

---

# **Storage Architecture for Development Teams – Citrix Platform**

**CloudByte Reference Architecture**

---

# Table of Contents

---

<b>1</b>	<b>Executive Summary .....</b>	<b>3</b>
<b>2</b>	<b>Design Objectives .....</b>	<b>4</b>
<b>3</b>	<b>Solution Architecture.....</b>	<b>5</b>
<b>4</b>	<b>CloudByte Software Components .....</b>	<b>8</b>
<b>5</b>	<b>Benefits .....</b>	<b>9</b>
<b>6</b>	<b>About CloudByte .....</b>	<b>10</b>

## 1 Executive Summary

For any company, Developer productivity is critical. For hardware and software companies, however, Developer productivity is paramount. Yet, lack of adequate storage performance and storage backup reliability can easily hamstring the entire development organization. Since predictable and reliable storage performance is a key feature of CloudByte's architecture, companies with development organizations have been requesting a reference architecture that can handle the needs of multiple development departments on a single physical storage array.

This document outlines a reference architecture that is based on a Citrix virtualization architecture composed of Citrix CloudPlatform, CloudStack, and XenServer. The hardware architecture can be a wide range of x86 server blades attached to JBOD chassis via direct SAS connect. Storage, even for multiple departments, for around 200 VMs can begin with a modest number (around a dozen) SAS drives. As the x86 server blades are configured in a High Availability layout, Developers can expect high reliability and availability. As workload ramps, performance can be increased non-disruptively simply by adding more SSD drives for caching and by adding more components to the existing configuration.

In summary, this reference architecture provides the following benefits:

- Development Cycle Accelerated
  - The productivity of each Developer is significantly enhanced through consistent, predictable development process execution. Further, by eliminating the possibility that a single “noisy neighbor” development operation such as a build would compromise an entire storage array and by adjusting storage resources “on the fly” to respond to development needs, Developer team and department productivity is maximized. Taken as a whole, CloudByte helps accelerate the overall development cycle.
- Very high VM density
  - On just a handful of physical drives, this reference architecture supports approximately 200 development Virtual Machines. This very high VM density is a highly cost-efficient solution for a development team or for an entire development organization with multiple departments / development teams.
- Developer Code Protection
  - With backup granularity at the VM level, critical development work is better protected at the VM level than ever before. As backups are VM-consistent and aligned with the server VM, CloudByte ensures that precious development work is protected and that rollbacks are convenient.

## 2 Design Objectives

Typically, the Developer's IT team will provides infrastructure to multiple internal departments for their development and test environments. This leads to the following issues with existing storage architectures:

- To avoid a single “noisy neighbor” development run (such as a compilation of a release) slowing down all development, each department typically requests a dedicated storage array.
- Yet, even within a department's dedicated array, a single development run could consume the array's resources and slow down other development jobs.
- While storage workload demands vary greatly based on the development cycle, legacy storage arrays are not dynamically reconfigurable and therefore slow overall workcycles.
- As the granularity of legacy storage arrays are typically at the volume or LUN level, VM-consistent backups are not reliably attainable. This exposes critical development work to potential loss in the event of a disaster or need for a code rollback.

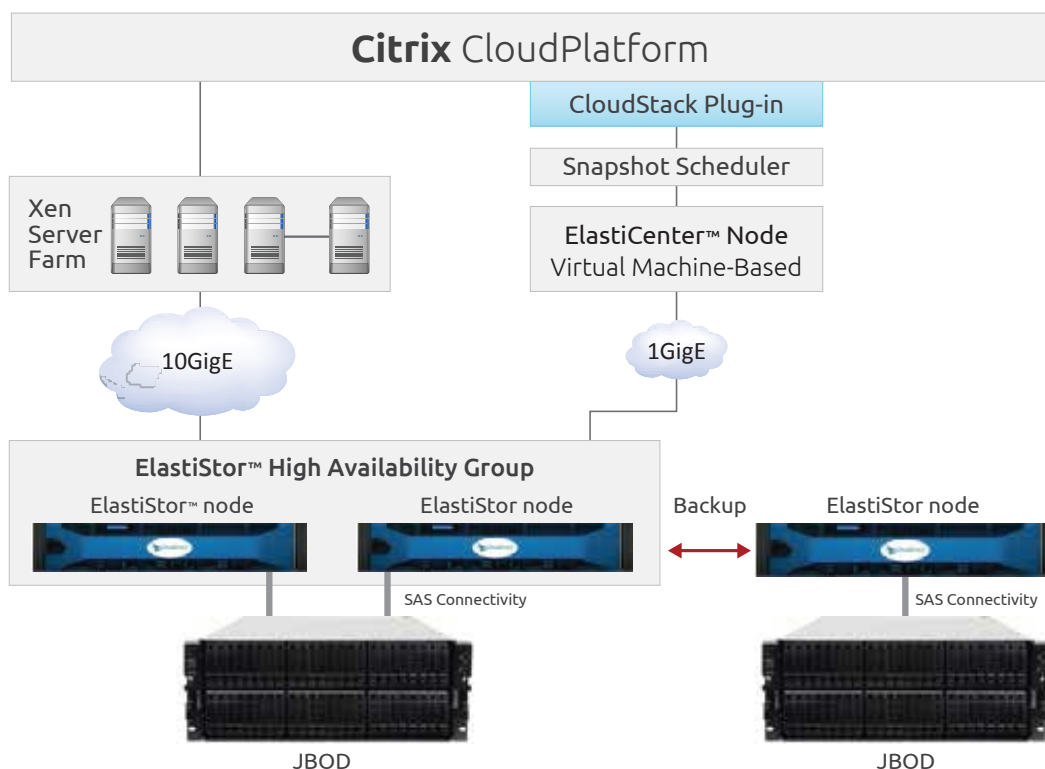
### 3 Solution Architecture

#### 3.1 Overview

Running on commodity hardware, CloudByte storage software products are the perfect fit for development private clouds. With CloudByte, IT teams can provide a multi-tenant shared storage infrastructure with guaranteed performance, VM-consistent backups, and flexible provisioning options tailored to each of their internal customers – developer departments.

From a hardware perspective, CloudByte ElastiStor™ storage nodes run on standard x86 server blades. An ElastiCenter™ management node can run on a standard x86 server blade or in a Virtual Machine. Storage JBODs are attached to the server blades via standard SAS cables. The primary and backup JBODs each contain a total of only 12 physical drives. Standard Ethernet switching is used to interconnect among all components.

From a software perspective, there is seamless compatibility with the Citrix CloudPlatform framework, Cloud Stack, and the Citrix XenServer virtualization platform. Requests for VM snapshots are made from the CloudPlatform framework via the CloudByte CloudStack Plug-in to the ElastiCenter management system. Storage requests for each VM are made from Citrix XenServer and then relayed via the CloudPlatform to ElastiCenter. This configuration is laid out in the diagram below.



### 3.2 Developer-specific Functionality

This reference architecture provides the following major functionality and features of primary significance to development teams:

- **High VM density:** The storage requirements of multiple development departments (approximately 200 VMs) can be met from a single platform with around a dozen physical drives.
- **VM-consistent Backups:** One key capability of CloudByte is granularity of backups by Tenant Storage Machine (TSM). Each TSM's backup parameters (hourly, daily, etc.) can be specified independently of others. When aligned with a corresponding Development Virtual Machine, this provides Developers with the capability to roll back to any previous version. This ensures that critical development productivity is maximized.
- **Discrete virtual arrays:** The solution provides storage isolation among departments both for compiler job integrity and for security.

### 3.3 Operational Functionality

This reference architecture provides the following major functionality and features that enhance operational convenience and reduce operational cost:

- **Empowered IT:** IT teams are able to access all management functionality – including performance analytics – via CloudByte's easy to use GUI.
- **Flexible provisioning:** Administrators can dynamically provision IOPS, throughput and latency for any development process or testing workload.
- **Self-Service:** Developers can also safely adjust some storage parameters themselves. This 'self-service' model enables the IT team to focus on overall system planning and performance.

### 3.4 Core Storage Functionality

This reference architecture provides the following major core storage functionality and features:

- **Storage Resilience:** With VM-consistent snapshots, replication to a backup array, and protection against data corruption, this reference architecture meets the resiliency requirements of IT teams. This solution utilizes CloudByte asynchronous replication to maintain a disaster recovery ElastiStor node separate from the production node. Should a disaster strike the primary node, then the backup node can easily be put into production, thereby avoiding disruption to critical development work.
- **Full storage functionality:** Enterprise-grade storage functionality such as RAID, compression, and protection against data corruption are provided.
- **Flexible Scalability:** The solution provides the ability to scale capacity (via adding JBODs and drives) and performance (adding more ElastiStor controller nodes) independently.

- No vendor lock-in: Performance and capacity can be independently scaled with industry-standard hardware. With CloudByte, resources can be seamlessly added to an existing storage array in a non-disruptive manner.

### 3.5 High Availability Group

A CloudByte High Availability Group 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, two controllers will be configured in an Active / Active pair. As the total estimated controller load is minimal, most x86 blades will handle the workload. This provides very good protection against the failure of a single server blade. Further, there is significant headroom for growth in terms of VMs and in storage capacity.

### 3.6 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 10 Gigabit Ethernet links.

Since control traffic does not consume any data bandwidth, this configuration maximizes performance and throughput for data traffic. Further, by using lower-capacity links for control traffic, cost is minimized.

## 4 CloudByte Software Components

This solution utilizes many components of the CloudByte product family.

### 4.1 Citrix-Specific Products

#### ElastiPlugins

ElastiPlugins provide integration and interoperation with virtualization environments and cloud stacks. The following products are part of this reference architecture:

- Apache CloudStack™ : The plugin allows storage provisioning and management in real time right from CloudStack and/or Citrix CloudPlatform.

### 4.2 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.

### 4.3 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.

### 4.4 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 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 Benefits

### Development Cycle Accelerated

The productivity of each Developer is significantly enhanced through consistent, predictable development process execution. Further, by eliminating the possibility that a single “noisy neighbor” development operation such as a build would compromise an entire storage array and by adjusting storage resources “on the fly” to respond to development needs, Developer team and department productivity is maximized. Taken as a whole, CloudByte helps accelerate the overall development cycle.

### Very high VM density

On just a handful of physical drives, this reference architecture supports approximately 200 development Virtual Machines. This very high VM density is a highly cost-efficient solution for a development team or for an entire development organization with multiple departments / development teams.

### Developer Code Protection

With backup granularity at the VM level, critical development work is better protected at the VM level than ever before. As backups are VM-consistent and aligned with the server VM, CloudByte ensures that precious development work is protected and that rollbacks are convenient.

## 6 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](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