



Delivering Predictable Performance for Enterprise IT At-Scale

Intelligent Infrastructure Succeeds Where Conventional and HCI Architectures Fail

Hyperconverged infrastructure (HCI) brings compute and storage together in a single chassis. The idea behind HCI solutions from vendors like Nutanix and Simplivity is that IT would be freed from integrating point solutions and to provide scalable, guaranteed performance with lower risk.

The thinking is that customers would be able to provision resources quickly, scale easily, and reduce costs significantly because of the integrated technology layers and features in the converged platform. The reality is more complicated—it's difficult for HCI to deliver scalability, simplicity, and cost advantages without sacrificing performance. And while HCI may seem simpler, the tightly coupled architecture also makes it really difficult to troubleshoot performance issues because everything is layered together on each node.

Although HCI has gotten a lot of attention in the industry press, HCI architectures with conventional storage lag behind best-of-breed external storage systems—both hybrid flash and all-flash offerings—across a number of important performance metrics:

- Latency
- IOPS (especially compared to all-flash systems)
- Predictability

As a result, enterprise data centers have quickly adopted all-flash storage as a means of delivering the IO performance needed to power applications of all kinds—especially analytics and new mobile and customer-facing applications.

The main reasons for choosing all-flash storage are:

- Significant reductions in the latency of each IO operation
- Massive increases in total IOPS
- More predictable performance for every IO

Challenges

- Most infrastructure vendors can't deliver predictable performance to smoothly scale enterprise workloads
- IOPS potential of flash SSDs can be wasted
- Latencies are often unpredictable under load
- Activating new features can further drain performance

Tintri VMstore effortlessly delivers all three of these results. In fact, Tintri's Intelligent Infrastructure architecture goes even further—it automatically assigns every virtual machine, database and container to its own IO lane to eliminate any conflict over resources. That makes it simple to set minimum and maximum quality of service (QoS) levels for individual virtual machines, and autonomously guarantee application performance for each one. If application or database performance is a primary consideration, you'll want to evaluate all options carefully and think twice before choosing an HCI solution.

Predictability

For many applications, performance that's predictable is just as important as low latency or raw IOPS, but it's often much harder to deliver. Predictable performance is even harder to achieve with HCI solutions that combine storage and compute and run both sets of activities on the same nodes. Why? Let's look at virtualized environments, where storage software usually runs inside a VM as a virtual appliance at the guest layer. In these cases, each IO operation flows through the hypervisor's CPU scheduler four times: twice for the IO, twice for IO acknowledgment.

This might not create a problem when system utilization is low, but when CPU resources are shared – as they are with HCI – it frequently results in a major bottleneck as utilization becomes moderate or heavy. One solution is to apply CPU reservations, but per-VM reservations introduce a new set of challenges that cascade to the point where cluster-wide HA-failover policies can be affected. And CPU reservations do not guarantee that the virtual appliance will have instant access to the CPU. If another vCPU is scheduled and running, it must be allowed to finish its operation, which can cause IO delays within the virtual appliance.

With or without CPU reservations, the result is latency that's less predictable, with unexpected spikes when a node or cluster becomes busy. Applications may experience latency that varies widely from one IO to the next, which can be disastrous for those that are particularly latency-sensitive. This problem is further exacerbated in an enterprise cloud environment where an organization may be managing thousands of virtual machines and/or containers. Despite claims of delivering "web-scale" capacity, HCI solutions are rarely able to meet performance expectations at scale – on-premises or in the cloud.

Latency

The latency of IO operations on conventional HCI storage implementations suffers in comparison to external storage systems. For example, when data is mirrored or copied to other nodes for data protection, multiple copies of each data block are stored across the network – which has a direct impact on latency. Some HCI vendors support erasure coding as an alternative to mirroring. But this technology, while offering more "nines" of availability, also comes with a high penalty on both performance and latency. There are workaround solutions that support post-process erasure coding, but for cold data only – which is not useful for dynamic enterprise workloads.

Both mirroring and erasure coding affect write latencies and may affect read latencies as well. Enterprise Strategy Group (ESG) recently compared the performance of several storage platforms under various conditions. The best latency achieved by any HCI solution was around 5ms, which is far slower than best-of-breed all-flash systems.

Apart from mirroring and erasure coding, activities such as VMware vMotion, HA events, maintenance on nodes and node failures can result in increased latency for workloads due to the noisy, resource-intensive nature of vMotion and HA events and the reduction in total available resources.

IOPS Performance

The IOPS performance that storage can deliver, especially all-flash storage, correlates directly to how much CPU power is available. Most standalone all-flash arrays use 28-40 cores per controller for 13-24 SSDs. (Though some arrays scale up even higher, this progression can negatively affect the IO density of all-flash media and, as a result, performance predictability.)

Some HCI vendor implementations limit the amount of CPU available for storage. Up to 8 vCPUs or 20% of available CPU are the typical limits. This is not enough horsepower to deliver full performance from the flash drives on each node (6-24), resulting in a lot of underutilized, wasted flash IOPS.

Other HCI vendors enable you to increase the amount of CPU dedicated to storage, but this can have a big impact on licensing costs. You don't want to get stuck paying for expensive hypervisor, SQL Server, and/or Oracle licenses on CPUs that are ultimately running storage functions.

Buyer Beware

Many storage vendors promise a range of advanced capabilities and associated benefits, but as with any major infrastructure decision, it's important that the buyer beware. Enterprise IT teams want—and in many cases need—to use all the latest functionality the infrastructure can deliver.

The reality with many solutions is that enabling new storage functionality can increase resource utilization beyond acceptable levels. As features like snapshots, replication, deduplication, compression, and so on are activated, you're faced with a choice: add more hardware or sacrifice the predictability and performance of your infrastructure. Making these trade-offs has become an almost daily event for many administrators.

Enabling data reduction features, for example, causes HCI platforms to consume even more CPU resources. That's why data reduction is optional on many HCI implementations.

The Tintri Takeaway

Intelligent Workload Mobility. When it comes to performance delivery and management, HCI solutions have a range of limitations. The additional convergence and sharing of storage and compute resources creates unique bottlenecks that affect latency, IOPS and performance predictability. HCI can provide better workload mobility compared to conventional standalone storage systems. Yes, you can move things to different nodes in your cluster, but you don't have the intelligence to move things around without causing a performance impact to other workloads. That's something that's uniquely available with Tintri VMstore.

Best of Both Worlds. Tintri Intelligent Infrastructure takes flash-based storage to another dimension, with unparalleled performance advantages versus HCI around visibility, analytics and quality of service. Tintri has taken a different approach than HCI and other array vendors by providing the best of both designs: all the simplicity of HCI and all the scale and performance of all-flash. Unlike HCI, we focus on optimizing storage while supporting tight, simple integration across multiple best-of-breed compute and networking platforms. Our VMstore architecture was built around modules to also enable seamless performance at scale with best-of-breed external infrastructure.

Auto-QoS for Predictable Performance. Our autonomous QoS capabilities ensure that each workload – whether that's an application, database, or even an individual virtual desktop – always gets the resources it needs when it needs it so that performance and latency are never an issue. This is true whether you've got tens of VMs or tens of thousands; whether you're using advanced features or not. You always get predictable performance, with latency that's consistently less than 1ms for each workload.

Non-disruptive, Elevated User Experience. And because we offer both broad and granular visibility and control from a single screen, you get intelligence that just isn't available with HCI – or any other solution. You can instantly figure out exactly where problems reside for any VM or app across compute, network and storage layers to avoid the slightest disruption and maintain an outstanding user experience.

Experience the difference. Experience Tintri Intelligent Infrastructure.

