



# Accelerating Swift with Intelligent SSD Caching

Dramatically improve Swift\* storage performance, greatly reduce data loss risk.

## Challenge

The explosive growth of unstructured data has driven the need for expanded safe and efficient data storage. Enterprises and cloud service providers (CSPs) are transitioning from traditional storage to software-defined storage (SDS) solutions, running on commodity servers, as a cost-effective way to meet ever expanding storage demands. The OpenStack Object Store project, known as Swift, is an SDS solution that was developed to address petabyte and beyond data growth and scalability.

OpenStack Swift is an object-based storage architecture in which the basic unit of stored data is called an object. The type of content in each object is determined by the user (e.g., text, videos, images, email attachments, etc.). Swift provides each stored object with a rich set of metadata that it uses for access, security and most importantly superior data mining. These objects are durably stored (generally on HDDs) within a cluster of machines.

Swift is a powerful storage solution, however it does have its own performance and reliability challenges. For example, in order to read and write objects Swift must retrieve the file system metadata from HDDs, which can cause performance delays. Similarly the Swift replication process which protects data in case of system failures relies on file system metadata to operate. This critical replication process can be time consuming, based on used capacity and size of the cluster.

## Solution

Intel collaborated with SwiftStack, a lead contributor to the Openstack Swift project, to address these challenges. By utilizing SwiftStack's deep insights into the Swift architecture, we've integrated and tested a high-performance PCIe\*/NVMe\*-based Intel SSD with Intel® Cache Acceleration Software (Intel® CAS) into Swift.

In this reference solution we obtained a **3x improvement** to the Swift cluster throughput, and a **>5x improvement in replica convergence time** – down from 6 hours (without CAS) to 1 hour and 5 minutes with an Intel® Solid State Drive (SSD) with Intel CAS.

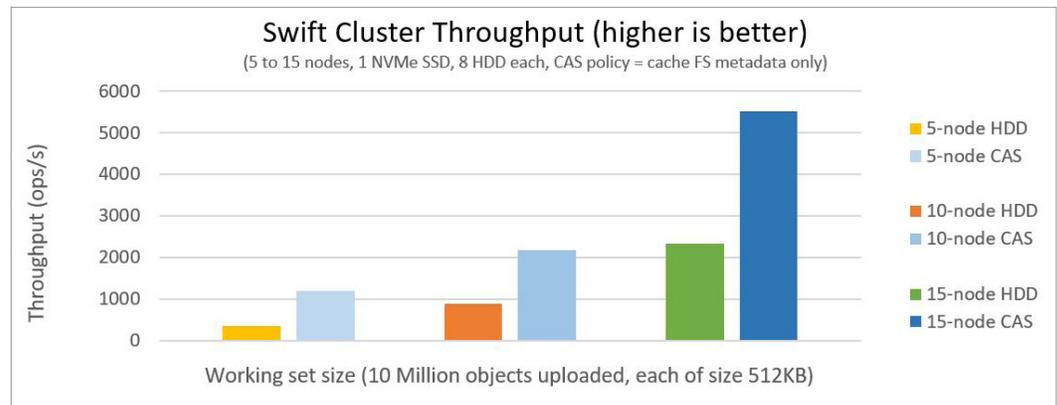


Figure 1: Swift Cluster Throughput

## Fix Bottlenecks with SwiftStack + Intel® SSD + Intel Cache Acceleration Software

This collaborative and high-performance solution was achieved by caching the file system metadata onto an NVMe-based Intel SSD with Intel CAS, on each Swift storage node.

Unlike conventional and less effective caching techniques that rely solely on the frequency that data is accessed, Intel CAS has an intelligent classification-based solution called differentiated storage services (DSS) that prioritizes I/O. Because of this intelligence, Intel CAS can identify the data types that lead to bottlenecks. In this case, performance sensitive metadata such as inodes were identified and cached on to an SSD, which accelerates the search and retrieval of an object.

Intel CAS on a high-performance NVMe-based Intel SSD improves the GET throughput and bandwidth by 2.5x to 3.5x (average of 3x) for Swift cluster sizes ranging from 5-15 nodes. The benefits of this approach are shown in Figure 1. These results are consistent across a range of object and cluster sizes.

## Improving Small File Latency

Small file performance is a known problem area for Swift storage clusters. Figure 2 shows latency improvements of utilizing an NVMe-based Intel SSD in a five node cluster with Intel CAS. This cluster continues to perform well across a range of object sizes, especially with small object sizes (64k), where up to 7x performance gains were realized.

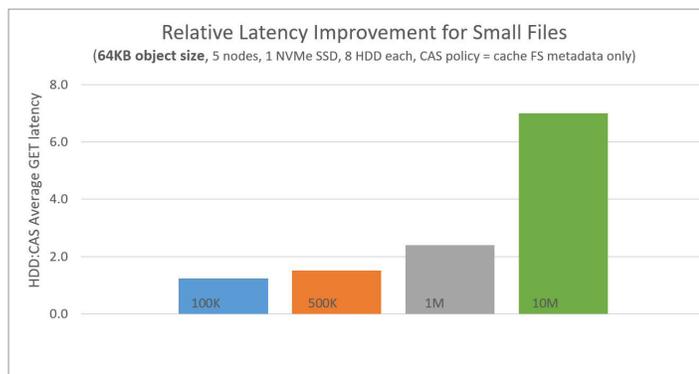


Figure 2: Relative Latency in Small Files

Intel Cache Acceleration Software: <http://intel.com/cas>  
 Intel SSD technical specifications: <http://intel.com/ssd>  
 SwiftStack: <https://www.swiftstack.com/try-it-now>

Testing performed by Intel on various Swift cluster sizes (5-15 object storage nodes and several proxy nodes)  
 Cluster configuration: Processors: 2x Intel® Xeon® E5-2699 (45MB cache, 2.3GHz, 18 cores), DRAM 128 and 256GB Memory, HDD 8 x 2TB Seagate® ST2000NX0403, NVMe SSD 1 x 2TB Intel® SSD DC P3520, Intel® RMS3CC080 RAID. Controller: SAS 3.0 Dell Mezzanine\* SAS/SATA 8 Port 1GB. Network: 2 x Intel® X540 10Gbe NICs, 2 x Niantic NIC, 2 x Intel® X520-DA2 and 2 x 10G NIC - Intel® X540. Operating system: Ubuntu 14.04.5 kernel revision: 4.4.0-47-generic, Swift Version: 2.9.0.2-4~trusty, SwiftStack Controller version 4.7.0.1. Cache Acceleration Software (CAS) 3.1.1, COSBench version - 0.4.2.c4  
 Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer. For copies of this document, documents that are referenced within, or other Intel literature, please contact your Intel representative.  
 Intel, Xeon, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

## Improving Swift Replicator Performance by Over 5x

Swift replication is designed to keep the data protected in the face of error conditions, such as network outages or drive failures. The replication process uses hash files to keep track of its multiple copies. When one of these copies disappears due to a lost HDD or disk corruption, a new version is replicated to ensure desired number of copies (typically three) of an object are available at all times.

In a typical Swift cluster with many storage nodes and a large number of uploaded objects, the replication process could take days or weeks, especially as cluster sizes increase. During this replication process data may be vulnerable to loss.

Figure 3 shows a >5x improvement in replica convergence time – down from 6 hours (for default HDD-only setup) to 1 hour 20 min (Intel CAS caching metadata only) to 1 hour 5 min (for Intel CAS caching metadata + hashes files).

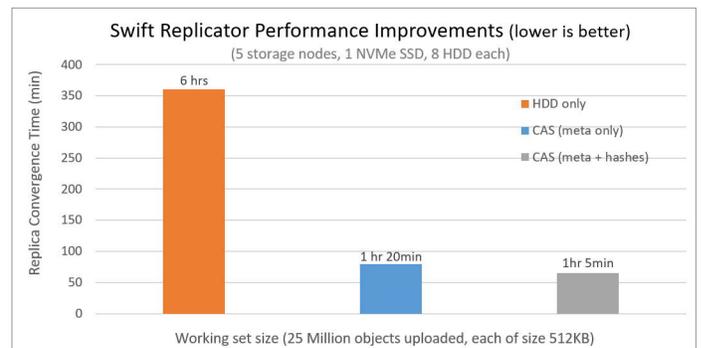


Figure 3: Swift Replicator Performance Improvements

## Conclusion

Deploying an Intel SSD with Intel Cache Acceleration Software provides superior performance improvements in a Swift cluster. As enterprises face exponential storage growth, adopting solutions such as Swift and a hybrid storage model, can keep costs down while meeting customer SLAs. Enterprises benefit from this high-performance solution, along with outstanding support and documentation from Intel, without the need for application modifications or an infrastructure overhaul.

