

This asset was published on January 2023 and may contain out-of-date information. Please visit <https://core.vmware.com/resource/selecting-workloads-vmware-cloud-flex-storage%E2%84%A2> for the latest version.

Selecting Workloads for VMware Cloud Flex Storage™

VMwareStorage

Table of contents

Selecting Workloads for VMware Cloud Flex Storage™	3
Solution Overview	3
Cloud Flex Storage Delivers Efficiency For Changing Application Requirements	3
Cloud Flex Storage Architecture	3
Testing Applications on Cloud Flex Storage	4
Conclusion	5

Selecting Workloads for VMware Cloud Flex Storage™

Solution Overview

Cloud Flex Storage Delivers Efficiency For Changing Application Requirements

The default storage of VMware Cloud on AWS is the Scale-out architecture of VMware vSAN. This enables powerful non-disruptive scale-up or scale-out capabilities. You can non-disruptively expand capacity and performance by adding hosts to a cluster (scale-out) or selecting from different node options for different clusters. As application workloads organically grow. Despite this some capacity heavy workloads may be challenging to fully optimize using vSAN alone. Examples include:

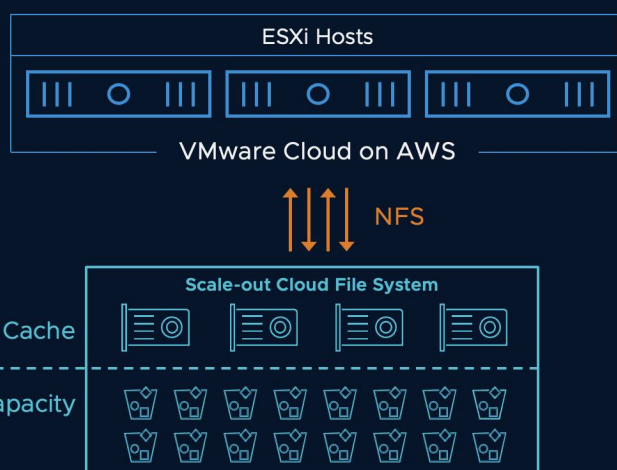
- Applications such as file servers, log retention, media storage, and other unstructured data that has a static CPU consumption, but growing storage requirement.
- An application refactor requires significant storage being added for logs.
- A new line of business for analytics may consume excessive storage.
- Changing regulatory requirements may demand significant unplanned storage retention.
- The deployment of Networking Virtualization, powered by NSX, has enabled the migration and consolidation of legacy stranded cluster resources.
- Certain workflows such as VDI, benefit from having Datastores that can be presented to multiple clusters within a SDDC.

Cloud Flex Storage Architecture

VMware Cloud Flex Storage™ is built on a mature, enterprise-class filesystem that has been developed and production-hardened over many years, dating back to Datrium’s DHCI storage product, which VMware acquired in July 2020. It is the same filesystem that has been backing the VMware Cloud Disaster Recovery service. The filesystem has a two-tier design that allows for independent scaling of storage performance and capacity, using a Log-Structure Filesystem (LFS) design. You can read more about the filesystem architecture in Sazzala Reddy’s (Chief Technologist and a founder of Datrium) [blog here](#). The combination of LFS with a 2-tier design, along with efficient snapshots and immutability, makes this a multi-purpose filesystem that unlocks many use cases, such as backup, disaster recovery, ransomware protection, and recovery. With VMware Cloud Flex Storage, we are extending this proven technology to primary storage and making it available in the public cloud, where it delivers exceptional storage performance, scalability, and cost efficiency for traditional and modern workloads.

Cost-Effective VM Storage for Capacity-Heavy Workloads

NVMe performance with S3-compatible durability



VMware managed NFS storage for virtual machines

Single support model as part of VMware Cloud

NVMe performance with optimized capacity pricing

- Low latency writes
- Variable latency reads

Compliments vSAN storage

A goal of this solution was to try to strike a balance between the performance of NVMe AWS hardware, and the durability and economics of object storage. The Write IO path has all incoming writes mirrored to redundant NVMe devices. Generous allocations of read cache in addition, help smooth out the higher latency response of S3 storage for cache friendly workloads. Do note that outliers of colder data from S3 that are not cached will respond in the tens of milliseconds. For workloads requiring more consistent reads on cold data such as transactional database workloads, vSAN is still recommended.

Configuring Cloud Flex Storage is simple!



<https://youtu.be/>

Testing Applications on Cloud Flex Storage

Cloud Flex storage provides supplemental storage that is significantly more cost effective on a cost per GB than adding an additional host to a existing SDDC purely for storage expansion. To validate that a workload will be suitable a number of options are available.

Review existing performance requirements

Using VMware Live Optics Virtual Assessments, vRealize Operations Reporting, or even vCenter performance graphs, the existing performance demands of an application can be measured. Due to all write IO being cached, the primary concern for latency response is going to come from random read heavy workloads that exceed the cache working set size. Large (multi-TB working set), random read heavy applications needing low latency should perhaps be prioritized to remain on the existing vSAN cluster, while write heavy, low IO, capacity hungry workloads should be prioritized for moving to Cloud Flex Storage.

Application Owner Interviews

Discuss with the application owner what the data set is used for. A reporting database that operates on a nightly process is unlikely to result in dissatisfied users as long as it can run through its generation of reporting within the required time frame.

Inversely a hospital EMR, or a ERP system that users often are "clicking and waiting" when there is disk latency should be prioritized to remain on vSAN. General purpose file shares, or file shares that store logs and other "Write once, Read Maybe" workloads are strong candidates for VMware Cloud Flex Storage.

Stack Ranking Based on Performance Demands

Given an existing vSAN datastore will be provisioned no matter what with a VMware Cloud on AWS SDDC, simply sizing the cluster for compute, and then adding Cloud Flex storage as needed for additional storage can be used as a first pass sizing method. Simply rank workloads by priority and performance demand and use this ranking to determine the first pass of "what goes on which storage option".

Test and Try Methodology

If the application owner can tolerate live testing, simply performing a storage vMotion of the workload to Cloud Flex Storage and following up with user experience interviews can provide a simple way to see if Cloud Flex Storage will work for a given workload. Given storage vMotion can be invoked to move the workload back to vSAN the "test and see if users notice" is sometimes an acceptable method of testing the lower cost Cloud Flex Storage.

Conclusion

It is worth noting that the Scale-out Cloud File System (SCFS) is highly performant due to its cache heavy architecture. Given the low friction to add a datastore to see if it will work for existing workloads, we strongly encourage customers to test it out and see if it can help remove roadblocks to application deployment or migration to the VMware Cloud on AWS SDDC.

