

CloudByte ElastiStor

Release Notes

Build: 2.0.0.1264

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 2.1 release highlights	4
Release Highlights	4
What's New in CloudByte ElastiStor 2.1	4
System requirements	6
Hardware Compatibility List	7
Host Interoperability matrix	7
Resolved issues	7
Limitations	8

CloudByte ElastiStor 2.1 release highlights

Welcome to the CloudByte ElastiStor™ 2.1 Release Notes.

CloudByte ElastiStor 2.1 is a stability release including major performance enhancements and various bug fixes. This document explains the release highlights.

Installing CloudByte ElastiStor 2.1

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

Build Description

Version	Build
2.1	2.0.0.1264

Release Highlights

Following are the high points of this release:

- Constant time resilvering
- ZFS Feature
- All new UI Design

What's New in CloudByte ElastiStor 2.1

Constant Time Resilvering

In this release, time taken for resilvering data onto new disk is almost independent on the capacity of the pool. Also, duration of resilvering process has been reduced.

New UI design and facilitation for StAAS portal

Usability of ElastiCenter component has been improved with new design of pages and improved performance, with LDAP support for user login.

Root-on-ZFS feature

The latest release enables administrator to use software based mirroring for ElastiStor and ElastiCenter installation instead of hardware mirroring.

Handling high latent disks for steady performance

One or few high latency disks bringing down the pool performance has been addressed in this release.

Improvements in High Availability functionality

The new version of Elastistor reduces the time taken to bring up the services on peer node during a node failure due to H/W or S/W.

Improved IOPS control

IOPS accounting is now done at 16KB IO size, considering a write as two standard IOPS.

Scalability and Feature improvements with NFS protocol

Scalability limits for NFS connections has been made as 20K per volume, with added functionality for file locking mechanism. Also now you have an option for controlling permissions on NFS volumes. Multi-port NFS server is made available in this release.

Support for backing up data to S3 or S3 compliant providers

You can now backup snapshots by sending them to cloud using S3 APIs.

Better hardware support

Quad-port SAS HBAs support has been added to enable administrator move away from the dependency of SAS switches. QLogic Quad-port FC cards are also supported by this version.

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)
- 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

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

Resolved issues

- Node crash when LAGG with LACP configured on Mellanox cards
- Issues related to reports have been fixed.
- Lock functionality in NFS is supported in this version
- Pool export issues are taken care when deletion of data is going on in NFS volume
- Alerts related to disks which doesn't have sg_vpd command support is fixed
- The new version allows E-Mail alerts to be sent with configured sender address
- Compression is enabled by default in this version
- Alert generation when meta-vdev become full by 80%
- Pool performance degradation due to latent disk in the pool is fixed in this release
- Throttling IOs based on 16KB IO size, and 1:2 penalty for read/write operations

Limitations

The following issues are known to exist in this release:

- During HA, assigning IP address for VSM can fail
Assign IP address and run start TSM script
- Deletion of data in NFS volume and destroying a NFS volume from UI at the same time can sometimes cause controller reboot.
- ‘Refresh’ button is not working for some storage capacity graphs.
- Alerts aren’t generated when FC ports are down.
- During disk labeling, UI hangs if peer node is down.
- Disk used for resilvering will be available to create raid group or pool or further bad disks replacements, until resilvering process completes.
- IOPS doesn’t get updated in VSM dashboard
IOPS details are available in Reports
- Alerts doesn’t get updated in VSM & volume dashboards
Alerts details are available in Reports
- The current release does not support Cloudstack plug-in.
- Backup QoS configured in UI is not honored
- When a node of Windows cluster with FC volumes is down, upon HA on storage controller, the node may crash.
- While doing “Cluster validation” from Windows on FC volumes, the node may reboot.
- Second cable may not be in ACTIVE state after both the cables pull/push in LACP configuration.
- Vlan interface creation/deletion may fail if one of the nodes is unreachable.
When one of the nodes in the HA group is down, Vlan creation/deletion fails.

-
- Account super admin could not edit the volume's properties:
 - Node health widget representation looks static always though one of the node is unhealthy
 - Issue while updating iSCSI initiator group for the first time
Change the initiator group to NONE, and set it back to required initiator group.
 - Clone creation is failing for FC volumes
-