

---

**Tiered IOPS Storage for Service Providers  
– Dell Platform and Fibre Channel protocol**

**CloudByte Reference Architecture**

---

# Table of Contents

---

<b>1</b>	<b>Executive Summary .....</b>	<b>3</b>
<b>2</b>	<b>Performance Specifications by Tier .....</b>	<b>4</b>
<b>3</b>	<b>Solution Architecture.....</b>	<b>5</b>
<b>4</b>	<b>CloudByte Software Components .....</b>	<b>7</b>
<b>5</b>	<b>Hardware BOM.....</b>	<b>9</b>
<b>6</b>	<b>Benefits .....</b>	<b>10</b>
<b>7</b>	<b>About CloudByte .....</b>	<b>11</b>

## 1 Executive Summary

Service Providers are increasingly being asked to host enterprise applications in their Clouds. Yet, enterprise applications typically require stringent performance at every level – from server down to storage. Since predictable storage performance is a key feature of CloudByte’s architecture, Service Providers have been asking for a reference architecture that includes performance tiers. This document provides a tiered storage solution:

- Capacity: 100 Terabytes of data
- Total IOPS: 56,750
- Tiers:
  - Gold: All SSD for caching and storage
  - Silver: Hybrid combination of SSD for caching / SAS for storage
  - Bronze: All SAS for caching and storage
- Protocol(s): Storage is accessible via the Fibre Channel protocol
- Virtualization: Supports both VMware and Citrix server virtualization environments
- Open Standards: Supports both OpenStack and Cloud Stack

From a hardware perspective, this reference architecture is based on a Dell hardware implementation. Off-the-shelf Seagate and STEC drives are utilized for caching and capacity. Control and data paths are on separate Ethernet links for performance and reliability. With this hardware platform, Service Providers can expect high reliability, availability, and performance.

In summary, this reference architecture provides the following benefits:

- Foundation for New Services
  - This architecture enables a wide variety of tiered services to be created and defined, thereby maximizing service provider flexibility.
- Performance Needed for Enterprise Applications
  - This reference architecture provides the performance and flexibility needed to handle enterprise-class applications – whether they require minimum transaction times or maximum throughput.
- Strong fit for even leading-edge customer environments
  - By supporting multiple virtualization, multiple standards for ‘stacks’, and multiple storage access protocols, this reference architecture maximizes the number of customers that can utilize the Service Provider’s services.

## 2 Performance Specifications by Tier

This reference architecture provides a unified storage pool of 100 TB with three performance tiers. This storage can be accessed by the Fibre Channel protocol.

There are three primary tiers with different performance objectives:

- **Gold:** High IOPS levels for compute-intensive operations. Example use cases:
  - High Performance Database Applications
  - Analytic Applications Acceleration
  - Highly Transactional Applications
  - Virtual Desktop Infrastructure (VDI)
- **Silver:** Medium IOPS levels for Read / Write intensive operations. Example use cases:
  - Virtualization
  - Shared filestores
  - Streaming
- **Bronze:** Low IOPS for Write intensive operations. Example use cases:
  - Backup
  - Disaster Recovery
  - Versioning for Virtual Machines and Development releases

Detailed specifications by tier are summarized in the table below.

Tier	Size / Data Capacity	IOPS Needed	% Read / Write IOPS
<b>Gold</b>	15% / 15 TB	3000 IOPS per TB / 45,000 IOPS	30% Read / 70% Write
<b>Silver</b>	50% / 50 TB	200 IOPS per TB / 10,000 IOPS	60% Read / 40% Write
<b>Bronze</b>	35% / 35 TB	50 IOPS per TB / 1,750 IOPS	30% Read / 70% Write
<b>Totals</b>	100 TB	56,750 IOPS	

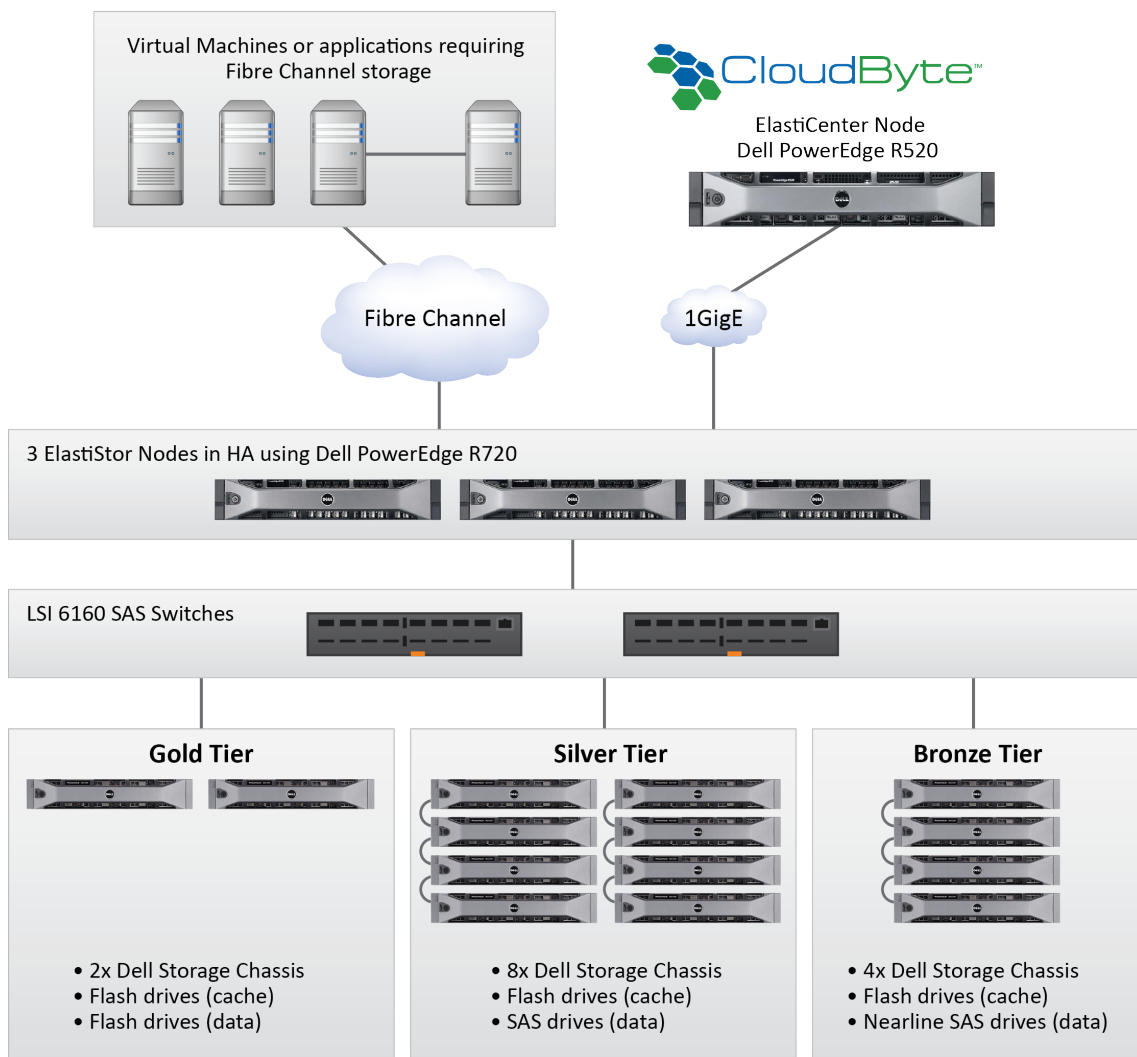
### 3 Solution Architecture

#### 3.1 Overview

The solution utilizes:

- Two CloudByte ElastiCenter™ management nodes (for redundancy).
- Four CloudByte ElastiStor™ nodes, configured in a High Availability Group (see following section).
- Three tiers of storage – Gold, Silver, and Bronze:
  - Each tier has either or both SSD and SAS drives. SSD drives are used for caching in all three tiers.
  - For future expansion, each tier can contain any combination of SSD, SAS, and SATA drives to realize the desired performance and capacity specifications.
- Fibre Channel connection to storage
- Plug-ins to VMware and Citrix virtualization environments
- Control from OpenStack, Cloud Stack, or from the Service Provider’s management platform via CloudByte’s REST API

The reference architecture is depicted below.



### 3.2 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, three controllers will be configured in an Active / Active group. As the total estimated controller load is around 40% of a single controller, loss of even two controllers will not result in an outage and will not impact performance. The additional reliability for the entire array is a significant benefit for the additional cost of a controller. Further, there is significant headroom for adding new, high IOPS tiers or for adding additional capacity on the existing tiers.

### 3.3 Control and Data Traffic

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

- Control traffic flows over 1 Gigabit Ethernet links.
- Data traffic flows over 8 Gigabit Fibre Channel links.

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

### 3.4 Fibre Channel Data Network

Providing additional detail on the Fibre Channel network for data traffic:

- Fibre Channel flows over two redundant Fibre Channel switches.
- Each switch has twenty four (24) ports. Assuming multipathing, six ports on each switch are utilized for connection to the ElastiStor nodes. The remaining ports are available for connection to application servers.

## 4 CloudByte Software Components

This solution fully 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

#### ElastiReplicate

The ElastiReplicate add-on enables remote asynchronous replication between two ElastiStor nodes. This feature provides disaster recovery capabilities that extend the core capabilities of ElastiStor.

#### 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 that protect against the failure of up to 3 physical drives without losing data. ElastiHA can be controlled from both ElastiCenter and our REST API.

#### ElastiFC

ElastiFC enables ElastiStor to act as a block-level Fibre Channel target device. With the plug-in, ElastiStor's storage resources (ZFS volumes, also called ZVOLS) are presented to the server as Fibre Channel or iSCSI disk(s). Fully integrated with ElastiCenter and REST APIs, ElastiFC delivers all the functionality, including dynamically selectable performance and an easy to use interface to provision, monitor and manage ElastiStor as a block-level Fibre Channel target. ElastiFC is also integrated with the ElastiHA add-on, delivering high availability for groups of up to 4 nodes.

## ElastiPlugins

ElastiPlugins provide integration and interoperation with virtualization environments and cloud stacks:

- VMware vCenter: The plugin allows admins to create and manage performance-selectable VMs in real time right from the vCenter console.
- OpenStack Horizon: The plugin allows storage provisioning and management in real time right from OpenStack.
- Citrix CloudStack™ : The plugin allows storage provisioning and management in real time right from CloudStack.



## 5 Hardware BOM

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

Tier / Component	Description	Quantity
<b>ElastiCenter</b>	Dell PowerEdge™ R520 Rack Server with 1P, 6 Cores, 32G RAM	1
<b>ElastiStor</b>	Dell PowerEdge™ R720 Rack Server with 2P, 8 Cores each, 256G RAM, 2 X 400GB MLC SSD	3
	10G Ethernet cards – Intel® Ethernet Converged Network Adapter X520	3
	8G HBAs – QLogic QLE2562 8GB Adapter	3
	SAS cards – LSI® SAS 9207-8e 6Gb/s SAS Host Bus Adapter – 2 per server for redundancy	6
<b>Control Plane</b>	Ethernet 1 GbE Fabric	
	Dell Networking S Series S60 1/10GbE access switch	2
	Dell 1GbE SFP+ transceivers (4 included with each switch)	8
<b>Data Plane</b>	Fibre Channel 8G Fabric	
	Brocade® 300 Switch, up to 24 ports (available from Dell)	2
	Brocade® 8G transceivers (multiple options, available from Dell)	12
<b>SAS Switch</b>	LSI® SAS 6160 6Gb/s SAS Switch	2
<b>Gold Tier</b>	Gold Storage Tier – Caching and Capacity	
	Storage Capacity and Read Cache – Seagate® 1200 SSD ST800FM0053 800GB	25
	Write Cache – Stec ZeusIOPS® 100 GB 3.5” Internal Solid State Storage Drive Z16IZF3E-100UCU	2
	Chassis – Dell™ PowerVault™ MD1220 direct-attached storage array	2
<b>Silver Tier</b>	Silver Storage Tier – Caching and Capacity	
	Storage Capacity – Dell 600 GB 10000 RPM Serial Attached SCSI Hard Drive	160
	Read Cache – Seagate® 1200 SSD ST800FM0053 800GB SSD	8
	Write Cache – Stec ZeusIOPS® 100 GB 3.5” Internal Solid State Storage Drive Z16IZF3E-100UCU	8
	Chassis – Dell™ PowerVault™ MD1220 direct-attached storage array	8
<b>Bronze Tier</b>	Bronze Storage Tier – Caching and Capacity	
	Storage Capacity – Dell 500 GB 7200 RPM Serial Attached SCSI Hard Drive	80
	Read Cache – Seagate® 1200 SSD ST800FM0053 800GB SSD	4
	Write Cache – Stec Z16IZF3E-100UCU 100GB SLC	4
	Chassis – Dell™ PowerVault™ MD1220 direct-attached storage array	4
<b>Cabinet</b>	Tripp Lite SR42UB 42U Rack Enclosure Server Cabinet	1

## 6 Benefits

### Foundation for New Services

This architecture enables a wide variety of tiered services to be created and defined, thereby maximizing service provider flexibility.

### Performance Needed for Enterprise Applications

This reference architecture provides the performance and flexibility needed to handle enterprise-class applications – whether they require minimum transaction times or maximum throughput.

### Strong fit for even leading-edge customer environments

By supporting multiple virtualization, multiple standards for ‘stacks’, and standard Ethernet networking, this reference architecture maximizes the number of customers that can utilize the Service Provider’s services.

## 7 About CloudByte

CloudByte is first to provide storage software for enterprise-grade clouds run by public and private service providers. Patent-pending storage software empowers service providers to define new classes of storage services, enables efficient customer onboarding without expert intervention, and elastically scales performance higher or lower on demand. Established in 2010 and managed by technology executives from companies such as NetApp, EMC, LSI, Cisco, Juniper, and Novell, 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](http://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