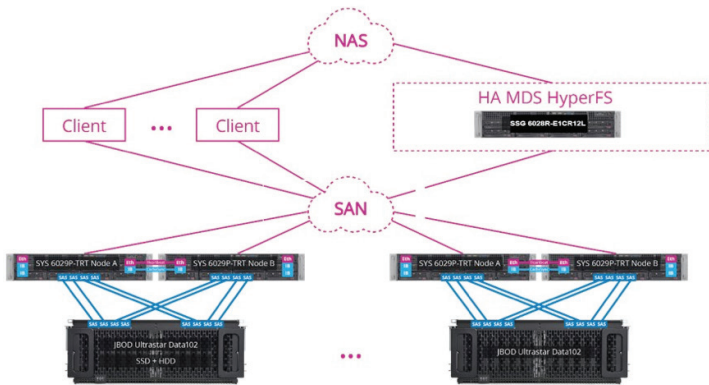


Scale-out Data Storage System



Storage servers: Supermicro SuperServer 6029P-TRT
Metadata servers: Supermicro SuperServer SSG-6028R-E1CR12L
JBOD: Ultrastar Data102 with HDDs and SSDs
SW: RAIDIX 4.6

To create a scale-out system based on HyperFS, two types of servers are required — one for data storage and one for metadata storage.

1. The data storage controller function will be performed by the SuperMicro® SuperServer® 6029P-TRT
2. To store metadata, several SSDs are installed in the Ultrastar® Data102 storage platform. They are configured into a RAID configuration and access to the MDS is provided via a SAN.

It is possible to connect up to four Ultrastar Data102 storage platforms to one storage server in cascade. The RAIDIX software controller function runs on top of the storage nodes.

Server Configuration

Item	Description	P/N	Qty
Platform	Supermicro SuperServer 6029P-TRT	SYS-6029P-TRT	2
CPU	Intel® Xeon® Silver 4112 Processor	Intel Xeon Silver 4112 processor	4
Memory	16GB PC4-21300 2666MHz DDR4 ECC Registered DIMM Micron MTA36ASF472PZ-2G6D1	MEM-DR416L-CL06-ER26	16
Persistent memory	Micron NVDIMM-N Module 16GB, MTA18ASF2G72PF1Z-2G6	MTA18ASF2G72PF1Z-2G6	8
System Disk	SanDisk® Extreme PRO 240GB SSD	SDSSDXPS-240G-G25	4
Hot-swap 3.5" to 2.5" SATA/SAS Drive Trays	Tool-less black hot-swap 3.5-to-2.5 converter HDD drive tray (Red tab)	MCP-220-00118-0B	4
HBA for cache-sync	Mellanox ConnectX-4 VPI adapter card, EDR IB (100Gb/s), dual-port QSFP28, PCIe3.0 x16	MCX456A-ECAT	2
HBA for JBOD connection	Broadcom HBA 9400-8e tri-mode storage adapter	05-50013-01	4
Ethernet patchcord	Ethernet patch cord for cache sync 0.5m		1
Cable for cache sync	Mellanox passive copper cable, VPI, EDR 1m	MCP1600-E001	2
HBA for host connection	Mellanox ConnectX-4 VPI adapter card, EDR IB (100Gb/s), dual-port QSFP28, PCIe3.0 x16	MCX456A-ECAT	2

Note: Alternatively, 2x AVAGO 9300-8e controllers or 1x 9405W-16e controller (requires availability of 16x PCIe slots) can be used for the JBOD connectivity.

Overall, 5.5 PB¹ of data is placed in one deep storage rack with one common namespace.

Target Applications:

- HPC (Life Sciences: DNA and RNA sequencing, climate research)
- Technical computing (CAS/CAM, Simulations)
- Video Surveillance, up to 10,000 Full HD cams
- Media & Entertainment (Content delivery and streaming, content archive, content management)

For purposes of server sizing, the PCIe bus is used to install backend and frontend controllers. Broadcom® 9400 8e controllers are used to connect the Ultrastar Data102 storage platform to the Supermicro servers. The synchronization channel slot is Infiniband, supported by Mellanox® ConnectX-4 VPI adapter card (PCIe3.0 x16) on each node.

The server is equipped with Micron® NVDIMM modules to protect the cache from power outage. Dirty cache segments will be synchronized via InfiniBand 100Gb.

The Ultrastar Data102 can be equipped with Ultrastar DC HC620 or HC6560 SAS HDDs, providing a data repository of up to 2.0PB in a 4U storage rack. Minimum configuration is 60 drives, providing an upgrade path of up to 102 drives. Up to 24 SAS/SATA SSDs can be installed in those drive slots to provide storage for the metadata.

Note: Each controller must have at least five x8 slots (without room for further scaling). Systems built for lower performance of 3-4 Gbps per node can be configured with only two slots.

Note: Ultrastar Data102 can also be switched out for Ultrastar Data60, if the depth of your rack is not sufficient to host the Ultrastar Data102 (Depth: 1047mm (41.25") or Depth in Rack: 1197mm (47.13") w/ CMA.

¹One MB is equal to one million bytes, one GB is equal to one billion bytes, one TB equals 1,000GB (one trillion bytes), and one PB equals 1,000TB when referring to hard drive capacity. Accessible capacity will vary from the stated capacity due to formatting and partitioning of the hard drive, the computer's operating system, and other factors.

Software

Item	Description	P/N	Qty
RAIDIX	RAIDIX 4.6 DC/NAS/iSCSI/FC/SAS/IB/SSD-cache/QoSmic/SanOpt/Extended 5 years support/unlimited disks	RX46DSMMC-NALL-SQ0S-P5	1

JBOD Configuration


Item	Description	P/N	Qty
JBOD	Ultrastar Data102 storage platform		1
SAS cable	Ultrastar Data102 cable IO HD mini-SAS to HD mini-SAS 2m 2 pack		8

MDC HyperFS

Item	Description	P/N	Qty
Platform	Supermicro SuperServer 6028R-E1CR12L	SSG-6028R-E1CR12L	1
CPU	Intel Xeon E5-2620v4 processor	Intel Xeon E5-2620v4	2
Memory	32GB DDR4 Crucial CT32G4RFD424A 32Gb DIMM ECC Reg PC4-19200 CL17 2400MHz	CT32G4RFD424A	4
System Disk	SanDisk Extreme PRO 240GB SSD	SDSSDXPS-240G-G25	2
Hot-swap 3.5" to 2.5" SATA/SAS Drive Trays	Tool-less black hot-swap 3.5-to-2.5 converter HDD drive tray (Red tab)	MCP-220-00118-0B	2
HBA	Mellanox ConnectX-4 VPI adapter card, EDR IB (100Gb/s), dual-port QSFP28, PCIe3.0 x16	MCX456A-ECAT	1

Note: The controller function for RAIDIX-based data storage systems can be played virtually by any x86 server platform, including Supermicro, AIC, Dell, Lenovo, HPE, and many others. This reference architecture chose to use the Supermicro platform as an example.


Note: This reference architecture is subject to maximum performance and functioning of all the options available.



IsoVibe

IsoVibe™ Vibration Isolation Technology

Precise cuts in the baseboard provide a suspension for the drives in the chassis, isolating them from transmitted vibration. The result is that consistent performance is maintained, even when all the drives are working hard.



ArcticFlow

ArcticFlow™ Thermal Zone Cooling Technology

By introducing cool air into the center of the chassis, drives operate at lower and more consistent temperatures than conventional systems. This results in lower fan speeds, reduced vibration, lower power consumption, quieter operation and ultimately higher reliability.

Western Digital.

5601 Great Oaks Parkway
 San Jose, CA 95119, USA
 US (Toll-Free): 800.801.4618
 International: 408.717.6000

www.westerndigital.com

© 2020 Western Digital Corporation or its affiliates. All rights reserved. Western Digital, the Western Digital logo, SanDisk, and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Dell is a trademark of Dell Inc. or its subsidiaries. Lenovo is a trademark of Lenovo in the United States, other countries, or both. RAIDIX and the RAIDIX logo are trademarks of RAIDIX. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Supermicro and SuperServer are trademarks of Super Micro Computer, Inc. or its subsidiaries in the United States and other countries. Broadcom is a trademark of Broadcom Inc. Micron is a trademark of Micron Technology, Inc. Mellanox is a trademark of Mellanox Technologies, Ltd. All other marks are the property of their respective owners.