
Media WorkFlow Solution – DataOn Platform

CloudByte Reference Architecture

Table of Contents

1	Executive Summary	3
2	Specifications	4
3	Solution Architecture.....	5
4	CloudByte Software Components	7
5	Hardware BOM.....	8
6	Benefits	9
7	About CloudByte	10

1 Executive Summary

Media editing workflows present a unique challenge to most storage architectures. Media projects are characterized by intense, overlapping workflows, often with teams of 30 or more editors at work concurrently. While very high network bandwidth is utilized for loading a track from disk and for storing an editing track to disk, actual disk I/O is relatively low.

In the past, this workload was addressed by overprovisioning the storage configuration, especially storage networking, resulting in very high costs. In contrast, this document provides a Media workflow solution that provides very high throughput in an efficient footprint:

- Workflows: 30 concurrent streams
- Capacity: 9 Terabytes of data
- Protocols: Storage accessible via NFS
- Reliability: High Availability pair

From a hardware perspective, this reference architecture is based on the DataOn Cluster in a Box product line. Off-the-shelf SSD and HDD drives are utilized for caching and capacity. Control and data paths are on separate Ethernet links for performance and reliability. With this hardware platform, Media producers can expect reliable editing performance, high bandwidth, and high availability.

In summary, this reference architecture provides the following benefits:

- Reliable performance
 - This architecture provides reliable performance for all editors by eliminating “noisy neighbor” issues.
- High throughput
 - By utilizing high-capacity 10 GbE links rather than relatively slow Storage Area Network links, throughput is maximized.
- Cost efficiencies
 - By utilizing a very dense architecture (dual processors, dual power supplies, and 12 drive bays in a 2U enclosure), cost efficiencies are maximized.

2 Specifications

2.1 Performance and Capacity

This reference architecture provides very high throughput for 30 concurrent media editing streams. The overall data capacity, at 9 Terabytes, provides sufficient storage for a large editing project. Should additional storage be required, an expansion storage chassis can be connected to existing SAS or PCI-e ports.

2.2 System Configuration

As media editing is performed on files, only NFS access is needed. Since the media editing cycle is a critical part of the production process, reliability is ensured via a high availability pair configuration.

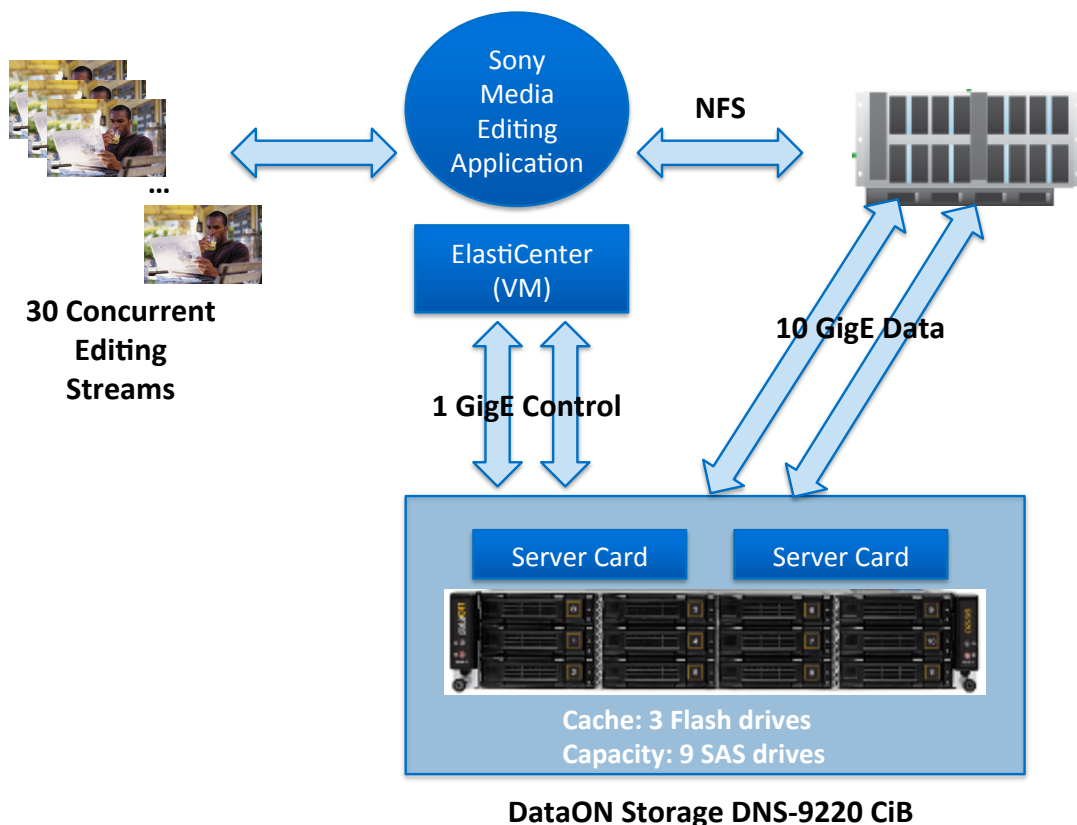
3 Solution Architecture

3.1 Overview

The solution utilizes:

- One CloudByte ElastiCenter™ management node running in a Virtual Machine (VM). For cost efficiency, the VM runs on the same server as the Sony media editing application.
- To maximize density, the hardware platform implemented is the DataOn Cluster-in-a-Box (DNS-9220 CiB) unit. The CiB takes up only 2U yet contains:
 - Dual Server cards (used to support two CloudByte ElastiStor™ nodes, configured in a High Availability Group).
 - 12 disk drive slots containing
 - 3 Flash drives for cache
 - 9 SAS drives for storage capacity
 - Dual 1 GbE management ports
 - Four 10 GbE data traffic ports

The reference architecture is depicted below:



1

3.2 Storage

Storage drives and configuration are:

- Three SSD drives are implemented for caching:
 - Read Cache: 1 x 800 GB SSD drive is used for the Read Cache
 - Write Cache: 2 x 200 GB SSD drives are used for the Write Cache. These drives are configured as mirrors to ensure that each server's memory cache is readily available to the alternate server should a failure occur.
- Nine x 1 Terabyte SAS drives are implemented for capacity. This size of disk was chosen to balance the storage performance needed without sacrificing density.

3.3 High Availability Group

CloudByte High Availability Groups can be configured in a variety of ways:

- Active / Passive Pairs: In this configuration, a node is on standby until the primary node fails. Then, the active node's load fails over to the passive node. Performance is typically not impacted by a single node failure.
- Active / Standby Group: CloudByte enables up to three nodes to be active with one standby node. Upon a failure, then the active node's load is taken by the passive node. In this configuration, performance is typically not impacted by a single failure.
- Active / Active Group: Up to four nodes can be configured to carry load. In the event of a failure, the load is distributed among the remaining active nodes. This configuration can offer protection against even three node failures. In this configuration, performance is typically not impacted by a single failure. Depending on the overall load, performance may not be impacted by up to three node failures.

For this reference architecture, the two controllers will be configured in an Active / Active group. Loss of one controller will not result in an outage and, depending on the editing workload cycle, may not even impact performance.

3.4 Control and Data Traffic

In this configuration, CloudByte control and data traffic flows over separate paths:

- Control traffic flows over 1 Gigabit Ethernet links. As the DataOn unit includes 2 x 1 GbE management ports, no external networking cards are needed.
- Data traffic flows over 10 Gigabit Ethernet links. As the DataOn unit includes 4 x 10 GbE RJ45 data ports, no external networking cards are needed.

This feature enables both fast performance and high throughput for data traffic.

4 CloudByte Software Components

This solution utilizes the capabilities of the CloudByte product family.

4.1 CloudByte ElastiStor

CloudByte ElastiStor™ is a full-featured storage software product. ElastiStor provides dynamically selectable performance to each application or tenant by continuously monitoring and adjusting key storage performance characteristics including IOPS, throughput, and latency. Based on ElastiStor's embedded analytics capabilities, an application's performance needs also can be dynamically adjusted by the administrator for the ultimate in flexibility. ElastiStor leverages the Zettabyte File System (ZFS) to create a unified storage pool for file, block, and JBOD storage. This capability not only eliminates storage silos but also enables scale-out growth up to cloud volumes. Further, ElastiStor runs on industry-standard hardware, allowing direct access to the most economical storage components from the vendor of choice. With ElastiStor, Enterprises and Service Providers can finally get ahead of data growth while providing exactly the performance needed for each application.

4.2 CloudByte ElastiCenter

CloudByte ElastiCenter™ is a web-based centralized management console that controls the storage environment worldwide and scales from a single site to many geographically dispersed clusters. ElastiCenter enables the definition of not only minimum selectable performance levels but also allocation of excess cycles to particular applications or tenants. Further, ElastiCenter provides efficient setup, customization, and tuning – all from a single interface. Lastly, ElastiCenter includes the capability to delegate administrative authority so that the organization can move authority for changes and updates close to the groups or customers that need support.

4.3 CloudByte Add-On Options

ElastiHA

ElastiHA provides high availability groups of up to 4 nodes. While conventional high availability only protects against the failure of 1 node, ElastiHA protects against the failure of up to 3 nodes. This capability is a perfect match for ElastiStor RAIDZ-3 data protection capabilities which protect against the failure of up to 3 physical drives without losing data. ElastiHA can be controlled from both ElastiCenter and our REST API.

5 Hardware BOM

The solution Bill of Materials provides details on the exact hardware configuration to achieve the specified performance and capacity requirements.

Tier / Component	Description	Quantity
ElastiCenter	Not applicable – installed in a Virtual Machine on the application server	0
ElastiStor	DataOn Cluster-in-a-Box (DNS-9220 CiB)	1
	Read Cache – 800GB SanDisk LB806M Lightning SSD	1
	Write Cache – 200GB SanDisk LB206M Lightning SDD	2
	Capacity – Western Digital HDD drive 1 Terabyte 7200 RPM WD1001FYYG	9
Networking	Utilizing the ports built into the default DataOn configuration: <ul style="list-style-type: none"> • 2 x 1 GbE ports are allocated to ElastiCenter / control traffic • 4 x 10 GbE ports are allocated to ElastiStor / data traffic 	0
Cabinet	Not applicable – utilizes 2U out of an existing rack	0

6 Benefits

Reliable Performance

This architecture provides reliable performance for all editors by eliminating “noisy neighbor” issues.

High throughput

By utilizing high-capacity 10 GbE links rather than relatively slow Storage Area Network links, throughput is maximized.

Cost efficiencies

By utilizing a very dense architecture (dual processors, dual power supplies, and 12 drive bays in a 2U enclosure), cost efficiencies are maximized.

7 About CloudByte

Targeted to the needs of service providers – public and private, CloudByte’s patent-pending storage software products uniquely deliver dynamically selectable performance for each and every application or tenant. By providing unified block, file, and direct attached storage, CloudByte eliminates inefficient and costly storage silos. CloudByte provides scale-out growth from a single site to distributed cloud environments without disruption. Established in 2010 and managed by technology executives from companies such as Cisco, HP, IBM, NetApp, Novell, and SanDisk, CloudByte is headquartered in the Silicon Valley and has a development center in India. CloudByte is venture-backed by Fidelity Worldwide Investment, Nexus Venture Partners and Kae Capital. For more information, visit www.cloudbyte.com or follow @cloudbyteinc.

CloudByte Inc.

20863 Stevens Creek Blvd
Suite 530
Cupertino, CA 95014
USA
Phone: +1-855-380-BYTE (2983)

CloudByte Technologies India Pvt. Ltd.

Plot No. 2799 & 2800, Srinidhi Bldg
3rd Floor, 27th Main, Sector – 1
HSR Layout, Bangalore 560102
India
Call: (91)-80-2258-2804