

The Role of Object Storage in Big Data Analytics

Prepared for RedHat
June 2016

TABLE OF CONTENTS

TABLE OF CONTENTS.....	2
INTRODUCTION.....	3
OBJECT STORES FOR BIG DATA.....	4
OBJECT STORAGE AND DATA LAKES.....	5
BOTTOM LINE.....	6
JUKU	7
WHY JUKU	7
AUTHOR.....	7

INTRODUCTION

Organizations of all sizes are experiencing a substantial diversification in the workloads served by their storage infrastructures. Consequently, requirements are no longer just about performance or capacity but about how storage resources are being consumed.

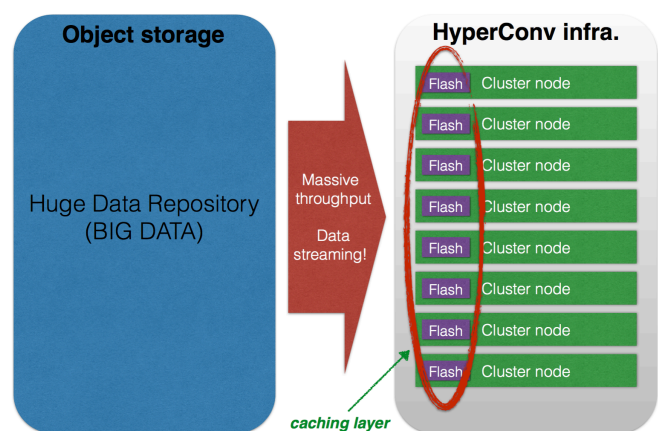
Mobile devices, Big Data applications, IoT are the most visible examples, they are drastically changing access patterns to large scale storage systems.

This is why Big Data analytics and, more precisely, in-memory applications, are posing another big challenge to traditional storage infrastructures where performance, throughput and scalability are involved. A tough problem to solve, especially since HDFS, a file system specifically designed to massively scale, is not suitable for all kinds of workloads and file sizes.

New radical infrastructure designs are surfacing and becoming more common, with memory-based ephemeral storage or sophisticated caching mechanisms used as first tier and coupled with massive capacity-driven data stores at the back-end. In this scenario, object storage (which is historically considered suitable only for cheap, durable and cold data archives), is in the right position to become the perfect foundation for this new class of storage. Now, throughput with scalability should not be an issue - S3/Swift APIs are supported by an increasing number of software vendors and HDFS/NFS/SMB gateways make it very easy to ingest data and remain compatible with legacy environments.

For object storage systems, the process of streaming data to the compute cluster is seamless, while storing original data sets or new results is less costly and more reliable than for any other kind of storage... and even more so now with the flexibility provided by cloud tiering - mechanisms implemented by most modern platforms.

Organizations of all sizes are experiencing substantial diversification in the workloads served by their storage infrastructure. Consequently, requirements are no longer just about performance or capacity but they depend on how storage resources are being consumed.



OBJECT STORES FOR BIG DATA

Object storage is no longer relegated to the second or third tier in storage infrastructures. Some of its characteristics are perfectly suited as the backend of modern IT infrastructures which need scalability, performance and availability demanded by developers, end users and cloud applications.

While many enterprises are embracing Big Data analytics, there are some important aspects to consider from the infrastructure standpoint:

- Even though HDFS is designed to scale, over time, cluster resources must be scaled differently than originally assumed. Storage could grow much more than compute or vice versa. And data retention can quickly lead to large capacity needs while size of active data sets remain quite unchanged.
- Furthermore, with consolidation in Big Data lakes, organizations are starting to use different compute resources: cloud-based and on-premises cluster with different applications.

Object storage is no longer to be considered only a second or third tier storage solution. Some of its characteristics are perfectly suitable to be the backend of modern IT infrastructures which need scalability, performance and security.

In this scenario, where flexibility and consolidation are key for top efficiency, two technologies are rapidly gaining attention for Big Data analytics: containers and in-memory databases.

- Big Data application containerization is happening: different application stacks use and access data differently and it is not unusual now to see activity running in real time alongside scheduled batch jobs. All of HDFS's limitations are exposed in these cases, especially when it comes to smaller files. An object storage system could be a good option here as it is much easier to access from remote compute resources, also enabling the expansion of the compute cluster using public cloud resources.
- Lately, in-memory databases have been chosen instead of Hadoop MapReduce-based application stacks for a greater number of use cases. This means the use of fewer large memory nodes. But the HDFS storage model, usually based on a large number of hard disks installed into the cluster nodes becomes unsustainable. These cluster nodes are alternatively filled with RAM and/or alternatively, NVMe devices acting as a ram extension, caching or first storage tier. As a direct consequence, data streaming capabilities provided by an external storage infrastructure become essential. Object storage is one of the best options in this case because of its ability to scale and sustain high throughputs at a relatively low cost while providing a highly reliable and durable persistent storage layer.

OBJECT STORAGE AND DATA LAKES

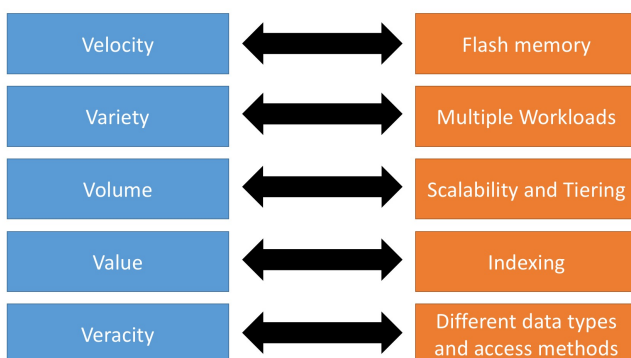
Many organizations are now taking great interest in new comprehensive data management strategies which involve building their own data lakes. These large-scale storage repositories are designed to save all kinds of data and to be accessed by different compute resources, aiming towards the elimination of data silos and enabling the creation of a wider range of applications capable of better insights.

Object storage is the best option for building the foundation of a data lake.

Object storage is the best option for building the foundation of a data lake. For example, data could come from the most disparate sources and can be easily tagged and indexed during the ingestion process. Once

data is safely stored as objects, it's available for metadata searching to all application stacks. Through specific API and commands, S3 based storage can also be leveraged as a primary store or as a secondary, but persistent, storage system. And if the object storage has tiering capabilities, it is possible to implement tiering (and retention) policies to move data across flash, disk and cloud or tape to improve the overall \$/GB of the entire storage infrastructure.

5 Vs of Big Data and Red Hat Ceph Storage



At the same time Big Data Storage infrastructures must be scalable and accessible through various protocols. This is why S3 is the perfect match for HDFS. They are complementary to each other:

- S3-based object stores are typically less costly (especially when based on commodity servers), more scalable and durable;
- While HDFS can be faster, S3 could be more suitable for specific use cases, (i.e. IoT, log analysis or in any other case where a large number of small files is involved);
- HDFS can be implemented as front-end of an object store, allowing for more simple data movement and access to data through different protocols.

HDFS is typically considered as part of the first storage tier while S3 is preferred for near-line storage and specific use cases, (as in all those cases where a large number of small files is involved).

BOTTOM LINE

Red Hat Ceph Storage, thanks to its flexibility and scalability, could be the best suited for building data lakes of any size. It's modern architecture, the improved performance of the latest version (version 2.0 released in June 2016) and its security features allow you to start off small and grow when needed. It's new policy engine associated to S3 protocol support will enable the construction of a multi-tier infrastructure that can span from low latency flash memory to the cloud to cover a very wide spectrum of big data applications, workloads and data types. Its flexibility allows enterprises to build storage infrastructures suitable for primary or secondary workloads depending on other Big Data technology already in use or planned.

By adopting Red Hat Ceph Storage, enterprises can effectively face the challenges posed by data growth and diversity with a sustainable and cost effective infrastructure. It can be the ideal solution for building the core of sustainable Big Data infrastructures for today's and future needs.

By adopting Red Hat Ceph Storage, enterprises can effectively face the challenges posed by data growth and diversity with a sustainable cost effective infrastructure.

The multi-protocol front-end, the object storage architecture back-end, it's modern distributed design, hardware independence and open, community based development model, all contribute to making Red Hat Ceph Storage a unique solution in the market, covering a growing number of use cases.

In a modern, two tier storage strategy which sees latency-sensitive workloads served by All-Flash Arrays and capacity-driven applications covered by high throughput scale-out storage, Red Hat Ceph Storage can play a perfect role in the latter category and, thanks to the latest improvements with Red Hat Ceph Storage 2.0, it can most definitely aspire to something more in the future.

JUKU

WHY JUKU

Jukus are Japanese specialized cram schools and our philosophy is the same. Not to replace the traditional information channels, but to help decision makers in their IT environments, to inform and to discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in the IT environment need to be aware of their surroundings: things are changing rapidly and there is a need to be constantly updated, to learn to adapt quickly and to support important decisions - but how? Through our support, our ideas, the result of our daily global interaction on the web and social networking with vendors, analysts, bloggers, journalists and consultants. But our work doesn't stop there - the comparison and the search is global, but the sharing and application of our ideas must be local and that is where our daily experience, with companies rooted in local areas, becomes essential in providing an honest and productive vision. That's why we have chosen: "think global, act local" as a payoff for Juku.

AUTHOR



Enrico Signoretti is an analyst, trusted advisor and passionate blogger (not necessarily in that order). He has been immersed in IT environments for over 20 years. His career began with Assembler in the second half of the 80's before moving on to UNIX platforms until now when he joined the "Cloudland". During these years his job has changed from highly technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm deeply focused on supporting end users, vendors and third parties in the development of their IT infrastructure strategies.

He keeps a vigil eye on how the market evolves and is constantly on the lookout for new ideas and innovative solutions. You can find Enrico's social profiles here: <http://about.me/esignoretti>

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources Juku Consulting srl (Juku) considers to be reliable but is not warranted by Juku. This publication may contain opinions of Juku, which are subject to change from time to time. This publication is covered by [Creative Commons License \(CC BY 4.0\)](#): Licensees may cite, copy, distribute, display and perform the work and make derivative works based on this paper only if Enrico Signoretti and Juku consulting are credited. The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Juku consulting srl has a consulting relationship with Red Hat. This paper was commissioned by Red Hat. No employees at the firm hold any equity positions with Red Hat. Should you have any questions, please contact Juku consulting srl (info@juku.it - <http://juku.it>).