

## CBTS OnX shows: Intel Optane



### How Intel Optane SSDs more than double SDS cache performance

Data center storage infrastructures must meet stringent requirements for availability, reliability, and performance. That's why software-defined storage (SDS) solutions such as VMware's Virtual SAN (vSAN) continue to grow at a rapid pace, delivering major value in configuration flexibility and scale.

Most SDS solutions employ a two-tier storage configuration: the back-end storage tier holds stored data, while the caching tier handles active data entering and leaving the SDS node. Reliability is what's most important in the storage tier, whereas caching is where the heavy lifting gets done, requiring higher levels of performance to continually serve all active data for the storage node.

In a standard SDS node configuration, traditional HDDs are sufficient for the storage tier, while the caching tier consists of higher speed SSDs. This is fine for moderate workloads such as branch office applications and basic VDI deployments, but it's less suited for high-demand environments.

#### Accelerating the cache tier with Intel Optane SSDs

Now, Intel has introduced a new kind of data center storage medium, the Intel Optane SSD. This technology provides an order of magnitude improvement over older SSDs in performance, availability, and reliability — so much improvement that it's begun to impact the design of all types of storage infrastructures, including SDS solutions. But OnX wanted to know: how much difference could the new Intel Optane SSD make in SDS cache performance?

In a recent proof of concept test in our Client Briefing Center, we measured the performance difference between two VMware vSAN configurations, with and without the Intel Optane SSD. The test results were impressive, revealing a clear advantage for accelerating cache speeds to attain a new class of SDS performance.

We ran a series of test workloads using VMware HCIbench to measure performance in two test beds. Each consisted of an HPE ProLiant XL450 Gen9 (two 3-node clusters), two Intel Xeon processors E5-2650 v4, eight 1.2Tb spinning hard drives, 128GB RAM, and 10Gb networking, running vSphere 6.5, vSAN 6.6 and vCenter 6.5. The test beds differed only in the cache media used: one configuration consisted of a standard non-Intel SSD (450GB per node) for the cache tier, while the second one utilized an Intel Optane SSD DC P4800X (375GB per node).

## See how our test findings can apply to your SDS needs

Contact CBTS OnX to learn more about the results from this latest proof of concept and how we help our clients evaluate and optimally deploy technology to support their business needs.

## One simple change generated more than 2X the performance

The test results from the Intel Optane SSD were so impressive, we had to go back and double-check to verify the performance data. By simply swapping out the standard SSD with the Intel Optane SSD DC P4800X for the cache tier, we increased IOPs by more than 2.6X and decreased overall latency by over 60%. This faster cache is significant because, in addition to providing faster overall speed and lower latency at the SDS node, it also supports a larger vSAN storage configuration, since a faster cache tier will support a larger storage tier.

Intel Optane SSDs provide so much improvement over older SSDs in performance, availability, and reliability that we expect it to continue to fundamentally change the design of storage technologies, including SDS solutions such as VMware's vSAN. IT managers can extend the life and increase the capabilities of their current vSAN configuration simply by replacing the caching tier with Intel Optane SSDs. IT Managers can also design a new class of SDS nodes utilizing Intel Optane SSDs at the caching tier, and combining them with Intel 3D NAND SSDs at the storage tier, creating an all-flash vSAN node with performance and capacity that exceeds traditional vSAN node capabilities. This will enable vSAN solutions that are suitable for an entirely new class of high performance, high capacity, workloads.