

White Paper

Next-Generation Data Platform for Hyperconvergence

Next-Generation Data Platform for Hyperconvergence

White Paper
March 2016

A Platform for a New Generation of Applications and Data


In the VMware vSphere environment, the controller occupies a virtual machine with a dedicated number of processor cores and amount of memory, allowing it to deliver consistent performance and not affect the performance of the other virtual machines on the cluster. The controller can access all storage without hypervisor intervention through the VMware VM_DIRECT_PATH feature. It uses the node's memory and SSD drives as part of a distributed caching layer, and it uses the node's HDDs for distributed capacity storage. The controller integrates the data platform into VMware software through the use of two preinstalled VMware ESXi vSphere Installation Bundles (VIBs):

- IO Visor:

The log-structured file system assembles blocks to be written to the cache until a configurable-sized write log is full or until workload conditions dictate that it be destaged to a spinning disk. When existing data is (logically) overwritten, the log-structured approach simply appends a new block and updates the metadata. When the data is destaged, the write operation consists of a single seek operation with a large amount of sequential data written. This approach improves performance significantly compared to the traditional read-modify-write model, which is characterized by numerous seek operations, with small amounts of data written at a time.

When data is destaged to a disk in each node, the data is deduplicated and compressed. This process occurs after the write operation is acknowledged, so no performance penalty is incurred for these operations. A small deduplication block size helps increase the deduplication rate. Compression further reduces the data footprint. Data is then moved to HDD storage as write cache segments are released for reuse (Figure 4).

Hot data sets—data that is frequently or recently read from the persistent tier—are cached both in SSD drives and in memory (Figure 5). Having the most frequently used data in the caching layer helps make Cisco HyperFlex Systems perform well for virtualized applications. When virtual machines modify data, the data is likely



Data Services

