

Enterprise Object Storage Performance Benefits of vADC Load Balancing

**Authors: Russ Fellows and
Mohammad Rabin**

March 2019



Executive Summary

Object storage has emerged as the preferred method of storing and retaining large data repositories, which may be utilized for a variety of purposes, including data mining, artificial intelligence training data, or mandated data retention. Object storage systems are able to scale to multiple petabytes of data, with billions of objects without performance degradation.

The use cases for Object storage systems differ from traditional application workloads due to the nature of their design and the requirement for close application interaction with these systems. Object storage systems maintain limited data about each object, which requires applications to maintain data via other mechanisms. Additionally, these applications and storage systems are designed for massive scale with deployments using on premise clouds, public clouds or some combination of both.

Hitachi Content Platform (HCP) is an object-based storage system designed for long-term archiving or retention of fixed content data. HCP supports configurations up to multiple Exabyte capacity. HCP's scalability includes multitenant management, delegation and provisioning features, along with scalable front-end and back-end storage components to enable the scaling of object access and storage capacity independently. Additionally, a number of HCP data services exist, including capacity balancing, compression, garbage collection, replication, storage tiering and geo-distribution of data.

However, with multiple object processing or front-end nodes, applications may overload a particular node while leaving other nodes idle. An important consideration in achieving optimal performance is balancing workloads across the entire object storage deployment. Workload variations may arise from a few high-demand applications or may be due to object-based applications serving as a repository for hundreds or thousands of applications. In both instances a web application load balancer can help alleviate hot-spots or overloading of individual nodes and evenly distribute workloads across systems in an object storage cluster.

In this Lab Insight, Evaluator Group analyzed HCP using multiple different workloads and configurations, with a primary focus on evaluating the effect of a vADC load-balancer with each configuration.

The benefits of using vADC included:

- **Performance 3X – 4X** that of systems without vADC for single applications (200% - 300% increase)
- Increased object access rates for all tests, measured by the number of objects accessed per second
- Increased throughput rates for all tests, measured by data transfer rate per second

Additionally, Evaluator Group Research offers an extensive analysis of HCP's capabilities, which is available for their subscribers at www.evaluatorgroup.com

Evaluation Overview

Hitachi Content Platform is an object storage system with a modular design, enabling the use of multiple object processing or “front-end” nodes along with multiple “back-end” storage nodes. The system tested was an HCP G10 system, which supports configurations from 4 to 80 front-end nodes and up to 80 back-end storage nodes. HCP front-end nodes may be utilized without captive storage, enabling the use of networked storage for back-end storage. The test configuration consisted of 8 “front-end” HCP G10 nodes along with 8 “back-end” HCP S30 storage nodes. This configuration of 8 application “front-end” nodes and 8 storage capacity “back-end” nodes is referred to as the HCP system.

The HCP system also provides a number of advanced features, including compression, single instance storage of objects, data encryption, versioning and archival protection, search and a number of other capabilities. The focus of this evaluation was to specifically ascertain the benefit of utilizing a web application load-balancer, specifically the vADC load-balancer for a variety of workloads.

Configurations Evaluated

The tests were designed to determine the effectiveness of vADC with HCP using the IOmark-OBJ benchmarking tool. An 8-node HCP configuration was tested with a variety of access profiles and number of clients. The focus was on testing workloads with a vADC load balancer without using DNS redirection for load-balancing. Testing utilized the following:

- 8 HCP G10 nodes with 8-enclosure HCP S30 storage node
- Testing utilized configurations with vADC and without vADC
- Test access with 4-clients and 1-client configurations
- Multiple access profiles (Get and Put operations) and object sizes

Testing focused on measuring performance gains including bandwidth, throughput, response time and scalability.

Measuring Performance

The IOmark-OBJ tool utilizes configuration files to specify the size of objects generated and the percentage of access protocols. Similar to block storage, object storage may either be written with a “Put” operation, or read using a “Get” operation.

Several different object sizes were utilized, along with a 100% write or “Put” workload and a 100% read or “Get” workload. Small objects were defined as 1 MB or smaller, with large objects defined as being larger than 1 MB. A total of 2 different hardware configurations were tested with 6 different object sizes for both read and write testing, producing a total of 24 test configurations. Each configuration was tested using multiple number of access rates for both Put and Get operations, yielding approximately 100 test configurations in total.

Workload configurations included:

- Two configurations:
 - HCP object access via a vADC load-balancer
 - HCP object access without the vADC load-balancer
- Object Size:
 - Large Objects : 1 MB, 10 MB, 100 MB and 1 GB
 - Small Objects: 4KB, 64KB and 1 MB (note 1 MB shown for both)
- Multiple outstanding operations
 - The number of operations or threads varied from 8 per IOmark-OBJ instance up to 256, depending upon the size of the object (large objects utilized fewer threads)
- 100% “Put” and 100% “Get” workloads
- Two different client configurations
 - 4-client configuration (i.e., 4 workload drivers all accessing the target HCP object store)
 - 1-client configuration (a single driver or application accessing the target HCP object store)




Figure 1: Test Environment Configuration

Test Results

As indicated previously, testing was performed utilizing the IOmark-OBJ tool to generate workloads. This tool was developed specifically to test object storage systems utilizing a variety of object sizes and request rates across multiple workload drivers, all accessing the same object storage target. The reported results indicate the total, or aggregate of all workload drivers running during the testing period.

Note: All test results provided have been normalized, meaning that all values provided have been set to the same scale. As a result, the absolute values are not pertinent, only the relative differences between results. By normalizing results, focus is removed from the absolute results and instead placed upon the area of testing, namely the performance differences between configurations that utilized a vADC and those that did not. All graphs and results shown utilized the first data point as the “normalization” value, with the factor set to 100. Moreover, all graphs and reported values represent the first data point with a value of “100” with all other data points normalized accordingly.

Test Summary

As shown below in Table 1, the benefits of utilizing a vADC load-balancer rather than using DNS lookup for access to multiple nodes were significant regardless of access methods, such as “Get” or “Put” operations or the size of the objects. For large-scale workloads, utilizing 4 clients for simultaneous access, the level of improvement was typically around 20% better with vADC. The exception to this is for “Put” or write operations with small objects under 1 MB, for which the differences were negligible.

For smaller applications, as represented by the 1-client workloads, the improvement was significantly greater, with improvements of between 2X and 3X as measured by bandwidth for workloads above 1 MB object size.

Test Cases	Object Sizes	Performance Metric	Average Increase
8 Node HCP – 4 Clients 100% “Get” Read	1 MB - 1 GB	Bandwidth	14.7%
8 Node HCP – 4 Clients 100% “Get” Read	4 KB – 1 MB	Ops / sec.	28.8%
8 Node HCP – 4 Clients 100% “Put” Write	1 MB - 1 GB	Bandwidth	21.3%
8 Node HCP – 4 Clients 100% “Put” Write	4 KB – 1 MB	Ops / sec.	1.9%
8 Node HCP – 1 Client 100% “Get” Read	1 MB - 1 GB	Bandwidth	291.6%
8 Node HCP – 1 Client 100% “Put” Write	1 MB - 1 GB	Bandwidth	229.5%

Table 1: Summary of vADC Performance Benefits

Large Object Tests With 1 Client

Below in Figures 2 and 3 are the results of 8-node tests utilizing 1 client with large object transfer workloads. These tests are typical of workloads that are designed to maximize the total throughput capability of a single object storage application. Although the client workload itself utilized multiple threads for access to the storage system, the test itself was not designed to saturate the object storage system. The client workload saturated its resources when accessing the HCP cluster while still maintaining reasonable response times.

When testing an 8-node HCP configuration using 1 client, a significant performance increase occurred when utilizing the vADC load-balancer rather than DNS load-balancing for all object sizes. As seen below in Figure 2, the system throughput averaged more than 4X that of systems without vADC for all read or "Get" testing with objects above 1 MB in size.




Figure 2: Normalized - 1 Client "Get" with/without vADC

As shown in Figure 3, for the 8-node HCP configuration using 1 client, "Put" performance was also significantly greater when utilizing the vADC load-balancer rather than DNS load-balancing for all object sizes. The system throughput was more than 4X that of systems without vADC for all read or "Get" testing with objects above 10 MB in size.




Figure 3: Normalized - 1 Client “Put” with/without vADC

Evaluator Group comments: *The performance for a single client accessing an HCP cluster via a vADC load-balancer is 3X – 4X that of configurations without the load balancer, with gains of more than 200% improvement in throughput and with lower response times. These results are likely typical of most applications utilizing an object repository that are not designed to utilize the entire object cluster resources. Even when the HCP cluster was being utilized to its full potential, an additional 10 – 20% performance increase is possible by using a vADC load-balancer.*

Large Object Tests

Below in Figures 4 and 5 are the results of 8-node tests utilizing 4 clients with large object transfer workloads. These tests are typical of workloads that are designed to maximize the total throughput capability of a given HCP configuration. That is, the tests were designed to completely saturate the HCP cluster while still maintaining reasonable response times.

Using vADC was found to improve read “Get” operations performance for all large object size workloads, even when the HCP system was operating near its limits as shown below in Figure 4.




Figure 4:Normalized - Large object "Get" with/without vADC

Using vADC was found to improve write "Put" operations performance for all large object size workloads, even when the HCP system was operating near its limits as shown below in Figure 5.




Figure 5: Normalized - Large object "Put" with/without vADC

Small Object Tests

Testing utilized small object sizes of 4KB, 64KB and 1 MB, utilizing 4 multi-threaded clients to access an 8-node HCP cluster. Using vADC was found to improve performance in terms of the number of object operations per second. Results for “Get” and “Put” operations are shown in Figures 6 and 7, respectively.




Figure 6: Normalized - Small object “Get” with/without vADC




Figure 7: Normalized - Small object “Put” with/without vADC

HCP Object Storage System Overview

Hitachi Content Platform (HCP) is an object storage solution that enables object-based applications and devices to natively access a scalable storage platform. HCP enables IT organizations and cloud service providers to securely and cost effectively store, share, protect, preserve, and analyze data. Beyond efficiency, and ability to store data at massive scale, HCP provides automation for IT operations like data governance and protection with the ability to modify rules for regulatory compliance, applications and technologies over the life of data. HCP also automates the governance of data to ensure proper retention, access control, encryption and disposal of data as well as simplifying e-discovery and search. In IT environments where data grows quickly or must live for years or even indefinitely, these capabilities are invaluable.

HCP can be delivered in a variety of deployment models ranging from a fully integrated appliance, nodes connected to arrays, commodity based, software only, and more. As an appliance, HCP is available with pre-configured access node (HCP G series node) and dense storage nodes (HCP S series node), both of which feature self-contained disk storage. HCP access nodes can also connect to one or more Hitachi arrays simultaneously as storage targets. The HCP architecture allows customers to scale capacity and performance independently, which provides flexibility to support a wide range of workloads.

HCP can handle all kinds of data and almost any application. It offers high reliability, massive scale, seamless data mobility and storage across private clouds and public cloud services, encryption, access control, easy provisioning, chargeback measurement and more. The HCP G series access nodes allow organizations greater flexibility to support mixed workloads with varying performance and scale requirements. These nodes virtualize capacity from HCP S series nodes, local drives, Fibre Channel storage arrays, NFS shares and leading public cloud providers. HCP drastically reduces total cost of ownership and provides cost-effective storage with geographically dispersed erasure coding data protection for content that must remain behind the firewall. Such attributes enable IT to take advantage of cloud and deliver a whole new range of IT services, without compromising security and control of information.

HCP scales by clustering nodes and storage independently. The minimum is four physical nodes with the ability to scale up to 80 nodes in the cluster. Data is distributed among the nodes in the cluster for load balancing. The data stored is referred to as objects, where each object is comprised of a file and the metadata associated with that file. The metadata includes information added by an application or system information added by HCP. Files stored in HCP are accessed over an Ethernet network with remote file system protocols.

Objects may be stored, viewed and retrieved using HTTP/REST to access a namespace within the HCP. Furthermore, HCP also provides native Amazon S3, OpenStack Swift NFS, SMB, SMTP and WebDAV protocol support. Objects can also be accessed by users or applications as a Uniform Resource Locator (URL) with access by HTTP/HTTPS.

The high density of HCP storage is enhanced with built-in compression, single instancing and support for a variety of media to keep storage costs in control. With dynamic data protection, data integrity checks, data retention enforcement, erasure coding and many other technologies to preserve and protect content, HCP delivers compliance-quality archiving and data protection.

HCP Characteristics

- **Scale** – Up to 800PB can be stored in the largest system (HCP S30 with 80 nodes), supporting billions of objects. HCP supports objects up to 5TB in size. Compression and single instancing are implemented to reduce the back-end storage capacity required.
- **Protection / Durability / Resiliency** – Protection from device failure uses the RAID protection from the attached storage system or erasure coding for the HCP S series storage nodes. Multiple copies across nodes can be made via synchronous remote replication to another HCP system.
- **Index and Search** – Content and metadata indexing with search is included using an embedded scale-out Lucene implementation.
- **Access Methods** – A native REST-based interface as well as Amazon S3 and OpenStack Swift compatible interfaces. HCP also supports the NFS, SMB, SMTP and WebDAV protocols, and offers dual-stack support for IPV4 and IPV6.
- **Geographic Access** – Geographic access is through the multi-copy or replication capability or through geo-dispersion using erasure codes.
- **Security and Compliance** – Multi-tenancy is supported with multiple system administrative and user access roles. Encryption of data at rest depends on attached storage with all communications encrypted using SSL.
- **Metadata** – System and user metadata is added to each individual object. System metadata includes policies and compliance controls. User or application metadata is written as a companion XML file, with user metadata embedded into each object.
- **Integrity and Verification** – A digital fingerprint is created using a selectable (at installation) hash code. The hash code is verified on retrieval.
- **Longevity of Object Data** – Migration of object files is controlled according to user-defined policies. Data can be transparently migrated between nodes in the cluster or else tiered to a choice of leading public cloud services, including Amazon S3, Google Cloud, Microsoft Azure and more. Nodes can have data drained and moved or redistributed automatically. Automatic repair of damaged objects is done by leveraging replicas, thus enabling long term storage of data.
- **Billing and Chargeback** – Both predefined reports and ad hoc reports may be generated. Data may also be exported for billing systems.

Hitachi Vantara has recently enhanced the HCP portfolio capabilities through acquisitions and has integrated the Hitachi Content Platform software capabilities into other offerings, resulting in a portfolio of object access products that includes HCP, HCP Anywhere edge model (cloud file gateway), HCP Anywhere (enterprise file sync and share), and Hitachi Content Intelligence (search, analytics and data quality).

Evaluation Summary

A growing number of applications are able to utilize object storage systems to retain and access large amounts of data. These modern “cloud-scale” applications are designed differently from more traditional applications, utilizing designs that incorporate a new type of storage, designed to scale beyond traditional storage systems that were not designed to work effectively with data sets or content that exceeds 1 PB of data.

The new object storage systems were designed to operate as content repositories and scale up to handle web-scale applications. With object storage systems becoming increasingly popular, it is important to understand how to optimize application access, in order to improve performance and ultimately deliver better application efficiency. One method of improving application response is to utilize an application load-balancer in between the applications and the object storage system, thereby balancing the workload across multiple nodes of the object store.

The vADC load-balancer was tested in a variety of configurations and workloads, with results showing that in most cases, large scale applications achieved more than a 20% performance increase. Perhaps most importantly, typical object applications were able to provide 3X – 4X, an increase of between 200% and 300%, with the average performance increase of approximately 290% for “Get” operations and approximately 230% for “Put” operations.

Applications that utilize object storage systems vary significantly, with some running as scale-out designs that utilize multiple systems for access to object storage. However, many object applications operate independently, thereby running on a single system. For these workloads in particular, an application load-balancer can provide significantly higher throughput and object access rates as shown by the testing conducted.

Business efficiency is important for all companies, regardless of application or size. However, as the scale of applications running on a single system and application-based data, grow, the benefit from increased efficiency can translate into significant savings. Efficiency gains of 20% - 30% when applied to systems consisting of tens of nodes and multiple Petabytes of capacity, represent a significant value. For single applications that experienced performance of 3X – 4X compared to configurations with vADC are very significant improvements. Clearly, utilizing a vADC application load-balancer can provide financial and performance benefits for a wide variety of workloads and should be a consideration when designing modern, cloud-scale applications.

About Evaluator Group

Evaluator Group Inc. is dedicated to helping IT professionals and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver in-depth, unbiased analysis on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and education. www.evaluatorgroup.com Follow us on Twitter @evaluator_group

Copyright 2019 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.

This document was developed with Hitachi Vantara funding. Although the document may utilize publicly available material from various vendors, including Hitachi, it does not necessarily reflect the positions of such vendors. Hitachi is a trademark or registered trademark of Hitachi, Ltd. Content Platform Anywhere is a trademark or registered trademark of Hitachi Vantara Corporation. The analysis in this document reflects Evaluator Group independent testing and analysis.