

VMware Cloud Director Object Storage Extension 2.0 – Reference Design

Table of contents

Introduction	3
Audience	3
Use Cases	3
Storing Content	3
Persistent Storage for Application	7
Storing vApp Templates and Catalog	7
Architecture	9
High-Level Architecture View	9
OSE 2.0 Deployment View	11
Deployment Options.....	12
Small Deployment	12
Medium Deployment	12
Large Deployment	12
Multisite Deployment	14
Test Environment Reference Benchmark	16
Cloudian HyperStore Test Setup	16
Cloudian HyperStore - Bill of Materials	17
Cloudian HyperStore Test Results	17
Dell EMC ECS Test Setup	21
Dell EMC ECS - Bill of Materials	21
Dell EMC ECS – Test Results	22
AWS S3 Test Setup	25
AWS S3 - Bill of Materials	25
AWS S3 – Test Results	26
Abbreviations	29

Introduction

The VMware Cloud Director® Object Storage Extension™ allows VMware Cloud Providers™ who are using VMware Cloud Director to offer object storage services to their end-users. The extension acts as middleware which is tightly integrated with VMware Cloud Director to abstract 3rd party S3 API compatible storage providers in a multitenant fashion.

This guide provides information on how to properly design VMware Cloud Director Object Storage Extensions on top of a VMware Cloud Director infrastructure.

This document is specific to version 2.0 of VMware Cloud Director Object Storage Extension and its integration with Cloudian HyperStore, Dell EMC ECS, and AWS S3. Information about how Object Storage Extension can utilize other S3-compatible storage through the Object Storage Interoperability Service (OSIS) can be found in a separate whitepaper.

Audience

This document is intended for VMware Cloud Provider architects and technical leads responsible for planning and executing the deployment and upgrades of a VMware-based cloud environment.

Use Cases

VMware Cloud Director natively provides Infrastructure as a Service (IaaS) by integrating with the underlying VMware vSphere platform. All native storage services such as storage for virtual machines, named (independent) disks, and catalog storage for virtual machine templates and media are using storage attached to vSphere ESXi hosts such as block storage, NFS, or VMware vSAN.

There is, however, the need for highly scalable, durable, and network-accessible storage that could be utilized by tenants or their workloads without the dependency on the vSphere layer. The VMware Cloud Director Object Storage Extension (OSE) provides access to the object storage either through VMware Cloud Director UI extension or via standardized S3 APIs. This allows existing applications to easily access this new type of storage for various use cases.

Storing Content

Through the VMware Cloud Director User Interface, users can create storage buckets and upload and tag unstructured files (objects) of various types. These files can be easily accessed with Uniform Resource Locator (URL) links or directly previewed from the OSE plug-in. They can also be synced with the attached S3 object storage and shared by defining their Access Control Lists. Version control can also be applied to objects of a bucket, which helps you track their changes. Objects can also be downloaded and tagged. For protection, object lock can be applied to them. You can also update the metadata of the stored objects, thus providing more information for the stored file. In addition, you can manage the lifecycle of the available buckets and their stored objects and apply bucket read/write policies (See: [Figure 4. OSE Figure 4. ACL and Bucket Policy](#)). Server-side tenant-level encryption of buckets' content is also possible with OSE 2.0 (See: [Figure 5. Server-side Tenant Level Encryption](#)). However, it is only applied to new objects.

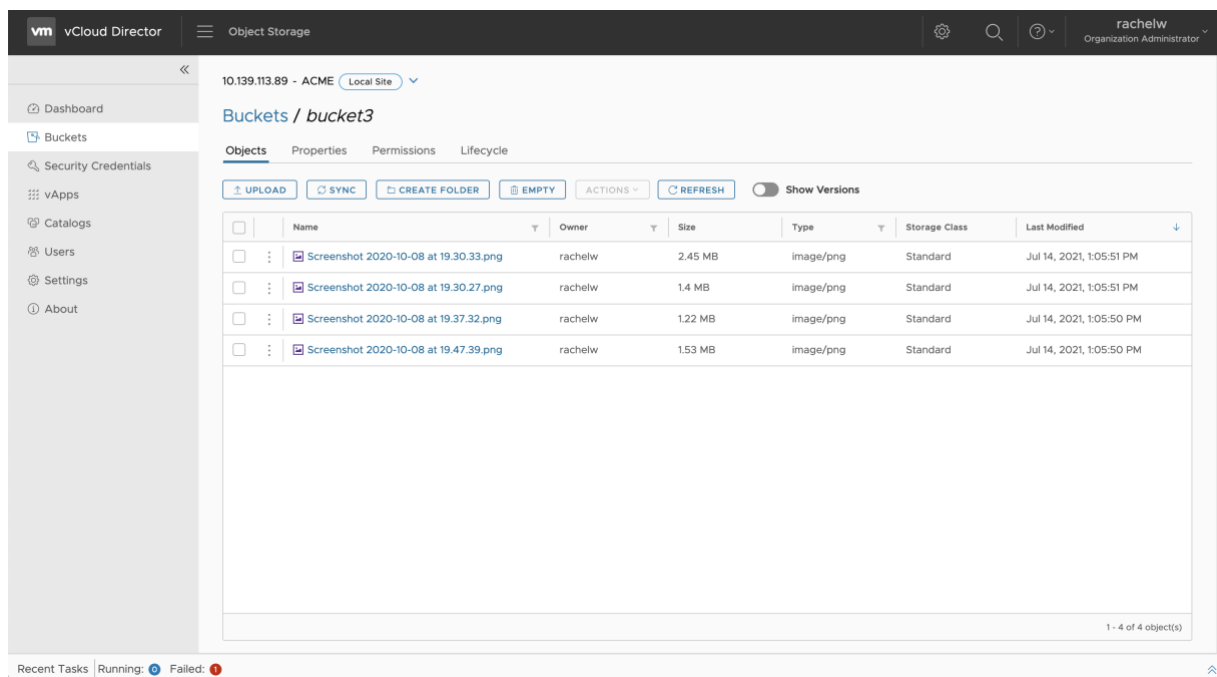


Figure 1. Object Upload to S3 Bucket

Thanks to the full S3 API compatibility, it is also possible to utilize existing 3rd party applications to upload and manage the files. The following figures display a freeware S3 browser for viewing a bucket and its API access configuration.

Edit Account [online help](#)

Edit account details and click Save changes

Account Name:

Assign any name to your account.

Account Type:

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Signature Version:

Choose the supported signature version. Default value is Signature V2.

Access Key ID:

Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

Secret Access Key:

Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

☐ Encrypt Access Keys with a password:

Turn this option on if you want to protect your Access Keys with a master password.

☒ Use secure transfer (SSL/TLS)
If checked, all communications with the storage will go through encrypted SSL/TLS channel

[Click here to sign up for Amazon S3...](#)

Figure 2. S3 Browser Configuration

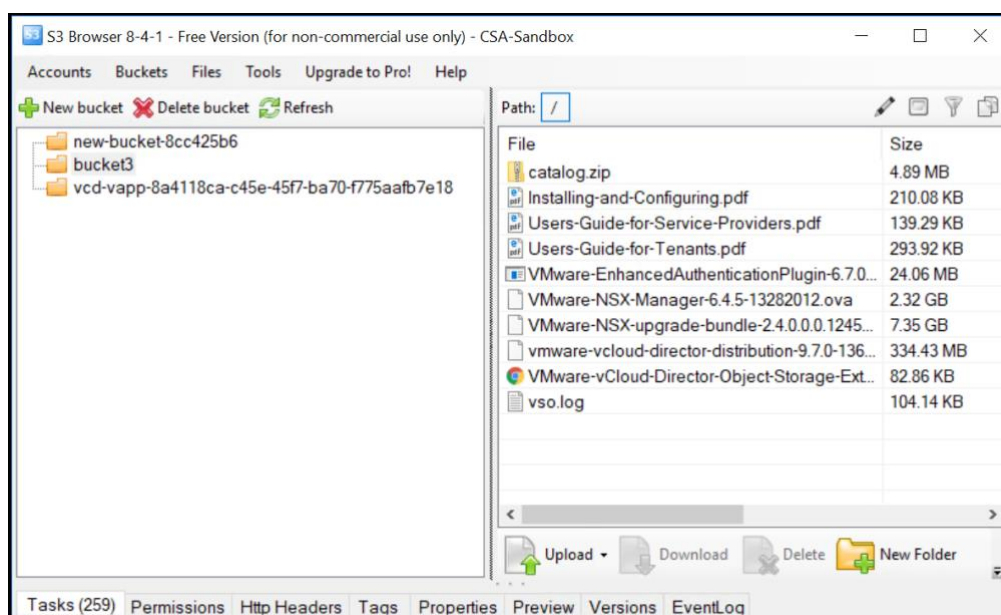


Figure 3. S3 Browser Application

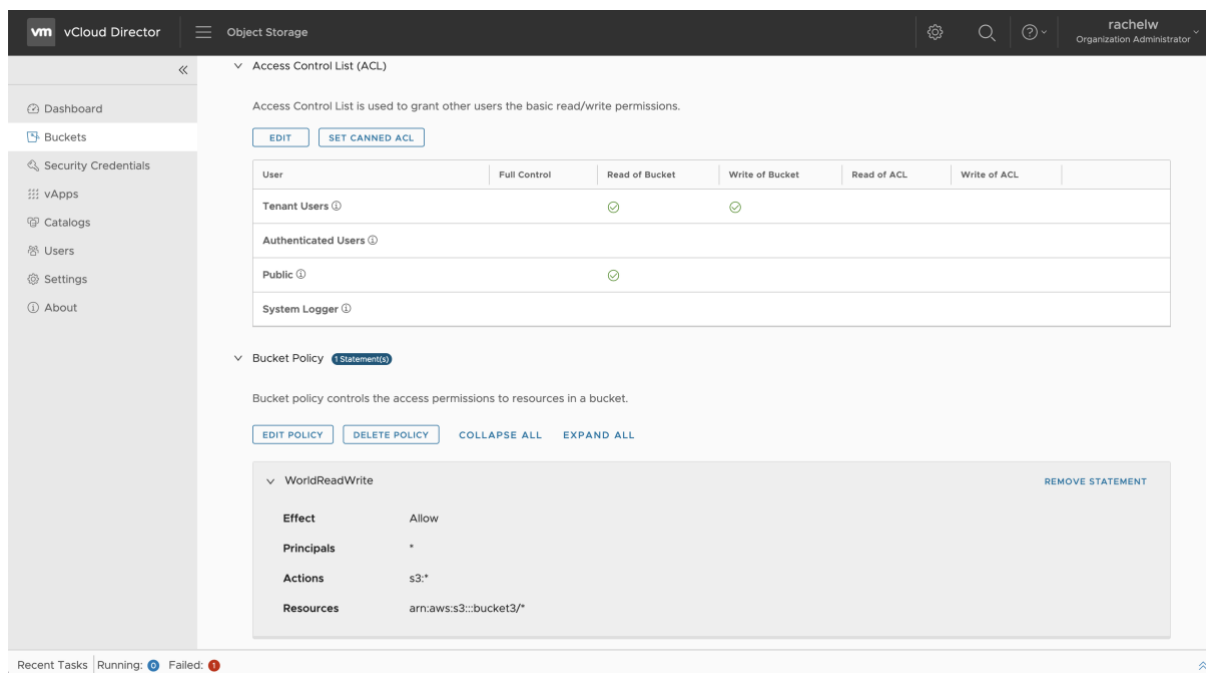


Figure 4. OSE Figure 4. ACL and Bucket Policy

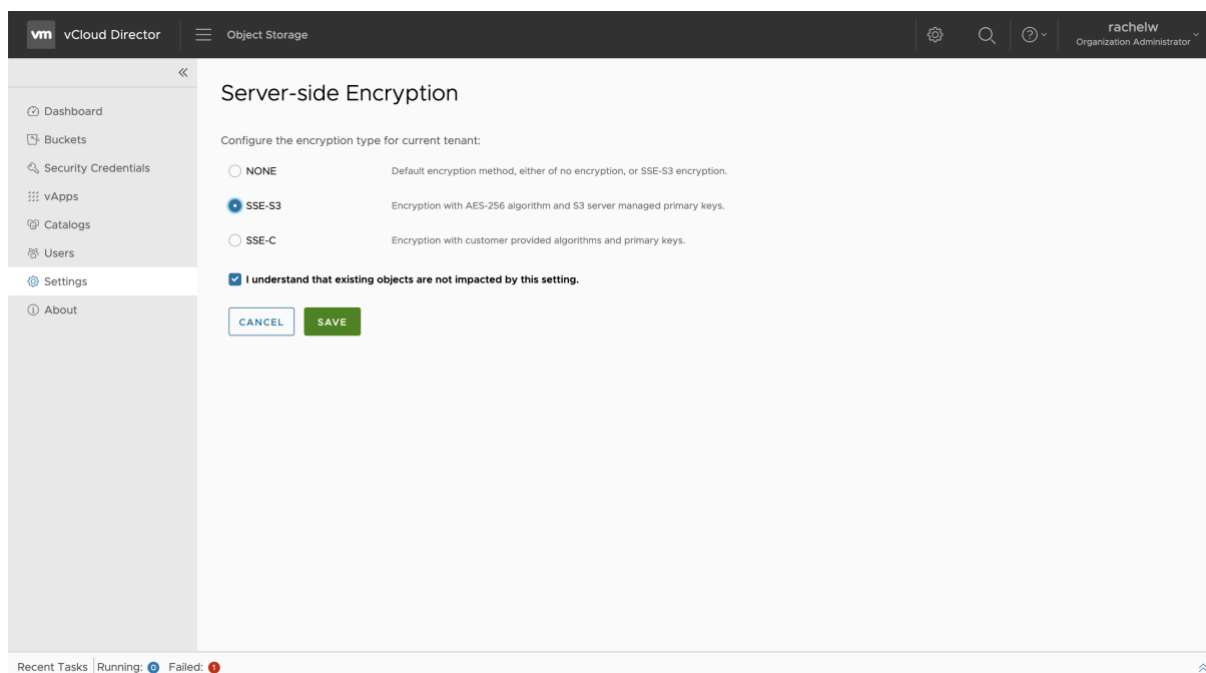


Figure 5. Server-side Tenant Level Encryption

Persistent Storage for Application

Users can create application credentials with limited access to a specific bucket. This allows (stateless) applications running in VMware Cloud Director (or outside) to persist their content such as configurations, logs, or static data (web servers) into the object-store. The application is using S3 API over the Internet to upload and retrieve object data.

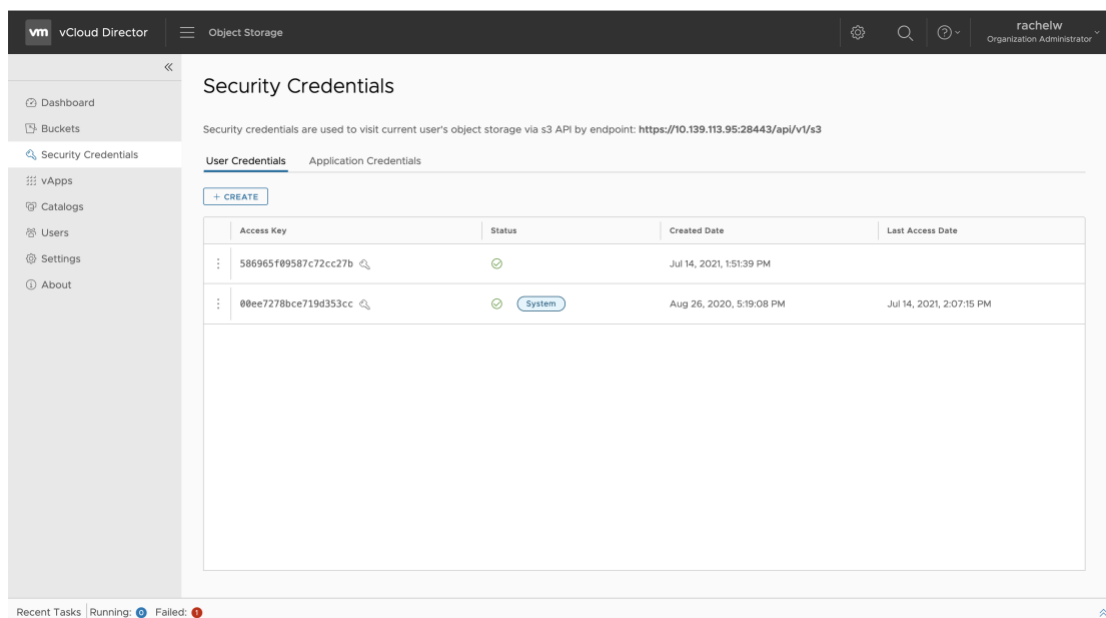


Figure 6. Application Credentials

Storing vApp Templates and Catalog

Because of the close integration with VMware Cloud Director, VMware Cloud Director Object Storage Extension can directly capture and restore user's VMware Cloud Director vApps. Users can also share these vApps with other users. Thus, VMware Cloud Director Object Storage Extension provides an additional tier of storage for vApp templates that can be used, for example, for archiving old images.

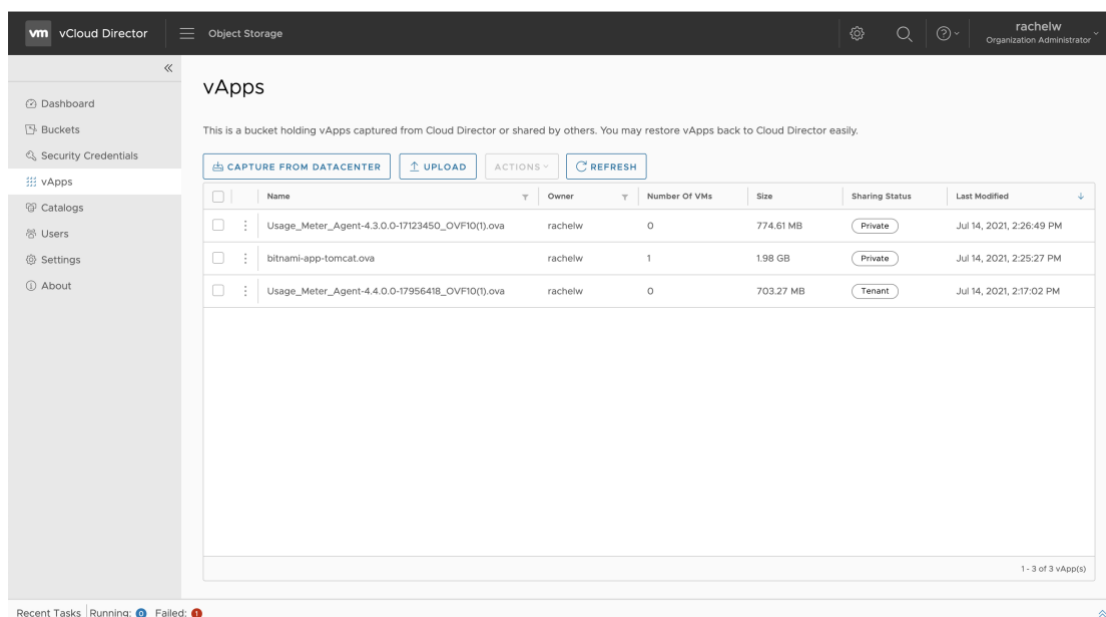


Figure 7. vApp Template Integration

A full VMware Cloud Director catalog (consisting of vApp templates and media ISO images) can be captured from an existing Org VCD catalog or created from scratch by uploading an individual ISO and OVA files to VMware Cloud Director Object Storage Extension. Then, the catalog can be published, which allows any VMware Cloud Director organization (from any VMware Cloud Director instance) to subscribe to the catalog. As a result, this OSE functionality enables easy distribution of specific catalogs publicly or geographically across VMware Cloud Director instances.

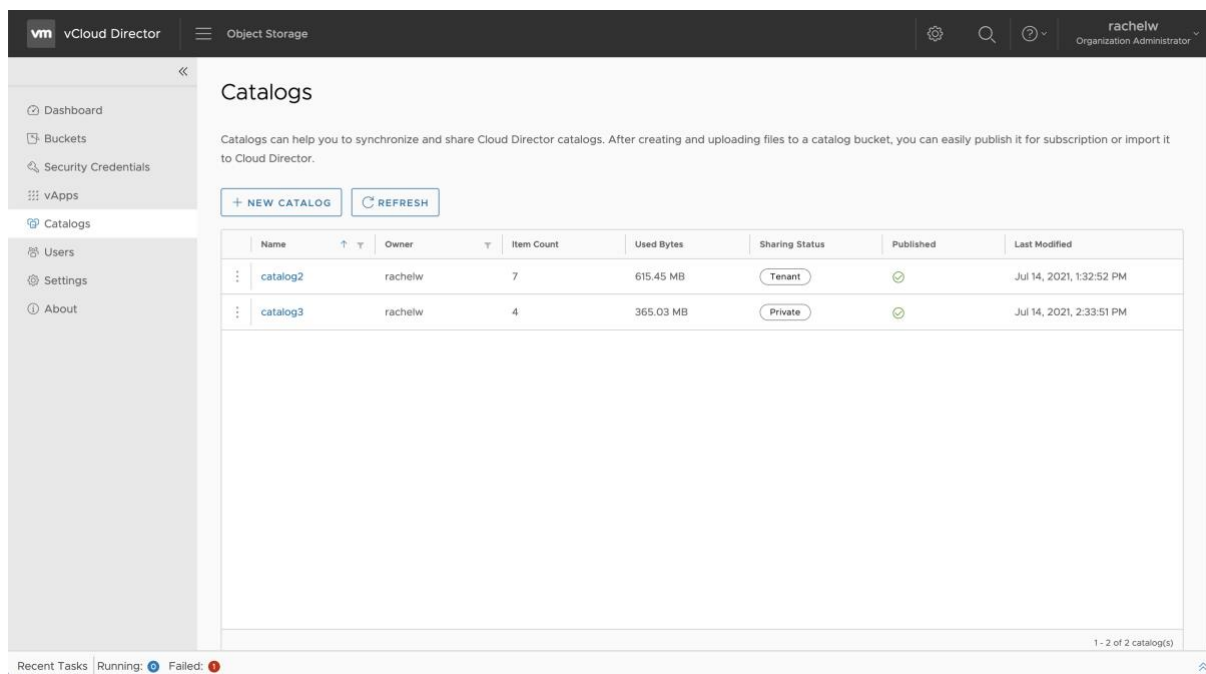


Figure 8. Catalog Integration

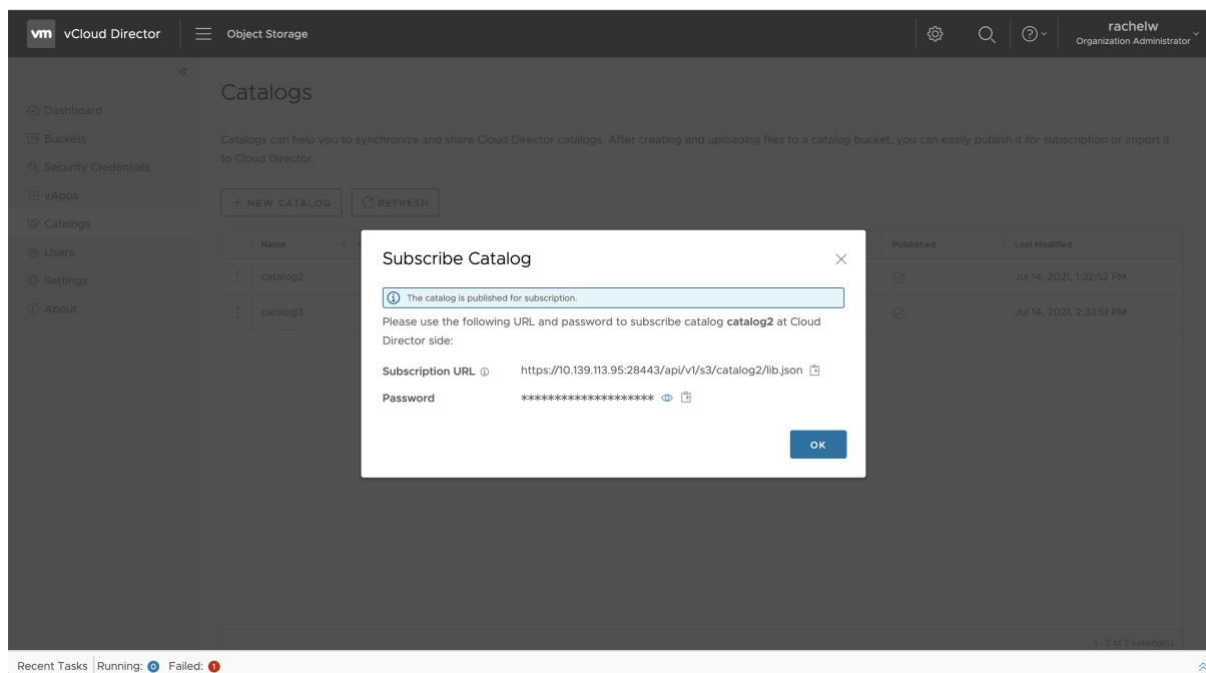


Figure 9. Catalog Published Directly from Object Store

Architecture

High-Level Architecture View

VMware Cloud Director Object Storage Extension 2.0 runs externally to VMware Cloud Director and integrates through a UI plug-in. It has a 1:1 relationship with the VMware Cloud Director instance. OSE 2.0 supports VMware Cloud Director version 10.0 and later.

In the current 2.0 release, VMware Cloud Director Object Storage Extension can be connected to the following storage providers: Cloudian HyperStore, Dell EMC ECS, AWS S3, or another S3-compatible storage platform). The provider can selectively enable VMware Cloud Director organizations to consume the service. The unique counterparts for organizations and users are created at the storage provider. The users authenticate to the service with VMware Cloud Director or S3 credentials and access it only through the UI plug-in. The provider can directly access the underlying storage appliance to set quotas or collect usage information for billing purposes.

You can switch between storage platforms that you use with VMware Cloud Director Object Storage Extension but cannot use two storage platforms simultaneously.

An instance of VMware Cloud Director Object Storage Extension can work with a single instance of VMware Cloud Director or a single VMware Cloud Director server group.

In addition to the storage platform that OSE will connect to Cloud Director, three or more (for high availability and scalability) RHEL/CentOS/Oracle Linux VM nodes that run OSE, provided as an RPM package, are required. The number of the OSE VM nodes depends on the used S3 storage. See for reference: [Deployment Options](#). These VMs are essentially stateless and persist all their data in PostgreSQL DB version 9.5 or later. This could be VMware Cloud Director external PostgreSQL DB (if available) or a dedicated database for VMware Cloud Director Object Storage Extension.

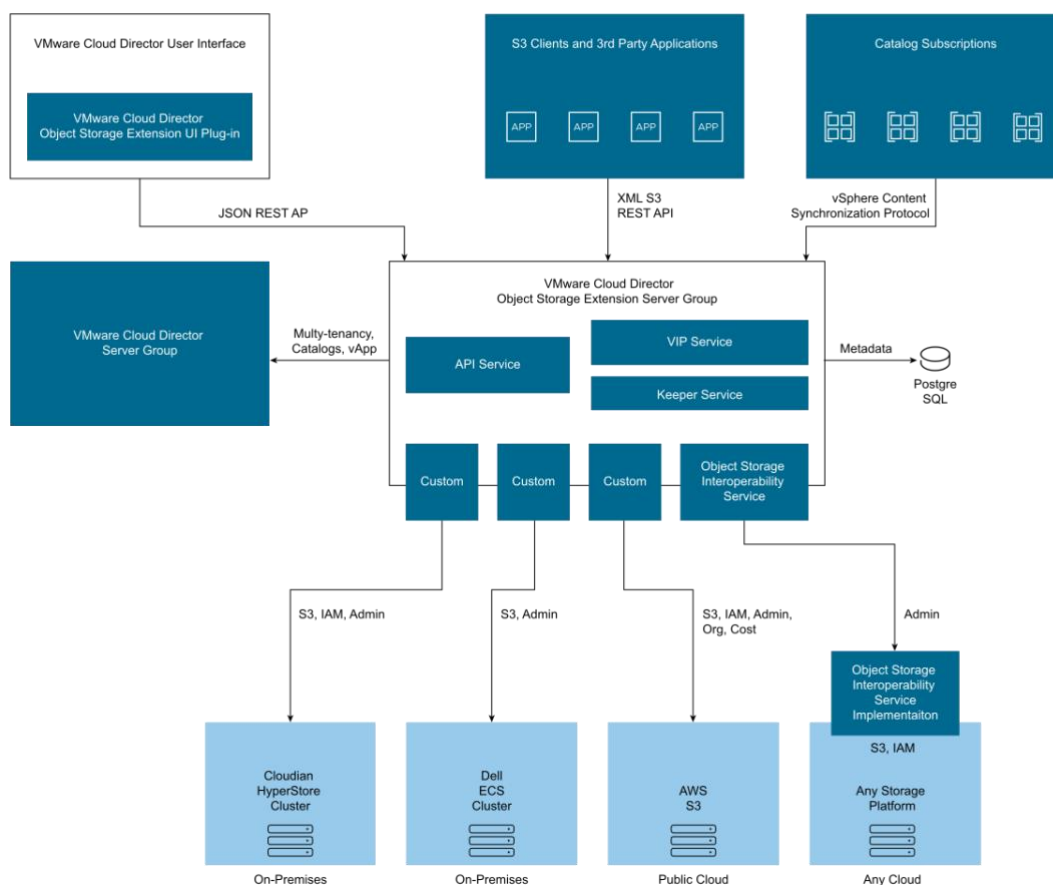


Figure 10: OSE 2.0 High-level Architecture View

VMware Cloud Director Object Storage Extension (OSE) enables Cloud Director tenant users to use object storage by native UI experience and support S3 clients to consume the object storage by S3 APIs.

The mapping between Cloud Director and Cloudian, ECS, and AWS is the following:

- VMware Cloud Director service provider is mapped to an ECS admin user, Cloudian admin user, or AWS master account.
- VMware Cloud Director tenant is mapped to an ECS namespace, Cloudian group, or AWS org unit.
- VMware Cloud Director user is mapped to an ECS user, Cloudian user, or AWS IAM user.

In 2.0, OSE adds Amazon Native S3 as another storage platform. It also supports deploying the AWS public cloud to work with Cloud Director Service (CDS) to offer object storage service.

The main architecture changes for this release are:

- Adds AWS as a new supported object storage platform
- Opens extension for any storage platform vendors to integrate with OSE
- For AWS support, we add built-in custom tenant and user mapping in OSE.

For platform extension, we propose an Object Storage Interoperability Service (OSIS) specification in Swagger. OSE certifies any storage platforms which implement the API specification.

OSE 2.0 Deployment View

The following example of an OSE deployment depicts the connections between OSE, Cloud Director, and the supported S3 storage – Cloudian, ECS, and AWS. It also outlines the communication flow between those products. OSE connects to Cloud Director as a plug-in, which shows either provider or tenant information depending on the type of logged-in user.

As the following diagram suggests, there are 2 UI views for OSE in Cloud Director – Provider and Tenant view. When a provider administrator logs in to Cloud Director and opens Object Storage Extension, they can connect to the Management Console of the underlying storage through a URL redirect. An internal load-balancer then distributes those requests to the storage nodes. Provider administrators can also send requests to the underlying storage through OSE, as shown on the diagram.

When a tenant user logs in to Cloud Director with enabled OSE and sends requests to OSE, then their requests go through an S3 client and are sent to an external network load balancer used for OSE nodes. Then the OSE load balancer distributes the requests to the OSE nodes. After the requests made by the tenant user reach OSE, they continue to their destination – which is the nodes of the S3 storage. The tenant requests are distributed by an internal load balancer used for the S3 storage. The results of the requests sent to OSE are recorded in the PostgreSQL database connected to the OSE

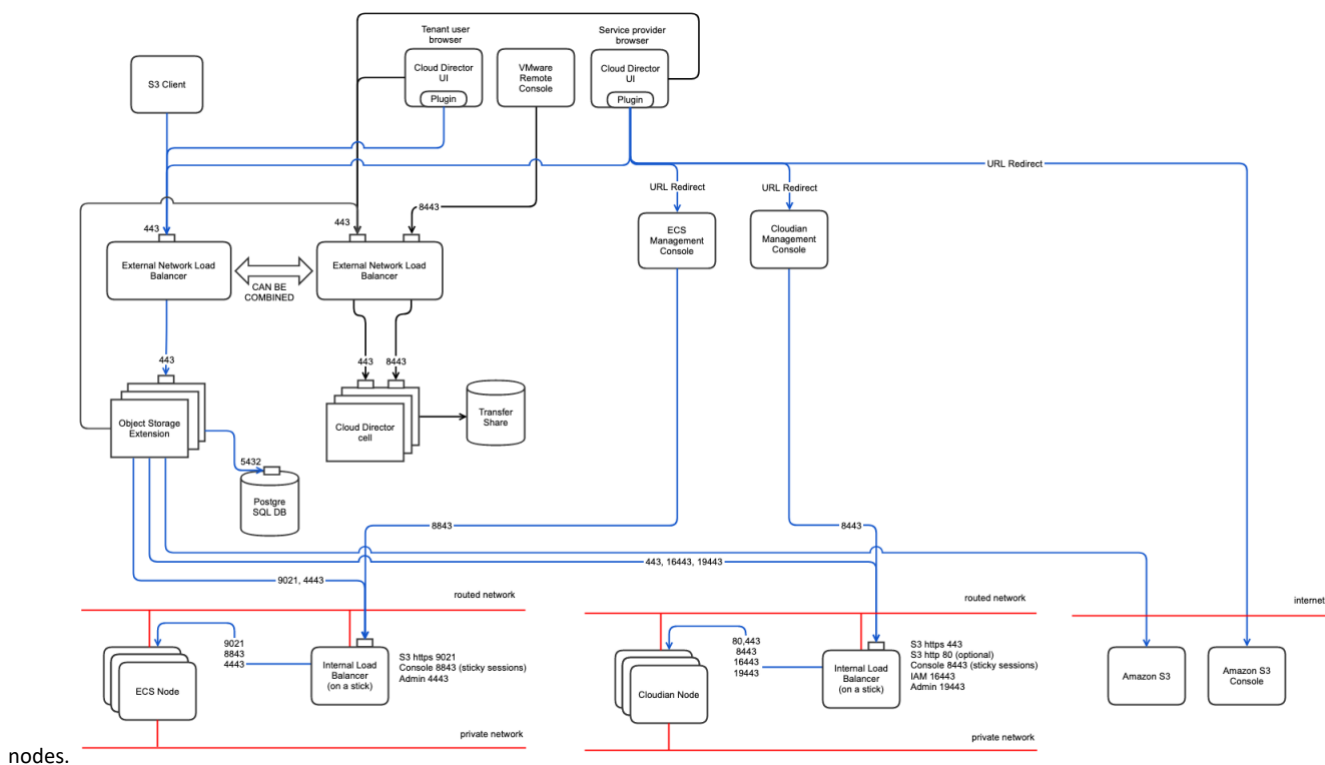


Figure 11: OSE 2.0 Deployment Diagram

Black lines - Existing Cloud Director network flows

Blue - New flows related to Object Storage Extension

Red – ECS, Cloudian or AWS networks

Note: The internode communication for Cloud Director cells, ECS nodes, Cloudian nodes and AWS is not shown.

Deployment Options

Based on the use case, user target group, and expected service parameters (SLA, scalability), the cloud provider can decide on the type of deployment.

Small Deployment

Usage: Niche use cases

- Requirement: Minimum resources required. High availability, supported for production.
- One or more RHEL/CentOS VMs for VMware Cloud Director. External PostgreSQL database (used both for VMware Cloud Director and VMware Cloud Director Object Storage Extension). NFS transfer share is needed when more than one VMware Cloud Director cell is used. Protected with vSphere HAs.
- One CentOS Linux 7 or 8/RedHat Enterprise Linux 7/Oracle Linux 7 VM: (4 vCPU, 8 GB RAM, 120 GB HDD) running VMware Cloud Director Object Storage Extension. Protected with vSphere HA.
- vSphere/NSX: As required for VMware Cloud Director resources.
- Storage provider: Three CentOS virtual machines running Cloudian HyperStore, or Five CentOS virtual machines running Dell EMC ECS (4 vCPUs, 32 GB RAM, 32+100 GB HDD on shared storage) or AWS S3.
- Load balancing: VMware Cloud Director cells and Cloudian HyperStore or Dell EMC ECS nodes load balancing provided by NSX.

Medium Deployment

Usage: typical use cases

- Requirement: High availability, supported for production.
- Multiple RHEL/CentOS or appliance VMs for VMware Cloud Director. NFS transfer share. For non-appliance form factor external PostgreSQL database.
- One or more CentOS Linux 7 or 8/RedHat Enterprise Linux 7/Oracle Linux 7 VMs: (8 vCPU, 8 GB RAM, 120 GB HDD) running VMware Cloud Director Object Storage Extension. Protected with vSphere HA and optionally load balanced. If VMware Cloud Director is deployed in appliance form factor, an external PostgreSQL database is needed; otherwise, the existing VMware Cloud Director database can be reused).
- vSphere/NSX: As required for VMware Cloud Director resources.
- Storage provider: Three CentOS virtual machines running Cloudian HyperStore, Five CentOS virtual machines running Dell EMC ECS on dedicated ESXi hosts with local disks (8 vCPUs, 64 GB RAM, 32 GB HDD + multiple large local disks) or AWS S3.
- Load balancing: VMware Cloud Director cells and Cloudian HyperStore, or Dell EMC ECS nodes load balancing provided by NSX or external hardware load balancer.

Large Deployment

Usage: large scale, low cost per GB use cases

- Requirement: High scale, performance, and availability, supported for production.
- Multiple RHEL/CentOS or appliance VMs for VMware Cloud Director. NFS transfer share. For non-appliance form factor external PostgreSQL database.
- Multiple CentOS Linux 7 or 8/RedHat Enterprise Linux 7/Oracle Linux 7 VMs (12 vCPU, 12 GB RAM, 120 GB HDD) running VMware Cloud Director Object Storage Extension. If VMware Cloud Director is deployed in an appliance form factor, an external HA PostgreSQL database is needed; otherwise, the existing Cloud Director database can be reused).
- vSphere/NSX: As required for VMware Cloud Director resources.
- Storage provider: Three or more dedicated bare-metal physical Cloudian HyperStore, Five or more physical Dell EMC ECS, or AWS S3.

- Load balancing: an external hardware load balancer

The following figures display how you scale out and load balance Object Storage Extension with Cloudian HyperStore, Dell EMC ECS, and AWS S3.




Figure 12. Example of Scale Out of Object Storage Extension Deployment with Load Balancing

Multisite Deployment

Object Storage Extension supports VMware Cloud Director multisite deployments where different VMware Cloud Director instances are federated (associated) with each other with a trust relationship. As these instances can be deployed in different locations, the end-users can deploy their applications with a higher level of resiliency and not be impacted by local datacenter outages.

Each VMware Cloud Director instance has its own VMware Cloud Director Object Storage Extension, which communicates with shared S3 object storage deployed in a multi-datacenter configuration. Objects are automatically replicated across all data centers, and VMware Cloud Director users can access them through either VMware Cloud Director or VMware Cloud Director Object Storage Extension endpoint.

Within a multisite architecture, you can configure VMware Cloud Director Object Storage Extension instances with a standalone virtual data center in each site. The following diagram illustrates the architecture.

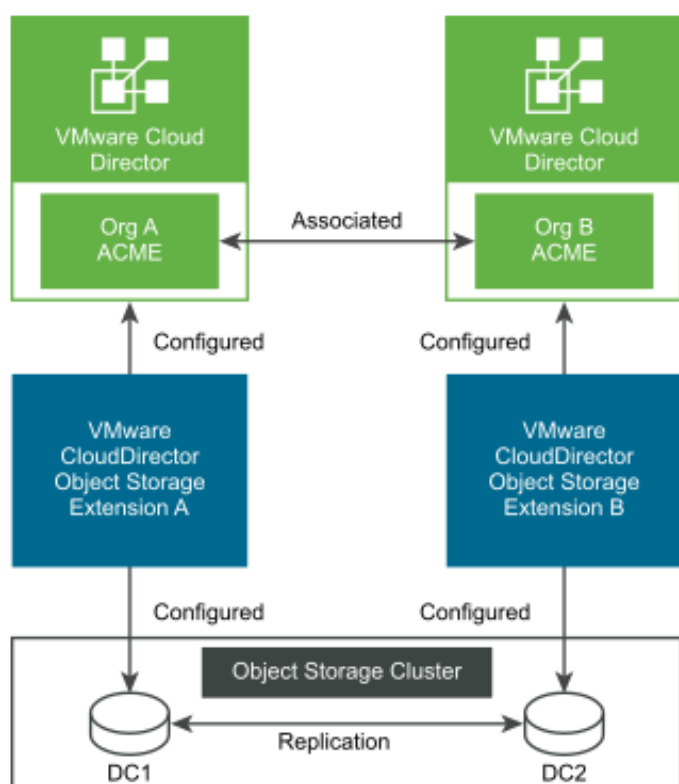


Figure 13. OSE Multisite Architecture: Single S3 Cluster for Multiple DCs

You can also configure VMware Cloud Director Object Storage Extension instances in different sites to use a single virtual data center. The following diagram illustrates the architecture.

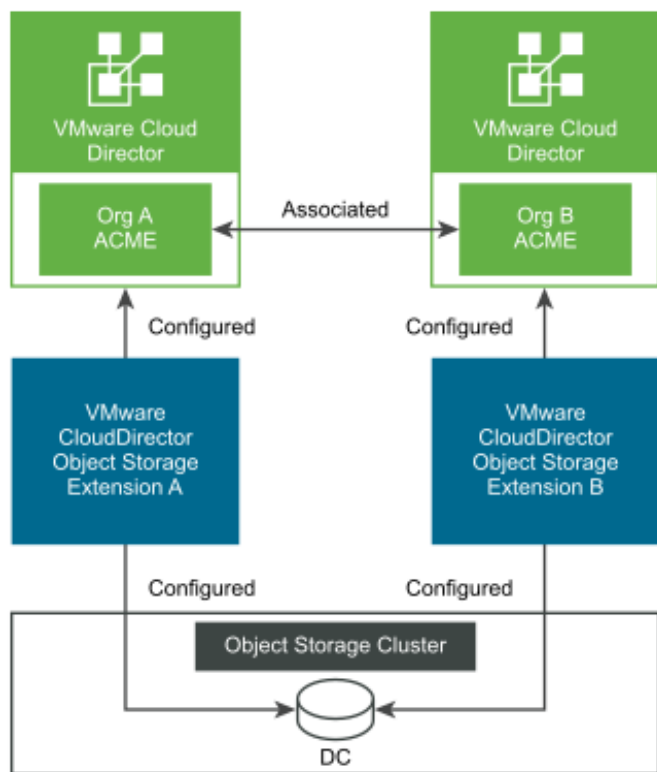


Figure 14: OSE Multisite Architecture: Single S3 Cluster for a Single DC

When you configure the multisite feature, you create a cluster of multiple VMware Cloud Director Object Storage Extension instances to create an availability zone. You can group the VMware Cloud Director Object Storage Extension instances together only in a single region. A region is a collection of compute resources in a geographic area. Regions are isolated and independent of one another. VMware Cloud Director Object Storage Extension does not support multi-region architectures.

You can share the same buckets and objects across tenant organizations within a multisite environment. To share buckets and objects across sites, map all tenant organizations to the same storage group. See [Edit Tenant Mapping Configuration](#).

Test Environment Reference Benchmark

Cloudian HyperStore Test Setup

A production grade setup of VMware Cloud Director and VMware Cloud Director Object Storage Extension was deployed in the lab. Both VMware Cloud Director and Object Storage Extension were deployed in a three-node configuration. The object storage platform also consisted of three-load balanced hardware appliances Cloudian Hyperstore 1508. The workloads were simulated by three VM nodes running COSBench software - the industry-standard benchmark tool for object storage. The effect of the front-end load balancer on the test results was eliminated by connecting each COSBench node to one Object Storage Extension node. Cloudian HyperStore nodes were load-balanced with NSX-V Load Balancer in L4 TCP accelerated mode.

To assess the impact of OSE proxying of S3 APIs, the same tests were performed directly to the Cloudian HyperStore (through the load balancer). The following diagram shows the network flows of the S3 API communication.

Note that HTTPS was used both for frontend traffic (COSBench to Object Storage Extension nodes) and backend traffic (Object Storage Extension to Cloudian HyperStore or COSBench to Cloudian HyperStore).

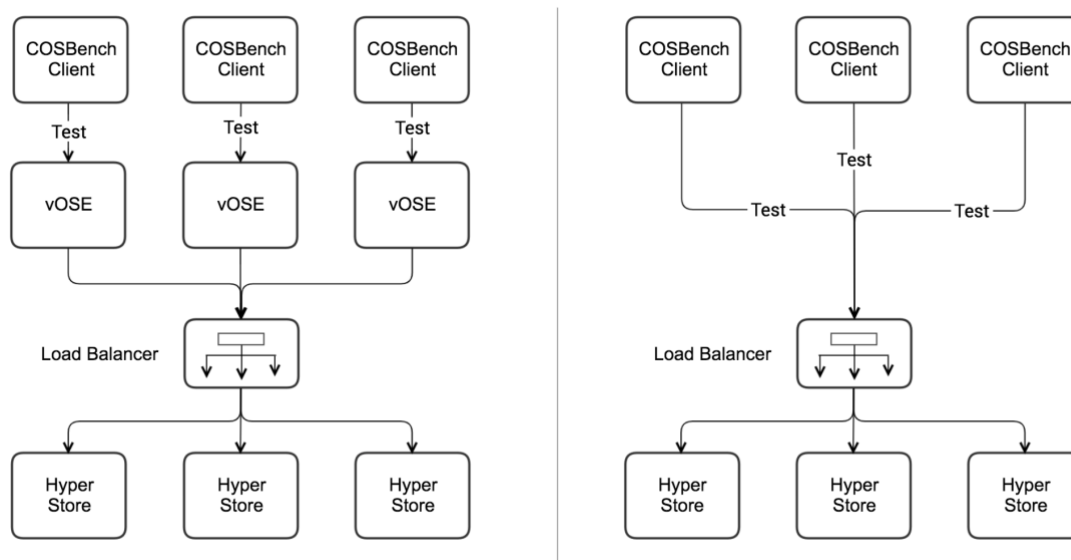


Figure 15. Cloudian HyperStore Test Topology

Cloudian HyperStore - Bill of Materials

The following table lists the software and hardware components used to collect performance results for VMware Cloud Director Object Storage Extension.

Table 1. Bill of Material

Component	Count	Specifications	Notes
vCenter Server	1	6.7.0U2 (14070654)	
ESXi	4	6.7.0U2 (13006603)	Dell PowerEdge R730xd Intel Xeon E5-2620, 128 GB RAM, NFS storage 5 TB
NSX-V	1	6.4.5 (13282012)	ESG LB: Quad-Large, accelerated SSL pass through
VMware Cloud Director cells	3	10.0.0.0 (14638910)	Appliance deployment (2 CPU, 12 GB RAM, 132 GB HDD)
NFS Transfer Share	1	CentOS 7	CentOS 7 VM (2 vCPU, 4 GB RAM, 532 GB HDD)
Object Storage Extension (OSE) nodes	3	2.0	CentOS 7 VM (8 vCPUs, 8 GB RAM, 128 GB HDD)
PostgreSQL	1	10	Could be separate or part of VMware Cloud Director installation
Cloudian HyperStore	3	7.1.6	Hardware appliance: Hyperstore 1508
COSBench	3	0.4.2	Ubuntu VM (6 vCPUs, 8 GB RAM, 240 GB HDD)

Cloudian HyperStore Test Results

Scenario 1 – Large Objects

Workloads: 100 workers doing writes and reads to 25 buckets with 10 MB objects

Step 0: Prepare data for reading

Step 1: Write for 5 mins

Step 2: Read for 5 mins

Step 3: Delete for 5 mins

Step 4: Clean up all buckets and objects

Table 2. Cloudian HyperStore - HTTPS Write/Read of 10 MB Objects by 100 Workers across 25 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE	Write	2030.84 ms	49.24 op/s	492.37 MB/S
Cloudian	Write	1889.99 ms	52.91 op/s	529.06 MB/S
OSE	Read	1681.79 ms	59.46 op/s	594.59 MB/S
Cloudian	Read	1626.88 ms	61.47 op/s	614.65 MB/S

Scenario 2 – Concurrency Comparison

Workloads: Write, read, and delete of objects size 100 MB for different concurrency level (10 – 200 workers)

Step 0: Prepare data for reading

Step 1: Write for 5 mins

Step 2: Read for 5 mins

Step 3: Delete for 5 mins

Step 4: Clean up all buckets and objects

Table 3. HTTPS 100 MB Objects with Various Concurrency 10, 50, 100 and 200 Workers

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 10 W	Write	5680.91 ms	1.76 op/s	175.91 MB/S
Cloudian: 10 W	Write	4552.46 ms	2.2 op/s	219.62 MB/S
OSE: 10 W	Read	5228.09 ms	1.91 op/s	191.24 MB/S
Cloudian: 10 W	Read	3245.85 ms	3.08 op/s	308.02 MB/S
OSE: 50 W	Write	10175 ms	4.91 op/s	491.35 MB/S
Cloudian: 50 W	Write	10115.48 ms	4.94 op/s	494.11 MB/S
OSE: 50 W	Read	9214.33 ms	5.42 op/s	542.48 MB/S
Cloudian: 50 W	Read	8858.65 ms	5.64 op/s	564.31 MB/S
OSE: 100 W	Write	23075.21 ms	4.33 op/s	432.85 MB/S
Cloudian: 100 W	Write	19947.78 ms	5.01 op/s	501.47 MB/S
OSE: 100 W	Read	18005.1 ms	5.54 op/s	553.96 MB/S
Cloudian: 100 W	Read	16472.98 ms	6.07 op/s	606.87 MB/S
OSE: 200 W	Write	38917.25 ms	5.13 op/s	513.18 MB/S
Cloudian: 200 W	Write	37003.96 ms	5.4 op/s	540.18 MB/S
OSE: 200 W	Read	33042.4 ms	6.06 op/s	605.57 MB/S
Cloudian: 200 W	Read	31668.87 ms	6.32 op/s	631.83 MB/S

Scenario 3 – Small Objects

Workloads: Write, read and delete for object size 1 MB with 1 00 workers across 30 buckets

Step 1: 30 buckets with each bucket having 100,000 objects

Step 2: Write for 1 hour and read for 1 hour

Step 3: Clean up all buckets and objects

Table 4. Read and Write of 1 MB Objects by 100 Workers across 30 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 100 W	Read	190.45 ms	525 op/s	525 MB/S
Cloudian: 100 W	Read	174.99 ms	571.22 op/s	571.22 MB/S
OSE: 100 W	Write	273.56 ms	365.38 op/s	365.38 MB/S
Cloudian: 100 W	Write	204.81 ms	487.85 op/s	487.85 MB/S

Scenario 4 – Object Size Comparison

Workloads: Write, read and delete for various object sizes ranging from 1 MB – 1 GB with 100 workers across 100 buckets

Step 1: Create 100 buckets with each bucket having 25 objects

Step 2: Do 1000 write operations

Step 3: Do 1000 read operations

Step 4: Clean up all objects and buckets

Table 5. HTTPs 1 MB – 1 GB Objects with Concurrency of 100 Workers

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 1 MB	Write	281.34 ms	368.66 op/s	368.66 MB/S
Cloudian: 1 MB	Write	252.05 ms	403.48 op/s	403.48 MB/S
OSE: 1 MB	Read	222.44 ms	463.3 op/s	463.3 MB/S
Cloudian: 1 MB	Read	198.74 ms	516.54 op/s	516.54 MB/S
OSE: 10 MB	Write	2204.59 ms	46.39 op/s	463.92 MB/S
Cloudian: 10 MB	Write	2082.8 ms	48.25 op/s	482.46 MB/S
OSE: 10 MB	Read	1847.61 ms	54.84 op/s	548.43 MB/S
Cloudian: 10 MB	Read	1750.51 ms	57.63 op/s	576.3 MB/S
OSE: 100 MB	Write	25137.85 ms	4.06 op/s	405.98 MB/S
Cloudian: 100 MB	Write	20184.66 ms	4.99 op/s	498.76 MB/S
OSE: 100 MB	Read	17110.77 ms	5.9 op/s	590.06 MB/S
Cloudian: 100 MB	Read	16932.9 ms	5.95 op/s	594.66 MB/S
OSE: 1 GB	Write	208321.21 ms	0.48 op/s	483.63 MB/S
Cloudian: 1 GB	Write	193002.93 ms	0.52 op/s	519.35 MB/S
OSE: 1 GB	Read	178607.12 ms	0.56 op/s	563.06 MB/S
Cloudian: 1 GB	Read	168282.92 ms	0.6 op/s	595.05 MB/S

Conclusion

As can be seen from the above test results VMware Cloud Director Object Storage Extension performance is much in line with the pure storage platform performance and does not add significant overhead for object sizes which are greater than 1 MB with maximums around 5 - 15%.

Dell EMC ECS Test Setup

The VMware Cloud Director Object Storage Extension allows VMware Cloud Providers who are using VMware Cloud Director to offer object storage services to their end users. The extension acts as middleware which is tightly integrated with Cloud Director to abstract 3rd party S3 API compatible storage providers in a multi-tenant fashion.

In this test setup, Object Storage Extension was deployed in five-node configuration. The object storage platform also consists of five load-balanced hardware appliances Dell EMC ECS. The workloads were simulated by three VM nodes running COSBench software, which is the industry standard benchmark tool for object storage. In order to assess the impact of the Storage Extension proxying of S3 APIs, the same tests were performed directly to the ECS nodes (through the load balancer). The following diagrams show the network flows of the S3 API communication.




Figure 16. Dell EMC ECS Test Topology

Dell EMC ECS - Bill of Materials

The following table lists the software and hardware components that were used to collect performance results for VMware Cloud Director Object Storage Extension.

Table 6. Bill of Materials

Component	Count	Specifications	Notes
vCenter Server	1	6.7.0U2 (14070654)	
ESXi	4	6.7.0U2 (13006603)	Dell PowerEdge R730xd Intel Xeon E5-2620, 128 GB RAM, NFS storage 5 TB
NSX-V	1	6.4.5 (13282012)	ESG LB: Quad-Large, accelerated SSL pass through
VMware Cloud Director cells	3	10.0.0.0 (14638910)	Appliance deployment (2 CPU, 12 GB RAM, 132 GB HDD)
NFS Transfer Share	1	CentOS 7	CentOS 7 VM (2 vCPU, 4 GB RAM, 532 GB HDD)
VMware Cloud Director Object Storage Extension nodes	3	2.0	CentOS 7 VM (8 vCPUs, 8 GB RAM, 128 GB HDD)

PostgreSQL	1	10	Could be separate or part of VMware Cloud Director installation
Dell EMC ECS	5	3.4	
COSBench	3	0.4.2	Ubuntu VM (6 vCPUs, 8 GB RAM, 240 GB HDD)

Dell EMC ECS – Test Results

Scenario 1 – Large Objects

Workloads: 100 workers doing writes and reads to 25 buckets with 10 MB objects

Step 0: Prepare data for read

Step 1: Write for 5 mins

Step 2: Read for 5 mins

Step 3: Delete for 5 mins

Step 4: Clean up all buckets and objects

Table 7. Dell EMC ECS - HTTPs 10 MB Objects with Concurrency of 100 Workers across 25 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE	Write	2654.28 ms	37.66 op/s	376.58 MB/S
ECS	Write	2005.08 ms	49.87 op/s	498.69 MB/S
OSE	Read	2432.37 ms	41.09 op/s	410.85 MB/S
ECS	Read	1677.03 ms	59.62 op/s	596.24 MB/S

Scenario 2 – Various Object Sizes Concurrency Comparison

Workloads: Write, read and delete for object size 100 MB for different concurrency level (10 – 100 workers)

Step 0: Prepare data for read

Step 1: Write for 5 mins

Step 2: Read for 5 mins

Step 3: Delete for 5 mins

Step 4: Clean up all buckets and objects

Table 8. Dell EMC ECS - HTTPs 100 MB Objects with Concurrency of [10-100] Workers

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 10 W	Write	5094.65 ms	1.96 op/s	196.28 MB/S
ECS: 10 W	Write	4573.61 ms	2.19 op/s	218.6 MB/S
OSE: 10 W	Read	4886.07 ms	2.05 op/s	204.62 MB/S
ECS: 10 W	Read	3596.51 ms	2.78 op/s	278.09 MB/S
OSE: 50 W	Write	13489.81 ms	3.71 op/s	370.61 MB/S
ECS: 50 W	Write	10108.52 ms	4.95 op/s	494.56 MB/S
OSE: 50 W	Read	11024.32 ms	4.51 op/s	451.12 MB/S
ECS: 50 W	Read	8600.87 ms	5.81 op/s	581.18 MB/S
OSE: 100 W	Write	26000.98 ms	3.79 op/s	379.02 MB/S
ECS: 100 W	Write	23602.16 ms	4.22 op/s	422.21 MB/S
OSE: 100 W	Read	23379.02 ms	4.18 op/s	418.45 MB/S
ECS: 100 W	Read	19714.67 ms	5.07 op/s	507.1 MB/S

Scenario 3 – Small Objects

Workloads: Write, read and delete for object size 4 KB with 200 workers across 30 buckets

Step 1: 30 buckets with each bucket having 10000 objects

Step 2: 50% read, 50% write for 1 hour

Step 3: Clean up all buckets and objects

Table 9. Dell EMC ECS - HTTPs 4 MB Objects with Concurrency of 100 Workers across 30 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 100 W	Read	15.82 ms	727.19 op/s	2.91 MB/S
ECS: 100 W	Read	13.83 ms	765.85 op/s	3.06 MB/S
OSE: 100 W	Write	121.11 ms	727.07 op/s	2.91 MB/S
ECS: 100 W	Write	116.56 ms	766.28 op/s	3.07 MB/S

Scenario 4 – Object Size Comparison

Workloads: Write, read and delete for various object sizes ranging from 1 MB – 1 GB with 100 workers across 100 buckets

Step 1: 100 buckets with each bucket having 25 objects

Step 2: Do 1000 write operations

Step 3: Do 1000 read operations

Step 4: Clean up all objects and buckets

Table 10. Dell EMC ECS - HTTPs 1 MB – 1GB Objects with Concurrency of 100 Workers across 100 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 1 MB	Write	327.92 ms	335.74 op/s	335.74 MB/S
ECS: 1 MB	Write	257.65 ms	399.24 op/s	399.24 MB/S
OSE: 1 MB	Read	234.94 ms	465.05 op/s	465.05 MB/S
ECS: 1 MB	Read	206.63 ms	503.81 op/s	503.81 MB/S
OSE: 10 MB	Write	2387.46 ms	42.72 op/s	427.25 MB/S
ECS: 10 MB	Write	2048.71 ms	49.03 op/s	490.28 MB/S
OSE: 10 MB	Read	2075.17 ms	50.59 op/s	505.95 MB/S
ECS: 10 MB	Read	1684.57 ms	59.45 op/s	594.54 MB/S
OSE: 100 MB	Write	24580.03 ms	4.22 op/s	421.51 MB/S
ECS: 100 MB	Write	19674.93 ms	5.1 op/s	510.33 MB/S
OSE: 100 MB	Read	21524.92 ms	4.73 op/s	472.63 MB/S
ECS: 100 MB	Read	16287.79 ms	6.18 op/s	617.89 MB/S
OSE: 1 GB	Write	235635.04 ms	0.45 op/s	445.13 MB/S
ECS: 1 GB	Write	198951.58 ms	0.5 op/s	504.42 MB/S
OSE: 1 GB	Read	205287.62 ms	0.51 op/s	507.21 MB/S
ECS: 1 GB	Read	163421.88 ms	0.62 op/s	615.88 MB/S

Conclusion

As can be seen from the above test results, VMware Cloud Director Object Storage Extension performance is much in line with the pure storage platform performance and adds overhead with maximums around 5-25%.

AWS S3 Test Setup

VMware Cloud Director Object Storage Extension was deployed in three-node configuration. The object storage platform is Amazon Simple Storage Service (Amazon S3). The workloads were simulated by three VM nodes running COSBench software, which is the industry standard benchmark tool for object storage. To assess the impact of VMware Cloud Director Object Storage Extension proxying of S3 APIs, the same tests were performed directly with the AWS S3 service. The following diagrams show the network flows of the S3 API communication.

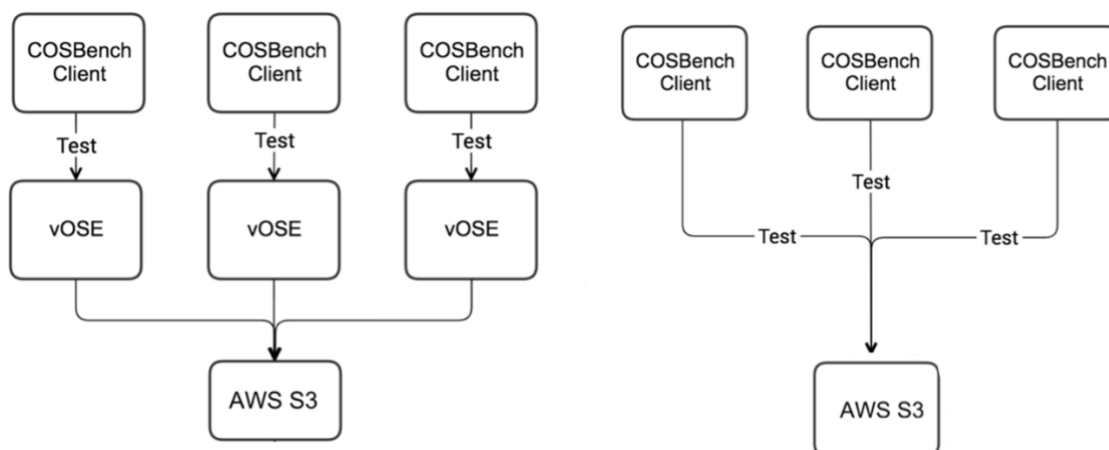


Figure 17. AWS S3 Test Topology

AWS S3 - Bill of Materials

The following table lists the software and hardware components that were used to collect performance results for VMware Cloud Director Object Storage Extension.

Table 11. Bill of Materials

Component	Count	Specifications	Notes
vCenter Server	1	6.7.0U2 (14070654)	
ESXi	4	6.7.0U2 (13006603)	Dell PowerEdge R730xd Intel Xeon E5-2620, 128 GB RAM, NFS storage 5 TB
NSX-V	1	6.4.5 (13282012)	ESG LB: Quad-Large, accelerated SSL pass through
VMware Cloud Director cells	3	10.0.0.0 (14638910)	Appliance deployment (2 CPU, 12 GB RAM, 132 GB HDD)
NFS Transfer Share	1	CentOS 7	CentOS 7 VM (2 vCPU, 4 GB RAM, 532 GB HDD)
VMware Cloud Director Object Storage Extension nodes	3	2.0	CentOS 7 VM (8 vCPUs, 8 GB RAM, 128 GB HDD)
PostgreSQL	1	10	Could be separate or part of VMware Cloud Director installation
AWS S3	1		
COSBench	3	0.4.2	Ubuntu VM (6 vCPUs, 8 GB RAM, 240 GB HDD)

AWS S3 – Test Results

Scenario 1 – Large Objects

Workloads: 100 workers doing writes and reads to 25 buckets with 10 MB objects

Step 0: Prepare data for read

Step 1: Write for 5 mins

Step 2: Read for 5 mins

Step 3: Delete for 5 mins

Step 4: Clean up all buckets and objects

Table 12. AWS- HTTPs 10 MB Objects with Concurrency of 100 Workers across 25 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE	Write	3219.13 ms	31.05 op/s	310.53 MB/S
AWS	Write	3213.18 ms	31.12 op/s	311.23 MB/S
OSE	Read	3614.7 ms	27.66 op/s	276.65 MB/S
AWS	Read	3189.3 ms	31.35 op/s	313.51 MB/S

Scenario 2 – Concurrency Comparison

Workloads: Write, read and delete for object size 100 MB for different concurrency level (10 – 200 workers)

Step 0: Prepare data for read

Step 1: Write for 5 mins

Step 2: Read for 5 mins

Step 3: Delete for 5 mins

Step 4: Clean up all buckets and objects

Table 13. AWS- HTTPs 100 MB Objects with Concurrency of [10-200] Workers

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 10 W	Write	3585.87 ms	2.79 op/s	278.56 MB/S
AWS: 10 W	Write	3211.07 ms	3.11 op/s	311.18 MB/S
OSE: 10 W	Read	3330.07 ms	3 op/s	300.09 MB/S
AWS: 10 W	Read	3186.34 ms	3.13 op/s	313.37 MB/S
OSE: 50 W	Write	16123.54 ms	3.1 op/s	309.76 MB/S
AWS: 50 W	Write	15941.99 ms	3.13 op/s	313.31 MB/S
OSE: 50 W	Read	17545.01 ms	2.85 op/s	284.77 MB/S
AWS: 50 W	Read	15839.36 ms	3.15 op/s	315.31 MB/S

OSE: 100 W	Write	31723.52 ms	3.14 op/s	313.93 MB/S
AWS: 100 W	Write	31557.58 ms	3.16 op/s	316 MB/S
OSE: 100 W	Read	33318.05 ms	3 op/s	300.2 MB/S
AWS: 100 W	Read	31271.15 ms	3.2 op/s	320.12 MB/S
OSE: 200 W	Write	63366.13 ms	3.15 op/s	315.4 MB/S
AWS: 200 W	Write	62448.93 ms	3.19 op/s	319.13 MB/S
OSE: 200 W	Read	70461.69 ms	2.84 op/s	284.14 MB/S
AWS: 200 W	Read	62421.21 ms	3.2 op/s	320.35 MB/S

Scenario 3 – Small Objects

Workloads: Write, read and delete for object size 1 MB with 1 00 workers across 30 buckets

Step 1: 30 buckets with each bucket having 100,000 objects

Step 2: Write for 1 hour and read for 1 hour

Step 3: Clean up all buckets and objects

Table 14. AWS- HTTPs 1 MB Objects with Concurrency of 100 Workers across 30 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 100 W	Write	370.41 ms	269.91 op/s	269.91 MB/S
AWS: 100 W	Write	369.81 ms	270.35 op/s	270.35 MB/S
OSE: 100 W	Read	369.21 ms	270.83 op/s	270.83 MB/S
AWS: 100 W	Read	367.04 ms	272.43 op/s	272.43 MB/S

Scenario 4 – Object Size Comparison

Workloads: Write, read and delete for various object sizes ranging from 1 MB – 1 GB with 100 workers across 100 buckets

Step 1: 100 buckets with each bucket having 25 objects

Step 2: Do 1000 write operations

Step 3: Do 1000 read operations

Step 4: Clean up all objects and buckets

Table 15. AWS- HTTPs 1 MB -1 GB Objects with Concurrency of 100 Workers across 100 Buckets

Test Type	Operation	Avg Response Time	Throughput	Bandwidth
OSE: 1 MB	Write	286.64 ms	352.07 op/s	352.07 MB/S
AWS: 1 MB	Write	255.88 ms	401.95 op/s	401.95 MB/S
OSE: 1 MB	Read	278.63 ms	361.63 op/s	361.63 MB/S

AWS: 1 MB	Read	243.58 ms	415.37 op/s	415.37 MB/S
OSE: 10 MB	Write	2789.34 ms	36.7 op/s	366.99 MB/S
AWS: 10 MB	Write	2509.52 ms	42 op/s	420 MB/S
OSE: 10 MB	Read	2735.01 ms	36.63 op/s	366.26 MB/S
AWS: 10 MB	Read	2604.47 ms	38.6 op/s	386.04 MB/S
OSE: 100 MB	Write	31179 ms	3.27 op/s	327.19 MB/S
AWS: 100 MB	Write	30227.8 ms	3.49 op/s	348.67 MB/S
OSE: 100 MB	Read	28733.64 ms	3.56 op/s	355.58 MB/S
AWS: 100 MB	Read	28201.35 ms	3.69 op/s	368.57 MB/S
OSE: 1 GB	Write	295057.91 ms	0.35 op/s	350.96 MB/S
AWS: 1 GB	Write	295723.67 ms	0.35 op/s	346.39 MB/S
OSE: 1 GB	Read	330634.61 ms	0.32 op/s	323.51 MB/S
AWS: 1 GB	Read	291255.77 ms	0.35 op/s	350.39 MB/S

Conclusion

As can be seen from the test results above, VMware Cloud Director Object Storage Extension performance is much in line with the pure storage platform performance and does not add significant overhead with maximums around 5- 13%.

Abbreviations

OSE	VMware Cloud Director Object Storage Extension
OSIS	Object Storage Interoperability Service
IaaS	Infrastructure as a Service
Cloud Director	VMware Cloud Director



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 vmware.com Copyright © 2020 VMware, Inc.
All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.