Low Storage Arrays Utilization Challenges HCI as a Storage Platform**

The storage industry refers to the storage model in an HCI as a shared nothing architecture. This definition loosely means the architecture is distributed and each node is independent and self-sufficient. The HCI platform provides some truly unique and beneficial capabilities in areas like granularity of deployment and scaling, per-VM data availability and service policies, and consolidation of operations; however the trade-off of HCI is low storage utilization due to the shared-nothing architecture.

Consider the example below (based on the EMC VSPEX BLUE).

# Hyper-Converged

is not a storage platform



16TBs

raw


1.2TB HDD 1.2TB HDD 1.2TB HDD 400GB SSD



4TBs

raw per node

23.9%  
usable per node



Rebuild Reserve Data Protection Snap Reserve VSAN Cache

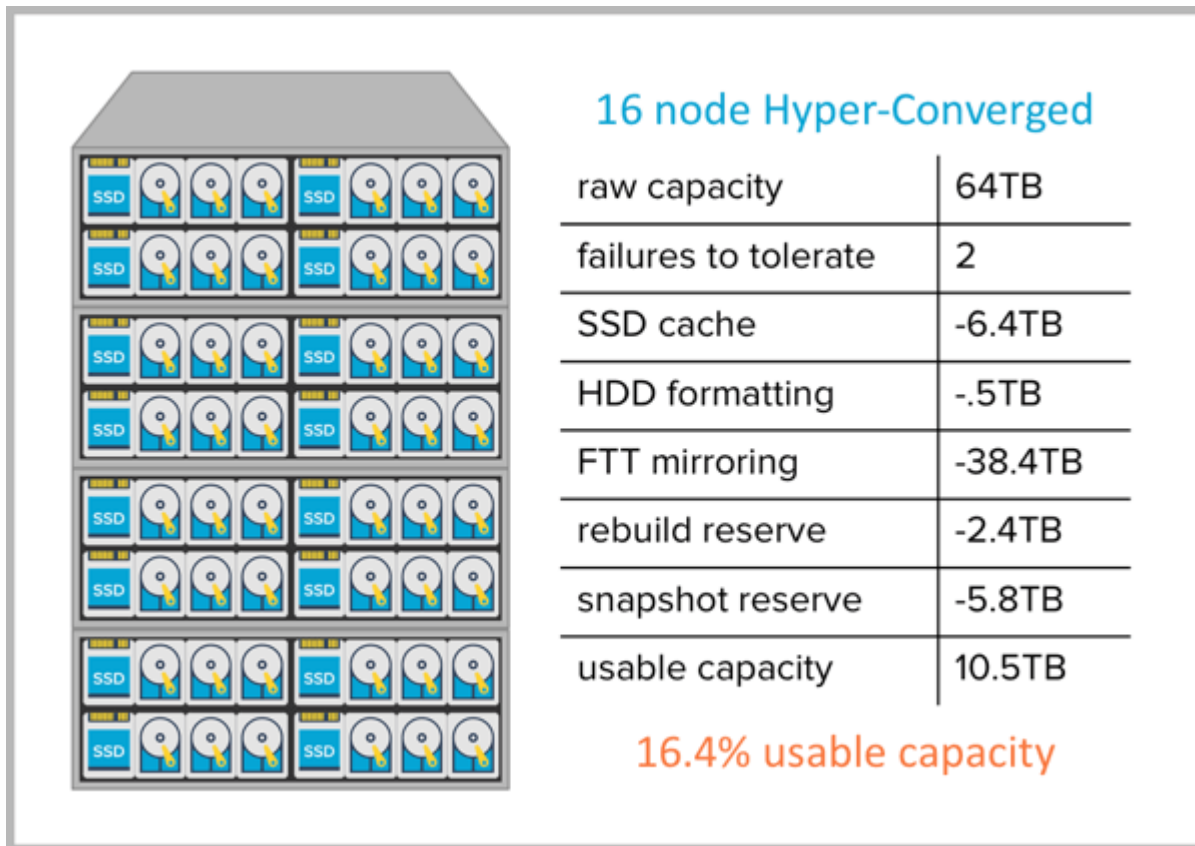
3.8TBs

usable




The largest area of capacity loss in the HCI model is the [data mirroring and node rebuild reserve](#) required to support the number of faults to tolerate. The overheads of both vary as the cluster scales. For example, set FFT=1 on a four node HCI and usable capacity is cut in half and one must reserve 25% of the capacity to support the rebuild of the data in the event of a single node failure.

Again, consider the example below (based on the EMC VSPEX BLUE).



As one adds nodes to an HCI Arrays deployment, the overhead of the node rebuild reserve reduces, which is good; however, the increase in raw storage capacity also increases the number of media errors. With such capacity increases, one needs to revisit resiliency as the setting of FTT=1 does not protect against media errors during a disk/node failure and the data rebuild process. Hopefully this example helps express why I view FTT=1 in the same vein as RAID 5 - appropriate for small deployments but fundamentally lacking for deployments at scale.



*It would be remiss of me not to mention that HCI Arrays architectures vary from vendor to vendor and some are more efficient than the platform I used as the example in this post. Consider the larger perspective: shared-nothing architectures trade off hardware efficiency in order to deliver value in other areas. Those that offer means to reclaim some of the loss through data reduction Arrays technologies.*

### Hyper-Converged Infrastructures Help Customers

I don't think we should include hyper-converged infrastructures in storage infrastructure discussions. Maybe the marketers who gave us product names like 'Virtual SAN' and marketing messages like 'No SAN'

should be commended for injecting HCI into the storage conversation... I guess they did their jobs rather well. With that said, it falls on the shoulder of the technically inclined to understand and position the advantages and trade-offs of HCI based on deployment requirements.

Customers have significant pain points delivering highly available, performant, and cost-effective storage. Until recently, storage platforms have not advanced in lockstep with the advancement of compute and networking, which has led many to seek innovation and new storage technologies.

It is my opinion that HCI is an ideal means to provide storage services in environments where the inclusion of a low-end storage array is simply a financial challenge. The requirements of this market are very different than what is expected from the core storage infrastructure within a data center - which is the crux of this post. I view all-flash array architectures as more capable and thus appropriate for next-generation storage infrastructure in part as leading AFA platforms deliver effective storage capacity that is greater than the raw capacity of the Arrays platform. This benefit directly translates into cost savings. Attempting to compare the storage utilization of an HCI vs an AFA should have the same impact of when you consider the fuel efficiency when purchasing an automobile. Are your purchase decisions influenced by similar vehicles when one gets 5 miles to the gallon of gasoline and the other 50?

I understand my views sound like a vendor bias - this is a fair critique - but I'm not alone in my view. Web-scale architectures 2.0, aka *data center disaggregation*, is composed of racks compute and shared all-flash Arrays storage. Sounds like a new take on an old model, but let's save the data center disaggregation discussion for a future post.

I fully expect HCI to displace a portion of the low-end storage array market. Face it, disruption is good for customers. Kudos to the hyper-converged infrastructure vendors; they have helped customers with a long-standing challenge and as a result they have created a net-new market. I'm excited to see where these platforms may take us.

What are your thoughts on HCI? Where does it fit well today and where may it go in the future? All comments welcome.