Western Digital.

Solution Test Report April 2022

Western Digital® Ultrastar® Data60 and MinIO Object Storage Solution Test Report

Abstract

This document provides a record of the testing performed to validate the compatibility and performance of MinIO solutions with the Western Digital Ultrastar Data60.

Notices

Western Digital Technologies, Inc., or its affiliates' (collectively "Western Digital") general policy does not recommend the use of its products in life support applications wherein a failure or malfunction of the product may directly threaten life or injury. Per Western Digital Terms and Conditions of Sale, the user of Western Digital products in life support applications assumes all risk of such use and indemnifies Western Digital against all damages.

This document is for information use only and is subject to change without prior notice. Western Digital assumes no responsibility for any errors that may appear in this document, nor for incidental or consequential damages resulting from the furnishing, performance, or use of this material.

Absent a written agreement signed by Western Digital or its authorized representative to the contrary, Western Digital explicitly disclaims any express and implied warranties and indemnities of any kind that may, or could, be associated with this document and related material, and any user of this document or related material agrees to such disclaimer as a precondition to receipt and usage hereof.

Each user of this document or any product referred to herein expressly waives all guaranties and warranties of any kind associated with this document any related materials or such product, whether expressed or implied, including without limitation, any implied warranty of merchantability or fitness for a particular purpose or non-infringement. Each user of this document or any product referred to herein also expressly agrees Western Digital shall not be liable for any incidental, punitive, indirect, special, or consequential damages, including without limitation physical injury or death, property damage, lost data, loss of profits or costs of procurement of substitute goods, technology, or services, arising out of or related to this document, any related materials or any product referred to herein, regardless of whether such damages are based on tort, warranty, contract, or any other legal theory, even if advised of the possibility of such damages.

This document and its contents, including diagrams, schematics, methodology, work product, and intellectual property rights described in, associated with, or implied by this document, are the sole and exclusive property of Western Digital. No intellectual property license, express or implied, is granted by Western Digital associated with the document recipient's receipt, access and/or use of this document or the products referred to herein; Western Digital retains all rights hereto.

Western Digital, the Western Digital logo, and OpenFlex are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Amazon and Amazon S3 are trademarks of Amazon.com, Inc. or its affiliates. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Lenovo® is a trademark of Lenovo in the United States, other countries, or both. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Mellanox and ConnectX are registered trademarks or trademarks of Nvidia Corporation in the U.S. and/or other countries. The NVMe and NVMe-oF word marks are trademarks of NVM Express, Inc. All other marks are the property of their respective owners.

Western Digital
5601 Great Oaks Parkway
San Jose, CA 95119

Table of Contents

Introduction

This Solution Test Report summarizes the results obtained during the validation of MinIO object storage application with a Western Digital® Ultrastar™ Data60. The objective of this exercise was to benchmark MinIO S3 performance on the Ultrastar Data60 using the WARP benchmark tool. This document includes guidelines for configuring MinIO Object Storage with an Ultrastar Data60.

MinIO is high-performance, S3 compatible object storage. Kubernetes-native, MinIO object storage software is available on every public cloud, every Kubernetes distribution, the private cloud, and the edge. MinIO is software-defined and is 100% open source under GNU AGPL v3. Customers depend on MinIO for AI/ML, IoT and big data workloads.

All tests were completed utilizing the MinIO Object Storage over mini-SAS HD interface to the Ultrastar Data60. These tests were performed by MinIO in Palo Alto, CA, USA.

This report is not an endorsement of MinIO by Western Digital, and no warranty of the product is either expressed or implied.

Objectives

There are several objectives for this test:

1. Validate environment configuration and compatibility between Minio object storage software and Ultrastar Data60.
2. Generate an idealized hardware configuration with OpenFlex Data24 and Minio to optimize for S3 performance.
3. Benchmark S3 read & write performance of Minio object storage when utilizing the OpenFlex Data24 as primary storage.

Prerequisite Table

Item	Notes
Servers with mini-SAS HD HBA	Western Digital External Storage Enclosures require mini-SAS HD connectivity, minimum of 2 ports for resiliency
25+ Gbps Network Switches	Minimum network connectivity to achieve ideal performance for designed solution

Technology Summary

Technology	Version
Lenovo® SR650 servers	
Linux® OS	CentOS 8.3.2011
Mellanox® ConnectX®-5 SmartNIC	Dual port 100 GbE, PCI Gen 3.0 x16, MOFED 5.3.1
Mellanox SN2700 Switches	32 port 100 Gbe Switch
Ultrastar Data60	FW: 3010-007
MinIO Object Storage	RELEASE.2021-06-17T00-10-46Z

Solution Overview

MinIO is a pioneer in the development of high-performance, cloud-native object storage, refining and perfecting many of the features, protocols and APIs that have come to define best in class. As the fastest object storage available, MinIO is often used for AI such as machine learning, IoT, data lake and other big data workloads. This solution combination of Ultrastar Data60 and MinIO object storage demonstrate some of the key benefits of distributed disaggregated infrastructure: flexibility, scalability, maintainability, and performance.

MinIO distributed object storage server Implements the Amazon S3™ API. MinIO Console, a browser-based GUI, incorporates all the functionality of MinIO Client in a design that feels familiar for IT admins and developers alike. Built to support cloud-scale deployments with minimal operational overhead, MinIO Console enables administrators and users to provision multi-tenant object storage as a service, visually inspect the health of the system, perform key audit tasks, and simplify integration (via webhooks and API) with other components.

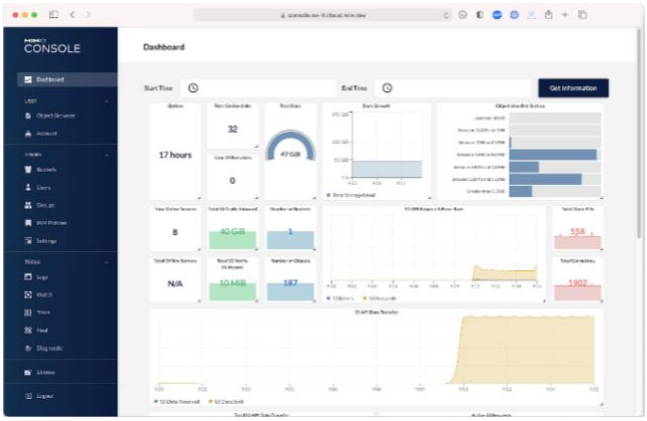

Diagram 1: MinIO Console Dashboard

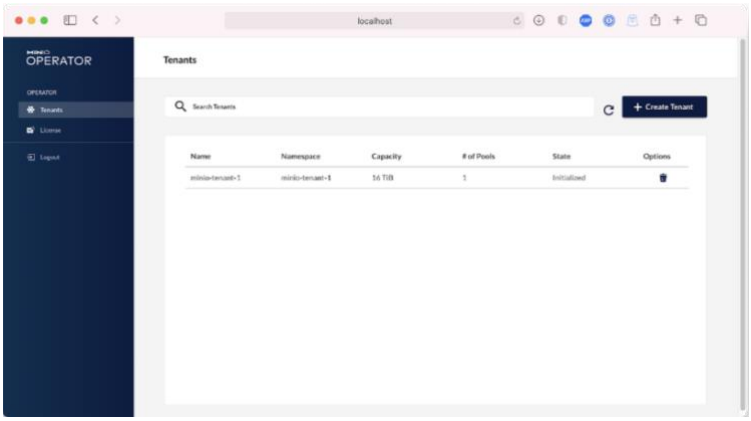

Diagram 2: MinIO Kubernetes Operator Dashboard

MinIO's singular focus on object storage has resulted in persistence layer that is performant, resilient, secure, and scalable across various hybrid-cloud deployments.

Key Features	Description
Erasure Coding	MinIO protects data with per-object inline erasure coding, written in assembly code delivering the highest performance. MinIO uses Reed-Solomon code to stripe objects into data and parity blocks, configurable to desired redundancy level.
BitRot Protection	MinIO's SIMD accelerated implementation of the HighwayHash algorithm ensures that it will never return corrupted data by capturing and healing corrupted objects on the fly. Integrity is ensured from end to end by computing hash on WRITE and verifying it on every READ from the application, across the network and to the memory/drive.
Identity and Access Management	MinIO supports the most advanced standards in identity management, integrating with the OpenID connect and LDAP compatible IDP providers. That means that access is centralized, passwords are temporary, and tokens rotated. Furthermore, access policies are fine grained and highly configurable at the API level, which means that supporting multi-tenant and multi-instance deployments is simple.
Encryption	MinIO supports multiple server-side encryption schemes to protect data. MinIO's approach ensures confidentiality, integrity, and authenticity with negligible performance overhead. Server side and client-side encryption are supported using AES-256-GCM, ChaCha20-Poly1305 and AES-CBC. Encrypted objects are tamper-proofed with AEAD server-side encryption.
Bucket and Object Immutability	MinIO supports object locking, retention, legal holds, governance and compliance for objects and buckets. Object locking is frequently combined with versioning to eliminate the risk of data tampering or destruction. Retention rules ensure that an object is WORM protected for a configurable period. This capability is critical for ransomware use cases and can be used in conjunction with leading backup vendors to ensure fast backup/restore across multiple workloads.
Data Lifecycle Management and Tiering	MinIO lifecycle management tools allow administrators to define how long data remains on disk before being removed. MinIO protects data within and across clouds with a wide range of policies built on object and tag filters to declare expiry rules. Bucket expiration rules are fully compliant with MinIO WORM locking and legal holds. MinIO can programmatically tier objects across storage mediums and cloud types to optimize for performance and cost.
Bucket and Object Versioning	MinIO's object-level versioning provides data protection and serves as the foundation for data lifecycle management, tiering and locking. MinIO follows Amazon's S3 structure/implementation to independently versioned objects. This allows users to retain multiple variants of every object at the bucket level, eliminating the need for a separate snapshot process.
Scalability	MinIO scales out or horizontally through server pools. Each server pool is an independent group of nodes with their own compute, network, and storage resources. In multi-tenant configurations, each tenant is a cluster of server pools in a single namespace, fully isolated from the other tenants' server pools. Capacity can easily be added to an existing system by configuring a new server pool and MinIO automatically prepares it for and places it in service.

Active-Active Replication	MinIO's active-active replication enables organizations to use object storage across multiple data centers and clouds in a manner that is resilient and scalable that can withstand a data center failure with no down time. Read-write operations can be conducted on either of the clusters and data is replicated bidirectionally between them. Replication of objects and their metadata is on a bucket level, using near-synchronous replication to update objects immediately after any change.
Metadata Architecture	MinIO has no separate metadata store. All operations are performed atomically at object level with strong consistency. This approach isolates any failures to be contained within an object and prevents spillover to larger system failures.
Cloud Native and Kubernetes Native	MinIO was designed and built as a multi-tenant and multi-user system that scales seamlessly from TBs to any size. Kubernetes plays a key role in MinIO's hybrid cloud functionality. As favored by DevOps teams, Kubernetes-native design requires an operator service to provision and manage a multi-tenant object-storage-as-a-service infrastructure. Each of these tenants run in their own isolated namespace while sharing the underlying hardware resources.

Using 12 Gb mini-SAS HD cables between the MinIO servers and the Ultrastar Data60 storage provides scalable performance by supporting both wide mode and redundant path SAS up to 12 connections per host. Further customizable zoning configurations allow up to 6 servers to connect with a single Ultrastar Data60 storage platform, all without sacrificing performance. These performance and scaling capabilities are also combined with our innovative ArcticFlow™ and patented IsoVibe™ technologies to reduce drive returns (up to 62%) compared to the previous generation while reducing both power and cooling requirements.

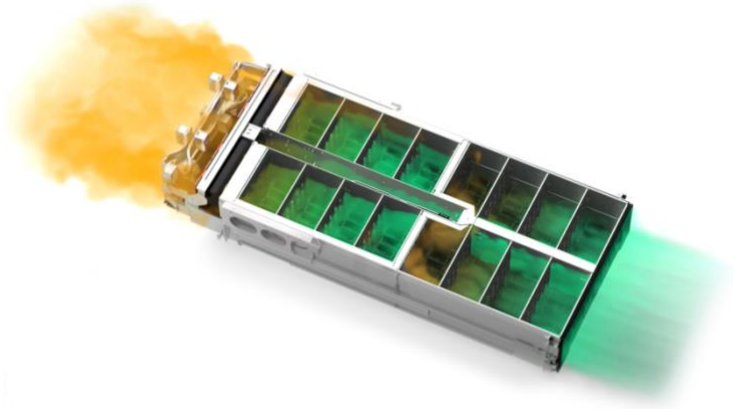

Diagram 4: Ultrastar Data102 Thermal Zones of ArcticFlow

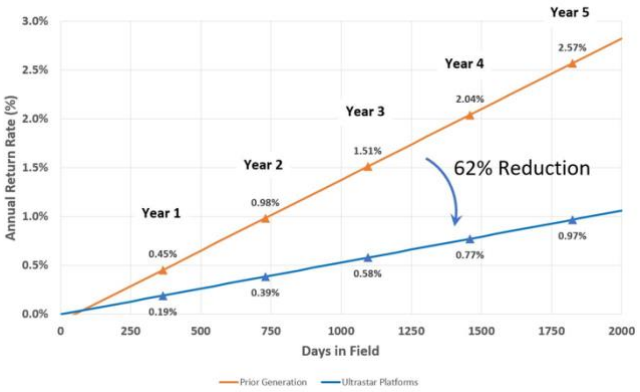

Diagram 3: Ultrastar JBOD and storage server annual return rate

ArcticFlow in the Ultrastar Data60 and Data102 is a thermal zone cooling technology that divides our enclosures into two separate thermal zones, each with dedicated paths for fresh cold air to reach the devices. The front zone is cooled via conventional airflow from the cold aisle and the higher temp exhaust is routed to the sides of the rear zone, preventing the typical heat buildup of other dense storage enclosures. The rear zone is then cooled by a dedicated cold aisle air duct, which is spread across all rear devices. This reduces what is typically the hottest devices temps which contribute to longer device lifespan.

IsoVibe is a collection of several patented innovations to reduce vibration transmission from both drives and fans that have caused longer seek times and possibly even errors in reading and writing. Isolating and dampening the vibration of all moving devices within the enclosure ensures that the Ultrastar Data 60 and Data102 can maintain expected performance for the full lifespan of the enclosure.

For this report, Lenovo SR650 servers were used for the MinIO storage servers, but MinIO object storage software only has minimal hardware requirements, enabling IT organizations to pair Ultrastar Data60 with their server platform of choice.

The MinIO Object Storage delivers the following solution(s):

☒ Backup and Archive ☒ Software Defined Storage ☐ NAS/SAN Services ☐ Storage Platform Manager

☐ Media Asset Management ☐ High Performance Computing ☒Other: Object Storage

The MinIO Object Storage is applicable to the following vertical(s):

☒ Media & Entertainment ☒ Life Sciences/HPC ☒ Financials ☒ Enterprise ☒ Government

☒ Other: Telecommunications

Architecture

Our testing environment was comprised of six MinIO object storage servers connected to two Ultrastar Data60s via two redundant mini-SAS HD cables per server. The release version of MinIO September 2021 was used for this report and was configured using standard setup procedures. Additional per server configurations steps are needed to enable Device Mapper (DM) multipath as a load balancer for the redundant device connections. To closely replicate a real customer deployment, we paired the storage cluster (MinIO Servers + Ultrastar Data60s) configuration with six load generation clients running the WARP S3 benchmarking tool to generate S3 traffic and measure the performance of the overall solution.

The illustration below shows an ideal solution configuration, where all networks are physically segmented for maximum availability and performance.

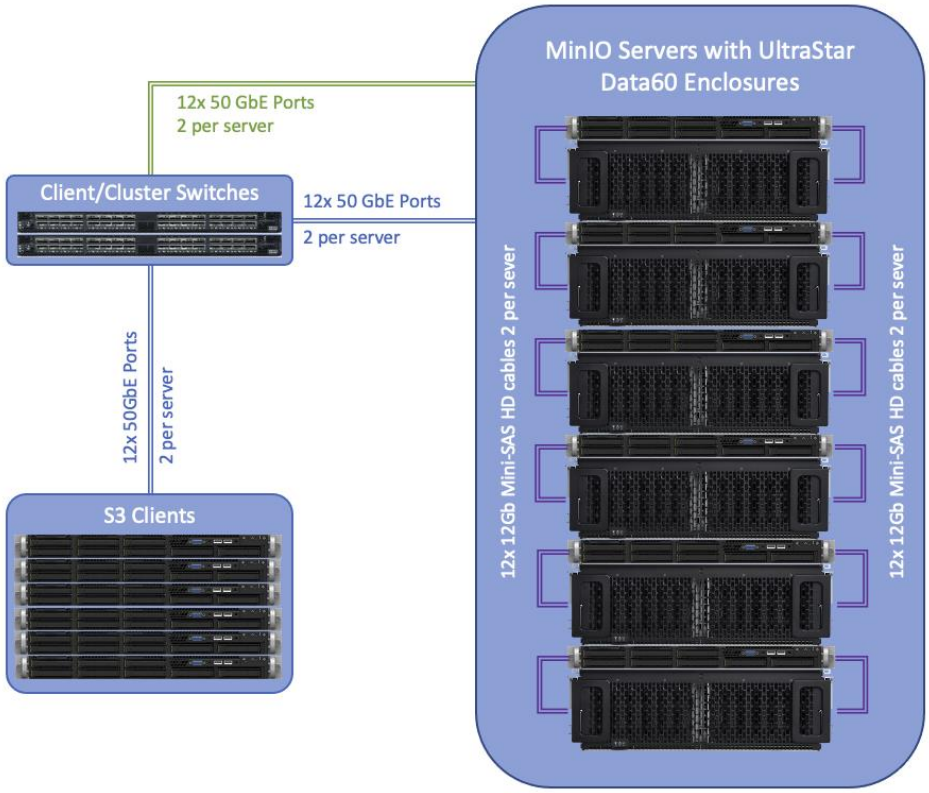

Diagram 5: Idealized minimal deployment configuration

Network Topology

A standard MinIO object storage configuration was set up with both a client and a cluster network. The client network is for servicing IO requests of clients and applications. The cluster network is for replicating or erasure encoding data and distributing it to storage nodes. The key to achieving ideal performance is to ensure that all interconnections (network and SAS) are as close to symmetrical bandwidth as possible. To this point we recommend using redundant 50 GbE networking or more across the entire solution.

Testing Environment

Each Ultrastar Data60 was configured into the "three zone" configurations so each server would be attached to 20 HDDs. In this configuration, 4 mini-SAS HD ports are available for each zone, two per IO module. Each server was then connected using two mini-SAS HD cables to a single IO module of Ultrastar Data60. Redundant connection to the second IO module was not done to bypass any performance loss due to DM multipathing since we are limited to only 20 HDDs per server.

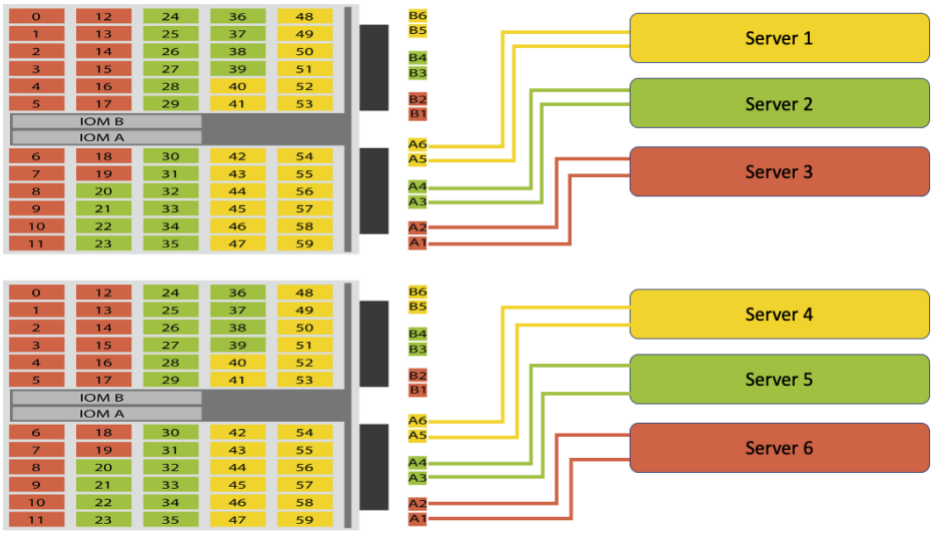

Diagram 6: Ultrastar Data60 three zone config and server cabling

The tested deployment is represented by the illustration below.

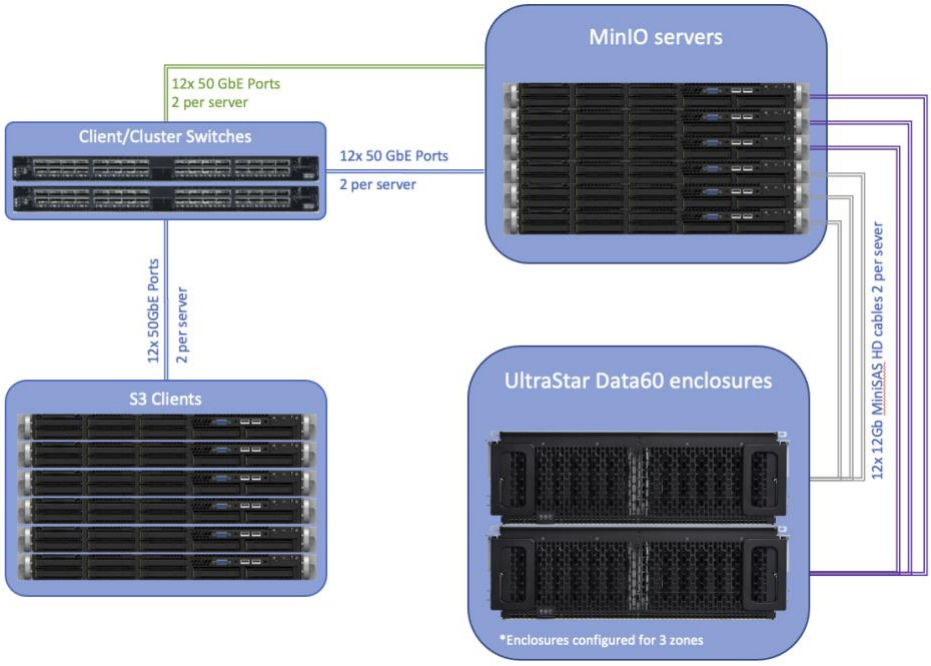

Diagram 7: Full connectivity of tested configuration

Ultrastar Data60 and MinIO Configuration

1. Configure Ultrastar Data60 into the desired multi-zone setup.
2. Cable each server to the appropriate ports for the specific zone.
3. Verify Ultrastar Data60 is connected to the servers running MinIO.
4. Verify all drives are available to the to host OS and MinIO.
5. Create and mount filesystems on all drives for each server running MinIO.
6. Configure MinIO to utilize all mounted filesystems for the multipath drives.
7. Configure MinIO for erasure coding EC:4.

Workflow Testing

A 6 node MinIO cluster was created utilizing storage from two Ultrastar Data60s. Each server was provisioned with CentOS 8.3.2001, a Mellanox ConnectX-6 100 GbE Network card, a Broadcom® LSI 9300-8E SAS HBA, Intel® Xeon® Gold 6150 CPU, and 384 GiB RAM.

We ran the WARP S3 Benchmark (https://github.com/minio/warp) for our performance tests. This tool conducts benchmark tests from one or more clients to one or more hosts. WARP is an open source S3 performance benchmark tool developed and maintained by MinIO.

We ran the WARP benchmark with 6 clients. Each client was provisioned with CentOS 8.3.2011, a Mellanox ConnectX-5 100 GbE Network card, Intel Xeon Gold 6150 CPU, and 384 GB RAM.

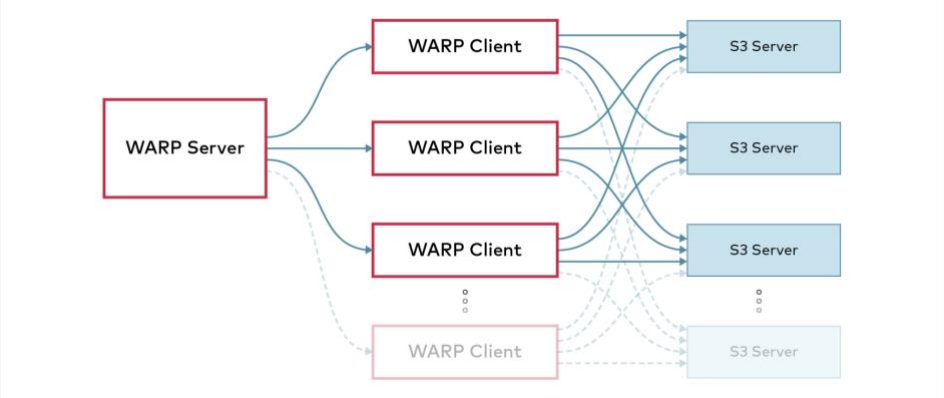

Diagram 8: WARP multiclient connection distribution

The benchmark ran with objects of various sizes from 1MiB to 1 GiB.

Results Summary

As a result of WARP benchmarking tests with the default erasure coding parity of 4 on 6 MinIO servers backed by the Ultrastar Data60 device, we obtained peak write/PUT bandwidth as 12067 MiB/sec. The read performance results were disregarded due to the system configuration enabling the OS to cache read data in the DRAM. Retesting is being planned and scheduled to provide a more comprehensive performance characteristics with newer versions of MinIO, Ultrastar Data60, and all other hardware.

The chart pictured below show WARP S3 write/PUT bandwidth for various object sizes on a MinIO cluster with 6 nodes.

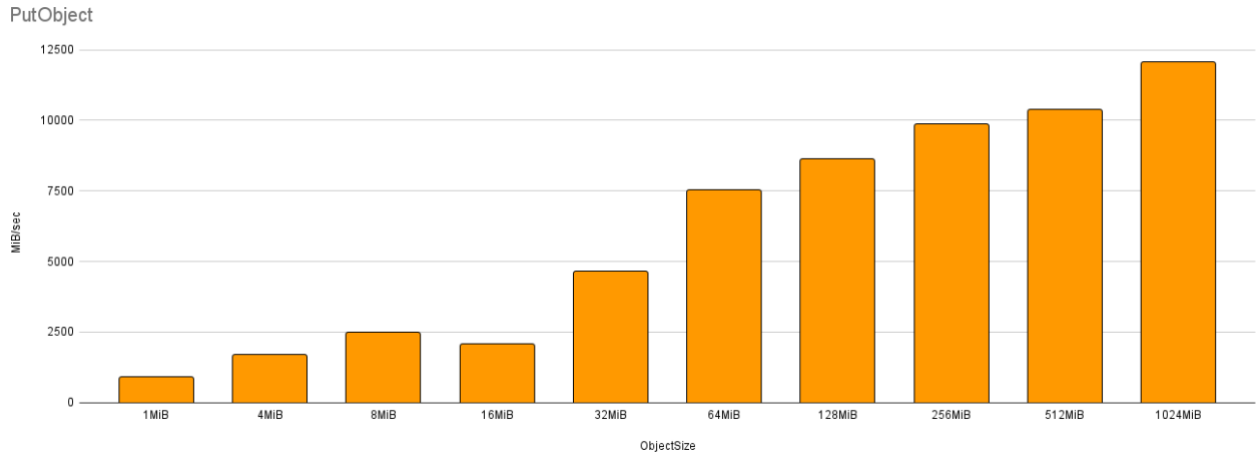

Diagram 9: Write/PUT throughput results accumulated across 6 MinIO servers

Testing was conducted by MinIO.

Appendix

Contributors

Name	Company	Title
David Tobin	Western Digital	Sr. Technical Marketing Engineer
Matt Sarrel	Minio	Director, Technical Marketing

References

Document Title	Date Delivered
Ultrastar Data60 User Manual	September 2, 2021
MinIO Administrator Guide	September 2, 2021
MinIO Installation Guide	September 2, 2021

Document Feedback

For feedback, questions, and suggestions for improvements to this document send an email to the Data Center Systems (DCS) Technical Marketing Engineering (TME) team distribution list at pdl-dcs-tm@wdc.com.

Version History

Version	Date	Notes
1.0	April 22, 2022	Initial release