# High Performance Object Storage for HPC, AI and Analytics

Madhu Thorat
IBM
India
madhu.punjabi@in.ibm.com

## 1. INTRODUCTION

Today a new world of modern application workloads has emerged as adoption of High Performance Computing (HPC), Artificial intelligence (AI) and Analytics solutions has accelerated. Applications in these domains are generating tremendous amounts of unstructured data. In fact, a study shows that 80% of data generated worldwide will be unstructured by 2025. This has led to the urgency of solving two big problems:

1) It has resulted in requirement of storage systems which can store and manage unstructured data.
2) In addition, the storage systems should be able to provide access to huge amounts of data at high speed and be scalable.

These problems can be addressed by High Performance Object Storage systems, which can store large volumes of unstructured data and meet the demands for high speed, low latency, and scale from storage. Moreover, object storage systems provide low-cost storage for large capacity of data and are suitable for HPC, AI, Analytics and Cloud native applications. Also, in the past few years' adoption and management of object storage has been simplified with Amazon S3 API which has become the standard protocol for leveraging object storage.

This poster will give information about High Performance Object Storage and its typical features which can help with the growing storage demands for HPC, AI and Analytics applications.

## 2. High Performance Object Storage Features for HPC, AI and Analytics

The poster will cover various features of a High Performance Object Storage solution, as shown in below figure.

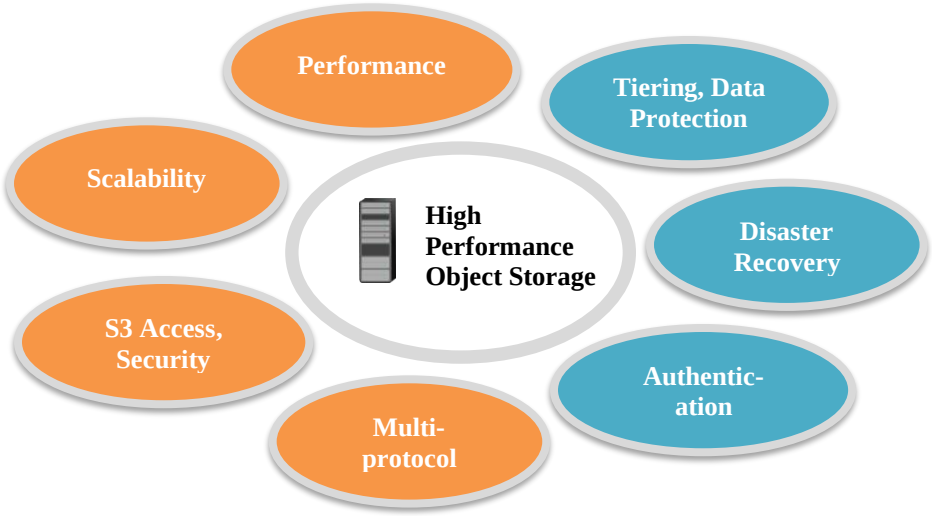


## 3. Use Case: Data Scientists accessing AI application data at high speed with High Performance Object Storage

The poster will cover a use case for a data scientist accessing High Performance Object Storage to process the data stored by AI applications using multiple protocols. The poster will also give an overview of the remarkable performance results delivered by a High Performance Object Storage solution offering more than 60 GB/s throughput for read workloads and more than 20 GB/s throughput for write workloads for HPC, AI and Analytics environments.

## REFERENCES

[1] https://www.ibm.com/docs/en/scalecontainernative?topic=516-spectrum-scale-data-access-services

[2] https://solutionsreview.com/data-management/80-percent-of-your-data-will-be-unstructured-in-five-years/