

CloudByte ElastiStor Release Notes

Build: 1.4.0.879

Copyright 2016 CloudByte Inc.. All rights reserved.

Information in this document is subject to change without notice. The software described in this document is furnished under a license agreement or nondisclosure agreement. The software may be used or copied only in accordance with the terms of those agreements. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of CloudByte Inc..

CloudByte Inc.

4675 Stevens Creek Blvd, Suite 122

Santa Clara, CA 95051

USA

+1-(408)-663-6900

info@cloudbyte.com

www.cloudbyte.com

CONTENTS

CloudByte ElastiStor 1.4.0 Patch 5 release highlights	4
Release Highlights	4
What's New in CloudByte ElastiStor 1.4.0	4
System requirements	6
Hardware Compatibility List	7
Host Interoperability matrix	7
Upgrading to CloudByte ElastiStor 1.4.0 Patch 5	7
Resolved issues	8
Limitations	10

CloudByte ElastiStor 1.4.0 Patch 5 release highlights

Welcome to the CloudByte ElastiStor™ 1.4.0 Patch 5 Release Notes.

CloudByte ElastiStor 1.4.0 Patch 5 is a stability release including major performance enhancements and various bug fixes. This document explains the release highlights.

Installing CloudByte ElastiStor 1.4.0 Patch 5

Go to the following URL to access the instructions to install: <http://www.docs.cloudbyte.com/install/>

Build Description

Version	Build
1.4	1.4.0.879

Release Highlights

Following are the high points of this release:

- Better iSCSI I/O handling
- HA optimizations
- Usability and supportability improvements

What's New in CloudByte ElastiStor 1.4.0

Assigning real IP spaces

Multiple client virtual machines can access ElastiStor Nodes on the same IP addresses.

Backing up interfaces from ElastiCenter

You can establish a highly available network connection by configuring redundant (backup) interface NIC and LAGG right from ElastiCenter. The redundant interfaces protect against network failures.

Command line utilities to

- Display FIB ID of VSM. The utility can be found in the location `/usr/local/cb/bin/jls.py`
- Collect logs. The utility can be found in the location `/usr/local/cb/bin/collectlog.py`

Space reclamation

In the current release, the space reclamation process (UNMAP) has been regulated to ensure there is no performance impact on other volumes in the Node. This is achieved by allowing you to configure the number of blocks to be reclaimed in one request.

Alert emails

You are alerted when

- There is a CDB error logged in `/var/log/messages`.
- You login/logout the iSCSI targets:
 - `es.volume.iscsi.connected`
 - `es.volume.iscsi.disconnected`
- The volume WRITE latency exceeds 15 ms: `es.volume.qos.latency.high`

Activating primary VSM

You can activate the primary DR VSM even if data transfer is not enabled.

Hardware-based encryption/ Support for self-encrypting drives (SEDs)

The current release of ElastiStor introduces a hardware-level disk encryption by supporting the self-encrypting drives.

Space reclamation

On deletion of files on the Client filesystems (Ext4/NTFS), you can reclaim space on the thin provisioned iSCSI LUNs. This is achieved by supporting the SCSI UNMAP CDB.

GELI Encryption

CloudByte ElastiStor 1.4.Patch 2 has introduced a software-level disk encryption mechanism using the block device-layer disk encryption system GELI.

NFS home directory

The latest release of CloudByte ElastiStor lets you use a NFS share like a home directory. The users on the common NFS share do not have access to delete the other user files.

QoS Alerts

You are alerted when there is a variation in the quality of service (QoS) you have opted for. The following two Global settings have been added to support the new alerts:

- `alert.qos.throttling.percentage`
- `alert.qos.throttle.count`

GUI changes to distinguish your ElastiStor

ElastiCenter uses different color schemes for the Community Edition and Enterprise Edition.

Support for LDAP authentication

You can use an external LDAP server such as Microsoft Active Directory to authenticate ElastiCenter users.

Global Settings to specify provisioning without specifying IOPS

The new global settings option `provisioning.uncontrolled.iops` lets you enable/disable provisioning with uncontrolled IOPS.

VSM Webshell Terminal on ElastiCenter

An option for advanced troubleshooting, where you might have to access the VSM terminal, for example, to check if your iSCSI service is running. This option lets you monitor various processes and services, snapshot list, volume list, etc .

Option to check if the Disks are Write Cache Enabled

The new feature lets you see if the disks are write cache enabled.

Faster and resilient HA

Significant improvement in HA takeover and giveback time while ensuring data continuity.

Improved QoS Control

Offers tighter control over QoS. Also, you can define throughput for a Storage Volume.

8Gbps FC target support

Now, you can configure 8Gbps FC targets (QLogic).

ElastiCenter HA

ElastiCenter is now available in High Availability mode.

Detect faulty disks

When a disk in a Pool goes bad, ElastiCenter displays the bad disk in Alerts.

Replace faulty disks

You can now replace a faulty disk from ElastiCenter.

Terminology changes

The following terminology change applies in this release: Tenant Storage Machine (TSM) to Virtual Storage Machine (VSM).

System requirements

The following topics explain the system requirements for installing CloudByte ElastiStor.

- CloudByte Node is built on commodity servers.

- The connectivity between Controller and Storage is either on Ethernet (iSCSI) or on SFF-8088 SAS cables.
- ElastiStor OS is built using FreeBSD as the base platform and therefore CloudByte ElastiStor's Hardware Supported List is similar to that of FreeBSD Release version.

Minimum hardware requirements to build CloudByte ElastiStor

- Supported AMD64, Intel 64 (x86-64), or i386 Processor and System Board (with Dual Core Processor and 8GB of RAM)
- Four 1 Gigabit Ethernet Controller Cards
- One or more SAS HBAs, depending on the number of disk arrays connected
- OS boot disk drive size of 20 GB. However, the recommended size is $2.4 * \text{RAM size} + 33\text{GB}$

Minimum hardware requirements for CloudByte ElastiCenter

- Supported AMD64, Intel 64 (x86-64), or i386 Processor and System Board (with Dual Core Processor and 8GB of RAM)
- OS boot disk drive size of 20 GB. However, the recommended size is $2.4 * \text{RAM size} + 33\text{GB}$
- 1 Gigabit Ethernet Controller Cards

For supported list of

- AMD64 Processors, see [FreeBSD 9.0 supported list of AMD64 processors](#)
- i386 Processors and System Boards of CloudByte ElastiStor, see [FreeBSD 9.0 supported list of i386 Processors](#)
- Disk Controllers of CloudByte ElastiStor, see [FreeBSD 9.0 supported list of Disk Controllers](#)
- Ethernet Interfaces of CloudByte ElastiStor, see [FreeBSD 9.0 supported list of Ethernet Interfaces](#)

Hardware Compatibility List

For the complete list of compatible hardware, see [Hardware Compatibility List](#).

Host Interoperability matrix

Access the host interoperability matrix at the following URL : <http://www.cloudbyte.com/support/host-interoperability-matrix/>

Upgrading to CloudByte ElastiStor 1.4.0 Patch 5

Versions from which you can upgrade

If you are using the following versions of ElastiStor, you can upgrade to ElastiStor 1.4.0 Patch 5.

- 1.4.0.7004
- 1.4.0.6018

Installation options

You have the following two installation options:

- **Upgrade:** Upgrades from any of the versions mentioned in "[Versions from which you can upgrade](#)" to ElastiStor 1.4.0 Patch 5.
- **New Installation:** Replaces the current version of ElastiStor with ElastiStor 1.4.0 Patch 5.

When you upgrade, ensure that you upgrade ElastiCenter first, and then proceed to upgrade the Node.

The workflow is as follows:

1. Back up ElastiCenter. For details see the section [Backing up ElastiCenter Configurations](#)
2. Upgrade ElastiCenter using the steps in the section [Upgrade ElastiCenter](#).
3. Log in to the Node as root and then run the following command to download the Volume configuration script: `fetch http://cloudbyte.com/downloads/updateVolSize.py`
4. Run the following command: `python updateVolSize.py`
5. Move the Node to Maintenance Mode. For details see the section Bringing Node to maintenance mode in [Configuring High Availability](#)
6. Upgrade the Node using the procedures in the following section [Upgrade ElastiStor](#).
7. Move the Node to Available mode. For details, see the section Making a Node available in [Configuring High Availability](#)
8. Refresh hardware for the upgraded Node.
9. Ensure that the Pools on the Node are online.
10. Repeat the procedures (from step 3 to step 7) for the other Nodes.

Resolved issues

The following issues that existed in the previous releases have been fixed as explained in the description:

EC reporting node offline intermittently

Issue fixed by allowing duplicate node offline alerts that are five minutes apart.

VMWare certification 6.0 NFS and iSCSI with VAAI

Certified with ESXi6.0 on NFS & iSCSI and notified on the VMWare HCL.

VMWare certification 5.5 NFS and iSCSI with VAAI

Certified with ESXi5.5 on NFS & iSCSI and notified on the VMWare HCL.

EC backup issue

Fix for backup failed alert in case of root partition is full.

EC and ES log rotation issue

Hardware acceleration UI Support

UI support for enabling/disabling ATS, XCOPY and WRITE-SAME.

Enhanced monitoring for the SCSI abort

Monitoring the VSM log for the ABORT TASK command sent by ESX and raising an alert.

Duplicate alerts suppression

Generate alert only when there is a state change

SAS cable disconnect alert

Disk removal event is generating SAS cable disconnect alert has been addressed.

Pool vdev metrics

Fixed by Purging vdev metrics periodically.

HA time optimization

The following improvements have been made:

- Use a maintenance thread for converting LUNs between Fake/Real.
- Split the zil replay of volumes outside of the zpool import.
- Provide granular timings of the import and export.
- Optimize the start of the services within the VSM to reduce the load on the CPU during HA.
- Use refresh instead of restart for converting LUNs from fake to real.
- Sending alert (node offline) in a different thread

Istgt Stability

- New troubleshooting command for listing the details of a LUN.
- User should be forced to set initiator group to None before disabling iscsi on Storage Volume.
- LUN Power-on RESET UA will be only sent from the LUN where it was received.
- iSCSI client connectivity status.
- iSCSI control refresh operation optimization.

TSM troubleshooting section in UI

All the commands will be executed inside the appropriate jails.

Chrome Browser Issue

Fixed the chrome browser issue,EC can be accessed by latest chrome versions.

IO timeout on disk firmware

Core Dump generation during system freeze

Implemented an IPMI command that sends a message to the system to dump the core in NMI context.

Bad disk Reporting

A bad disk indicated in the UI disk array page (both in nodes and HA group).

Requires additional time to start the NFS service in a scaled setup

In an HA scenario with a very large number of VSMS, additional time is required to start the NFS service in the VSM.

Cannot list iSCSI initiators

When a Pool is taken over by the peer Node, ElastiCenter does not display the iSCSI initiators of the Volumes in the taken over Pool.

Limitations

The following limitations are known to exist in the release:

Issues with takeover after VSM network failure

Assume that in an HA Group, VSM network for both the Nodes (Node 1 and Node 2) are down. Later, if the VSM network for Node 1 is up, it fails to take over the Pools in Node 2 (whose VSM network is still down).

Issue with Alerts in ElastiStor Appliance ESA-A50

Assume power failure on one of the Nodes in ESA-A50. In this case, ElastiCenter generates the power failure Alert for the Node that is configured as primary even if the power failure is on the secondary Node.

Storage configuration queue goes into the hung state

Changing user authentication from Active Directory to local, results in the storage configuration going into the hung state.

Workaround

1. SSH into ElastiStor Node.
2. Run the following command: `service cbc_storageconfiguration onerestart.`

Issue with validating secondary IP address of the VSM

The secondary IP address of a VSM is not validated. This may result in an IP address conflict making the VSM inaccessible through secondary IP address.

No alerts if management network of primary Node is down

In an EC-HA environment, ElastiCenter fails to generate an alert on failure of management network on the primary Node.

Cannot load Manage CIFS configuration page on NFS Clone

Assume you have cloned the Snapshot of an NFS Volume . In this case, the *Manage CIFS configuration page* in ElastiCenter for the NFS Clone does not load.

Workaround

Enable CIFS service for the NFS Clone.

Issue with ElastiCenter configuration data sync

Assume that a Node (that hosts the secondary ElastiCenter) in the EC-HA environment is deleted. Adding this Node back to the HA Group, fails to sync data with the primary ElastiCenter.

Workaround

1. Move the newly added Node to Maintenance state.
2. Move the Node to Available state.

Cannot assign permissions on Active Directory objects

The current release does not support access control in Active Directory.

Issues with VSM migration

- VSM name cannot contain the special character underscore (_).
- A delay may be experienced in activating the migrant VSM when there is a large file WRITE in progress. This may also disconnect I/Os on few connected clients.

Perform the following procedures to reduce the delay:

1. Stop the I/Os in progress.
2. Migrate VSM.
3. Restart the I/Os.

Requires restart of iSCSI service

For iSCSI custom configuration changes (such as LU workers threads, physical record length) to take effect, you have to manually restart the iSCSI services.

Change in disk ownership

In the case of ElastiStor high availability process performed in heavy load conditions, the ownership of Windows clustered disks changes automatically.

The changes in authentication of a CIFS Volume do not reflect in the clients

If the issue occurs, do the following.

Workaround

Do either of the following:

- Clear the Cache
- Re-login to the Client

CloudByte recommends you to clear your Windows cache frequently to get the latest updates from ElastiCenter.

Requires disabling of CIFS I/Os to activate Migrant VSM

You have to disable I/Os on CIFS Storage Volumes to activate a Migrant VSM.

Cloned Storage Volumes cannot inherit the properties from Snapshot

A cloned Storage Volume inherits the parent Volume properties such as Sync and READ-only instead of preserving the snapshot's properties.

Space reclamation is not supported for Windows 2008 R2 SP1 Client

Deleting data from a Storage Volume mapped to Windows 2008 R2 SP1 Client might not free up the physical storage.

Requires rediscovery of LUN to enable space reclamation

You have to rediscover iSCSI LUNs on the Client (to detect a disk as thin-provisioned) to reclaim dead space.

Remounting Linux Clients on iSCSI LUNs

To reclaim dead space for Linux Clients using Ext3/Ext4 filesystem on iSCSI LUNs, remount using the mount option Discard.

Issue with SED status in ElastiCenter

If you re-insert SED, ElastiCenter shows the drive status as removed.

Online VSM migration is not supported in the case of WinSCP

Online VSM migration fails when file is copied using WinSCP to a Linux client.

Issue with SES blink

In the case of SES blink , manual intervention is required to uncheck the blink option in ElastiCenter.

Setvaliddata operation fails on CIFS share

On Fsutil, *Setvaliddata* operation fails for CIFS share.

IOPS and throughput cannot be set on the same Storage Volume

Current release supports only setting either IOPS or throughput on the same Storage Volume.

ElastiCenter fails to dynamically update status of physical components after reconnection

In the case of physical operations such as disk pull or disk push, ElastiCenter does not show the status of physical components such as disk arrays.

Workaround

Refresh hardware

NLM services are not supported

Network Lock Manager is not supported in this release.

ACL not supported for NFS

Access Control List (the user list specified while exporting NFS Volumes) is not supported.

NFS client performance issue with burst I/O

NFS client performance might degrade with burst I/O.

You are recommended not to utilize more than 90% capacity of the NFS Volumes.

Access/IO on NFS Volumes with greater than 90% capacity utilization can result in slowdown.

Workaround

Increase the quota of NFS volumes after you receive capacity alerts.

Frequent configuration changes might interrupt I/O on NFS/CIFS Volumes

CloudByte recommends not to perform frequent configuration changes (at the tenant side) such as addition, deletion, or property changes of NFS/CIFS Volumes.

I/O on CIFS share are disrupted during failover and giveback

You might have to restart the I/O on the CIFS shares post a failover and giveback.

No CIFS support for Hyper-V

This release does not support VHD creation on CIFS share.

DOS Charset and timeserver configurations on ElastiCenter are not operational

DOS Charset and timeserver configurations made from ElastiCenter do not reflected on the client side for CIFS Volume.

Delayed resilvering

Resilvering takes longer duration to complete.

Volume creation with very low IOPS (for example, 1-10) is not recommended

Low IOPS might result in hung operations on the Storage Volume.

Workaround

Bring the Node 2 to maintenance mode.

SAS Multipathing

When the path or disk is back after a failure, you must refresh hardware at the node level for

- The path or disk to reflect in ElastiCenter
- Configuring Pools

FC LUNs require multipathing on the client side

ElastiStor FC LUN must be managed by multipathing software for HA to work.

MTU not supported for Realtek Interfaces

Maximum Transmission Unit (MTU) is not supported for Realtek Interfaces (for example re0).

Cannot revise MTU while mapped to VMware ESX

Revising MTU while connected to VMware ESX disconnects storage. You have to reconnect.

Requires hardware refresh when you add disks

When you add disks, they do not appear in ElastiCenter. You must refresh hardware for them to reflect.

Issues with graphs

- Changing global settings may not render graphs for VSMs and Pools in the specified time.
- The QoS graphs might not display correctly for cloned Storage Volumes.

Storage Volume backup resumability

In the case of disaster recovery, assume that there is an interruption in data transfer of a volume (block device such as FC and iSCSI) either due to reboot of the remote Node or due to HA take over. Data transfer starts from the beginning instead of resuming from the point of interruption.

Requires respecification of authentication group and methods in ElastiCenter (post DR)

When you configure iSCSI Volume after DR VSM activation, you have to respecify authentication group and authentication methods in ElastiCenter.

Issues with IP address entry in Host allow/deny list for CIFS

- Same IP addresses are accepted in both hosts allow and hosts deny list in ElastiCenter for CIFS.
- IP address changes in both hosts allow and hosts deny list take a while to reflect on the Client.

I/O interruption while adding users

Adding a new authentication group or users to CIFS share interrupts the I/Os in progress.

Issue with disabling recycle bin of CIFS share

Even with the recycle bin parameter disabled, the deleted files from the CIFS share are present in the `.recycle` folder.

Secondary ElastiCenter fails to take over on network failure

Assume that the management network failed on the primary ElastiCenter. In this case, the secondary ElastiCenter fails to takeover. You have to wait until the management network on primary ElastiCenter is active.