



Hewlett Packard
Enterprise

NimbleOS 5.3.0.0 Release Notes

Version 5.3.0.0

Published October, 2020

Legal Notices

© Copyright 2020 Hewlett Packard Enterprise Development LP. All rights reserved worldwide.

Notices

The information contained in herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated. Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Publication Date

Tuesday October 6, 2020 12:19:40

Document ID

old1592810548757

Support

All documentation and knowledge base articles are available on HPE InfoSight at <https://infosight.hpe.com>. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: support@nimblestorage.com

For all other general support contact information, go to <https://www.nimblestorage.com/customer-support/>.

Contents

| | |
|-------------------------------|----------|
| NimbleOS 5.3.0.0..... | 4 |
| Important Update Note..... | 4 |
| Special Notes..... | 4 |
| New Features in 5.3.0.0..... | 7 |
| Documentation..... | 8 |
| Verified Update Paths..... | 9 |
| Known Critical Issues..... | 12 |
| Resolved Critical Issues..... | 19 |
| Resolved Issues..... | 21 |
| Known Issues..... | 24 |

NimbleOS 5.3.0.0

| | |
|------------------|----------------------------------|
| Version: | 5.3.0.0 |
| Revision: | Tuesday October 6, 2020 12:19:40 |

The release notes describe the major changes, fixes, and known issues for this release of the NimbleOS. They do not include all individual fixes and internal changes.

For technical support, contact HPE Nimble Storage Support at:

<mailto:support@nimblestorage.com>

877-3-NIMBLE (877-364-6253), option 2.

Important Update Note

Updating NimbleOS can involve an update to component firmware on the standby controller. This can cause an email alert and automated case indicating "Standby Controller Not Available" when the firmware update process takes longer than five minutes. This is expected behavior and does not affect data services. At the end of the software update, you can check status of both controllers in the Web UI under **Manage > Hardware**. One controller will be ACTIVE and the other STANDBY under normal operating conditions following a successful software update.

All third-party software notices can be found on HPE InfoSight (<https://infosight.hpe.com>) on the **Resources > Documentation** page:

<https://infosight.hpe.com/resources/nimble/docs>

The Documentation page also includes the *General Terms and Conditions* document. You can display this document by performing the following steps:

- 1 In the navigation pane on the HPE InfoSight Documentation page, scroll through the Document Type list and select Support Policy.
- 2 In the page that appears, select General Terms and Conditions. This document opens in a browser tab.

Special Notes

| Note | Description |
|-----------------|---|
| CRITICAL | HPE Nimble Storage continues to qualify configurations between releases. The Validated Configuration Matrix provides information about validated configurations and is updated frequently. It is a good practice to check your system configuration against this online tool. The Validated Configuration Matrix tool is available on HPE InfoSight: https://infosight.hpe.com/resources/nimble/validated-configuration-matrix |
| CRITICAL | Arrays must be running NimbleOS 5.0.6.0 or later to update to NimbleOS 5.3.0.0. |
| CRITICAL | An extended data services outage may occur with MS iSCSI initiator and Intel NICs using the built-in Windows driver e1q60x64.sys (version 11.0.5.21/11.0.5.22). If you encounter this problem, please update your system to use the latest Windows driver. |

| Note | Description |
|------------------|---|
| CRITICAL | A service outage may occur on Windows 2012 R2 hosts using Emulex or Broadcom Fibre Channel HBAs with firmware/driver prior to 11.2. Update the Emulex or Broadcom firmware/driver to 11.2 or later |
| CRITICAL | <p>Due to a known Red Hat Enterprise Linux bug 1002727, while running virtualized in VMware ESX, manually rebooting the active controller in presence of heavy IOs using the reboot --controller command on a Fibre Channel array may trigger an incorrect retry initiated by RHEL guests running the following kernel versions:</p> <ul style="list-style-type: none"> • 6.4 and earlier • 6.5 without the patch • 7.0 without the patch <p>This incorrect retry logic may lead to unexpected application behavior. In these environments, we recommend the failover command instead.</p> |
| CRITICAL | <p>Due to a known Red Hat Enterprise Linux bug 3550561, unexpected application behavior may occur on RHEL 7.5 hosts with kernel-3.10.0-862.3.2.el7 or derivatives using Emulex FC FCoE HBAs (lpfc driver) and raw devices. To avoid this issue:</p> <ul style="list-style-type: none"> • If running RHEL 7.6, update to kernel-3.10.0-957.el7 or later. • If running RHEL 7.5z, update to kernel-3.10.0-862.25.3.el7 or later. |
| CRITICAL | <p>As outlined in the current Validated Configuration Matrix, HPE Nimble Storage fully supports Windows guest operating systems on Microsoft Hyper-V, including Virtual Fibre Channel (VFC) connectivity and multipathing with HPE Nimble Storage DSM and VSS support. However, Linux guest operating systems running in Hyper-V VFC configurations are not qualified.</p> <p>Running Red Hat Linux guest operating systems with the “Linux Integration Services” kit installed, or with hv_storvsc drivers in such configurations can lead to Red Hat bug 1364282, which can cause an unexpected service outage.</p> |
| Important | Starting with NimbleOS 5.1.1.0, the size of the software package now exceeds 2 GB, which may lead to lengthier software download times. Previously, the sizes of the NimbleOS 5.0.x download packages were approximately 1.6 GB, and NimbleOS 4.x packages were approximately 900 MB. |
| Important | After completing the NimbleOS update for array groups configured for Synchronous Replication, download the corresponding version of the Synchronous Replication Witness software, and update the witness host. |
| Important | Microsoft Offload Data Transfer (ODX) is not supported if the destination volume has synchronous replication enabled. |
| Important | As of vSphere 7.0, VMware has discontinued the flex client. Consequently, the HPE Nimble Storage vCenter Plugin no longer supports the flex plugin for vCenter 7.0. |

| Note | Description |
|------------------|--|
| Important | <p>You can enable deduplication for CS1000, CS3000, CS5000, CS7000, CS700, and CS500 arrays on a volume only if the corresponding storage pool has a Flash to Disk Ratio (FDR) greater than 4%. To calculate the FDR, obtain the "Total array capacity (MiB)" and "Total array cache capacity (MiB)" values by using the HPE Nimble Storage CLI command <i>pool_name</i>. This command returns the Pool capacity (MiB), which is the "Total array capacity (MiB)", and the Pool cache capacity (MiB), which is the "Total array cache capacity (MiB)".</p> <p>Then perform the following calculation:</p> <p>FDR = "Total array cache capacity (MiB)" / "Total array capacity (MiB)" * 100</p> <p>If the array has sufficient capability for deduplication, the pool --info command will also show a value for dedupe capacity (MiB).</p> <p>Note On the HF20H, HF20, HF40, and HF60 platforms, pool --info displays "N/A" as the value for dedupe capacity (MiB). This because you can enable deduplication for the entire array.</p> |
| Important | <p>For connections to the NimbleOS GUI, you must have port 5392 open for the Group Management IP address and both diagnostic IP addresses.</p> |
| Important | <p>Numerous host integration toolkits are supported in NimbleOS 5.3.0.0. It is strongly recommended that they be installed on all Windows, Linux, and VMware hosts. For more information about supported toolkits, refer to the Validated Configuration Matrix, which is available on HPE Nimble Storage InfoSight:</p> <p>https://infosight.hpe.com/resources/nimble/validated-configuration-matrix</p> |
| Important | <p>HPE Nimble Storage recommends that you update to HPE Nimble Storage Windows Toolkit (NWT) 7.0.1 or later if you are using Microsoft VSS Synchronization and NimbleOS 5.1.4.200 or later.</p> <p>Using application consistent snapshots with earlier versions of NWT and NimbleOS 5.1.4.100 may result in the following error messages:</p> <ul style="list-style-type: none"> In the host's VSS requestor log (C:\ProgramData\Nimble Storage\Logs\VssRequestor.log): <p>PID:1996 TID:5752 ERR reqcommon. cpp:683 Request-Status=QueryStatus(), Function=pAsync->QueryStatus(), Error=VSS_E_PROVIDER_VETO, rc=SystemError, ca=ContactSupport</p> In the Windows event viewer: <p>event id 4100: EndPrepareSnapshots method: failed to find LUN s/n <SERIAL_NUMBER> on connected arrays. Make sure that the Nimble array version is compatible with this version of Nimble Windows Toolkit.</p> <p>event id 4170: Nimble VSS provider is not compatible with the current version of the Nimble array software(). Install appropriate version of the Nimble VSS provider.</p> <p>NWT 7.0.1 resolves this issue.</p> |

| Note | Description |
|------------------|---|
| Important | <p>HPE Nimble Storage Connection Manager (NCM) for VMware 7.0 is signed by VMware for ESXi 7.x. It can be installed through the VMware Update Manager or esxcli command without the --no-sig-check flag.</p> <p>See the NCM for VMware Release Notes 7.0 or later and the latest <i>VMware Integration Guide</i> for further details.</p> <p>To locate the latest version of the guide, log in to HPE InfoSight. Choose Resources > Nimble Storage Documentation. In the left pane, click Integration Guide, then click Connection Manager (NCM) for VMware. From the list displayed, choose the version of the guide that you want.</p> |
| Important | <p>Various timeout values affect HPE Nimble Storage targets from Windows/Linux hosts. Before you update the NimbleOS, install the HPE Nimble Storage Windows Toolkit (NWT) or HPE Nimble Storage Linux Toolkit (NLT) on the host or tune the timeout values. Timeout details for various operating systems can be found on HPE InfoSight under Resources > Documentation. From the HPE Nimble Storage Documentation page, locate the article you want.</p> <p>The following Knowledge Base articles and Integration Guides explain how to configure and verify host timeout settings for the major supported operating systems (OS):</p> <ul style="list-style-type: none"> For Windows, refer to KB-000052: Windows Host Disk Timeout Values. In the context of Microsoft Windows, the following article should also be considered: KB-000246 MPIO Timeout Parameters for MSDSM and NimbleDSM in Windows 2012 R2 For VMware, refer to the Common Tasks and Best Practices > Host Timeout Values section of the <i>VMware Integration Guide</i>. For Linux, refer to KB-000304: Linux Host Disk Timeout Values. |
| Important | <p>vVol VMs cannot be claimed after they are deleted from the downstream array.</p> <p>A vVol VM can be protected and may subsequently be replicated to a downstream array (as configured in the storage policy). In the case where this vVol VM was deleted, a supported “claim” workflow allows us to claim this vVol VM on the downstream array.</p> <p>Due to validation failures in the vCenter, this workflow is not supported if it is performed in an environment where the vCenter version is 6.5 or later.</p> <p>See the following VMware DCPN Ticket Reference: https://dcpn.force.com/TechnicalRequestCaseRedesignPartner?id=5000H00001JRKhf</p> |

New Features in 5.3.0.0

NimbleOS 5.3.0.0 introduces support for the following feature:

External KMIP Integration

NimbleOS now supports using an external key manager to manage the master encryption key. External key management is done via the industry standard KMIP protocol, and any KMIP supporting key manager may be used.

Documentation

These Release Notes and other user documentation are available on HPE InfoSight:

<https://infosight.hpe.com/resources/nimble/docs>

You can manually reach the documentation page by logging onto HPE InfoSight and selecting **Resources > Nimble Storage > Documentation**.

Document Search Interface

There are several methods you can use to locate the documents you need.

The **Nimble Storage Documentation** page provides a search interface that allows you to search for information across all documentation, including support and knowledge base articles, best practices, solutions and integration guides, product documentation, and configuration matrices.

To go directly to a document, use the navigation pane on the left side of the **Nimble Storage Documentation** page. The navigation pane organizes documents into categories, including:

- Document Type
- Nimble Software and Solutions
- Software Version
- Integration
- Platform

You can use the page scroll bar to move up and down the navigation pane.

Third-Party Software Notices

All third-part software notices can be found in the Documentation Portal on HPE InfoSight.

Here are the steps to manually access the third-party software notices.

- 1 Log in to HPE InfoSight (<https://infosight.hpe.com>).
- 2 From the menu, select Resources Nimble Documentation .
- 3 In the left navigation pane of the Documentation Portal, scroll through the Document Type section and select Support Policy.
- 4 From the list of documents, select General Terms and Conditions. The document opens in a new browser tab.

Core User Documentation

The following is the core user documentation for NimbleOS:

- *GUI Administration Guide*
- *CLI Administration Guide*
- *SNMP Reference*
- *Command Reference*
- *REST API Reference*

If you are using an HPE Nimble Storage dHCI-enabled array, you should also check the dHCI Deployment Guides and Getting Started Guide.

Workflow Documents

There are several workflow guides that contain procedures you can perform using either the CLI or the GUI. Each workflow guide covers a specific, frequently performed task related to HPE Nimble Storage products. Each task described by a workflow document is explained in detail in the *GUI Administration Guide* and the *CLI Administration Guide*.

Hardware

Documentation for all hardware components is available on HPE InfoSight. Click the Hardware Guide link in the **Document Type** category. Hardware documentation includes array and expansion shelf installation quick start guides, installation, upgrade, and replacement guides, and comprehensive hardware guides.

Host Integration Guides

Host Integration Guides are available from HPE InfoSight. To locate these documents on the HPE InfoSight **Documentation** page, scroll down the navigation pane to the section called **Integration Guide**.

Note A single Host Integration Guide supports multiple version of NimbleOS and the companion Integration Toolkit software packages. The version number listed on the guide might be different from the version numbers of the NimbleOS and Toolkit software packages that it supports.

Verified Update Paths

Table 1: From Versions 5.x

| From Versions 5.x | |
|-------------------|------------|
| From Version | To Version |
| 5.2.1.300 | 5.3.0.0 |
| 5.2.1.200 | 5.3.0.0 |
| 5.2.1.100 | 5.3.0.0 |
| 5.2.1.0 | 5.3.0.0 |
| 5.1.4.200 | 5.3.0.0 |
| 5.1.4.100 | 5.3.0.0 |
| 5.1.4.0 | 5.3.0.0 |
| 5.1.3.100 | 5.3.0.0 |
| 5.1.3.0 | 5.3.0.0 |
| 5.1.2.100 | 5.3.0.0 |
| 5.1.2.0 | 5.3.0.0 |
| 5.1.1.0 | 5.3.0.0 |
| 5.0.10.0 | 5.3.0.0 |
| 5.0.9.100 | 5.3.0.0 |
| 5.0.9.0 | 5.3.0.0 |
| 5.0.8.100 | 5.3.0.0 |
| 5.0.8.0 | 5.3.0.0 |
| 5.0.7.300 | 5.3.0.0 |
| 5.0.7.200 | 5.3.0.0 |
| 5.0.7.100 | 5.3.0.0 |
| 5.0.7.0 | 5.3.0.0 |

| From Versions 5.x | |
|-------------------|------------|
| From Version | To Version |
| 5.0.6.0 | 5.3.0.0 |
| 5.0.5.300 | 5.0.10.0 |
| 5.0.5.200 | 5.0.10.0 |
| 5.0.5.0 | 5.0.10.0 |
| 5.0.4.0 | 5.0.10.0 |
| 5.0.3.100 | 5.0.10.0 |
| 5.0.3.0 | 5.0.10.0 |
| 5.0.2.0 | 5.0.10.0 |
| 5.0.1.100 | 5.0.10.0 |
| 5.0.1.0 | 5.0.10.0 |

Table 2: From Versions 4.x

| From Versions 4.x | |
|-------------------|------------|
| From Version | To Version |
| 4.5.6.0 | 5.1.4.200 |
| 4.5.5.0 | 5.1.4.200 |
| 4.5.4.0 | 5.1.4.200 |
| 4.5.3.0 | 5.0.10.0 |
| 4.5.2.0 | 5.0.10.0 |
| 4.5.1.0 | 5.0.10.0 |
| 4.5.0.0 | 5.0.10.0 |
| 4.4.1.0 | 5.0.10.0 |
| 4.4.0.0 | 5.0.10.0 |
| 4.3.1.0 | 5.0.10.0 |
| 4.3.0.0 | 5.0.10.0 |
| 4.2.1.0 | 5.0.10.0 |
| 4.2.0.0 | 5.0.10.0 |
| 4.1.0.0 | 5.0.10.0 |

Table 3: From Versions 3.x

| From 3.x Versions | |
|-------------------|------------|
| From Version | To Version |
| 3.9.3.0 | 5.0.10.0 |
| 3.9.2.0 | 5.0.10.0 |

| From 3.x Versions | |
|-------------------|------------|
| From Version | To Version |
| 3.9.1.0 | 5.0.10.0 |
| 3.9.0.0 | 5.0.10.0 |
| 3.8.1.0 | 5.0.10.0 |
| 3.8.0.0 | 5.0.10.0 |
| 3.7.0.0 | 5.0.10.0 |
| 3.6.2.0 | 5.0.10.0 |
| 3.6.1.0 | 5.0.10.0 |
| 3.6.0.0 | 5.0.10.0 |
| 3.5.4.0 | 5.0.10.0 |
| 3.5.3.0 | 5.0.10.0 |
| 3.5.2.0 | 5.0.10.0 |
| 3.5.0.0 | 5.0.10.0 |
| 3.4.1.0 | 5.0.10.0 |
| 3.4.0.0 | 5.0.10.0 |
| 3.3.0.0 | 5.0.10.0 |
| 3.2.1.0 | 5.0.10.0 |
| 3.1.0.0 | 5.0.10.0 |

Table 4: From Versions 2.x

| From 2.2.x, 2.3.x Versions | | From 2.1.x Versions | | From 2.0.x Versions | |
|----------------------------|------------|---------------------|------------|---------------------|------------|
| From Version | To Version | From Version | To Version | From Version | To Version |
| 2.3.18.0 | 4.5.6.0 | 2.1.9.1 | 2.3.18.0 | 2.0.8.0 | 2.1.9.1 |
| 2.3.16.0 | 4.5.6.0 | 2.1.9.0 | 2.3.18.0 | 2.0.7.0 | 2.1.9.1 |
| 2.3.15.0 | 4.5.6.0 | 2.1.8.0 | 2.3.18.0 | 2.0.6.* | 2.1.9.1 |
| 2.3.14.0 | 4.5.6.0 | 2.1.7.0 | 2.2.9.0 | 2.0.5.0 | 2.1.9.1 |
| 2.3.12.* | 4.5.6.0 | 2.1.6.0 | 2.2.9.0 | 2.0.4.0 | 2.1.9.1 |
| 2.3.9.* | 4.5.6.0 | 2.1.5.0 | 2.2.9.0 | | |
| 2.3.8.0 | 4.5.6.0 | 2.1.4.0 | 2.2.9.0 | | |
| 2.3.7.0 | 4.5.6.0 | 2.1.3.0 | 2.2.9.0 | | |
| 2.3.6.0 | 4.5.6.0 | 2.1.2.0 | 2.2.9.0 | | |
| 2.3.4.0 | 4.5.6.0 | 2.1.1.0 | 2.1.9.1 | | |
| 2.3.3.0 | 4.5.6.0 | 2.1.0.0 | 2.1.9.1 | | |
| 2.3.2.1 | 4.5.6.0 | | | | |
| 2.3.2.0 | 4.5.6.0 | | | | |

| From 2.2.x, 2.3.x Versions | | From 2.1.x Versions | | From 2.0.x Versions | |
|----------------------------|------------|---------------------|------------|---------------------|------------|
| From Version | To Version | From Version | To Version | From Version | To Version |
| 2.3.1.0 | 4.5.6.0 | | | | |
| 2.2.11.0 | 3.9.3.0 | | | | |
| 2.2.10.0 | 3.9.3.0 | | | | |
| 2.2.9.0 | 3.9.3.0 | | | | |
| 2.2.7.* | 3.9.3.0 | | | | |
| 2.2.6.0 | 3.9.3.0 | | | | |
| 2.2.5.* | 3.9.3.0 | | | | |
| 2.2.3.* | 2.2.11.0 | | | | |
| 2.2.2.0 | 2.2.11.0 | | | | |
| 2.2.1.0 | 2.2.11.0 | | | | |
| 2.2.0.0 | 2.2.11.0 | | | | |

Table 5: From Versions 1.x

| From 1.4.x Versions | | From 1.3, 1.2, 1.1 Versions | | From 1.0.x Versions | |
|---------------------|------------|-----------------------------|------------|---------------------|-----------------|
| From Version | To Version | From Version | To Version | From Version | To Version |
| 1.4.12.0 | 2.1.9.1 | 1.3.** | 1.4.6.0 | 1.0.7.* | Contact Support |
| 1.4.11.0 | 2.1.9.1 | 1.2.** | 1.4.6.0 | 1.0.6.* | Contact Support |
| 1.4.10.0 | 2.1.9.1 | 1.1.** | 1.2.2.0 | | |
| 1.4.9.0 | 2.1.9.1 | | | | |
| 1.4.8.0 | 2.1.9.1 | | | | |
| 1.4.7.0 | 2.1.9.1 | | | | |
| 1.4.** | 1.4.12.0 | | | | |

Known Critical Issues

| Known Critical Issues in NimbleOS version 5.3.0.0 | | | | | |
|---|--------------|--|---|---|--|
| ID | Component | Title | Description | Workaround | |
| AS-77607 | Data Service | Removing member array from multi-array group may cause IO disruption to scaled vVol environments | Scaled vVol environments with 500 vVol VDI VMs or more than 5000 Nimble vVol volumes may experience IO disruption when removing a member array from group. Symptom of problem would appear as vVol datastores being (inaccessible). Virtual Machine status would also appear as (inaccessible). | When planning to remove a member array from group, schedule a planned maintenance window and place all ESX hosts into maintenance mode to minimize impact to availability. ESX typically resumes connection to vVol datastores, and reconnects to VMs, after a period of 15-30 minutes automatically without a manual intervention. | |

| Known Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|---|--------------|---|---|--|
| ID | Component | Title | Description | Workaround |
| AS-101976 | Data Service | Volume move may result in latency if Nimble Connection Manager is not installed | Volume moves transfer data from one Nimble array to another. During this move, if the host sends I/O to the incorrect array, the I/O needs to be forwarded to the correct array owning the data. This results in higher than usual I/O latency. To avoid this issue in VMware environments, the Nimble Connection Manager for VMware needs to be installed on all hosts accessing the volume. | Not applicable |
| AS-90668 | Data Service | Data Service restarts when detecting metadata inconsistency | When the Data Service detects a metadata inconsistency, the service may restart repeatedly and hosts could experience unexpected application behavior. | Contact HPE Nimble Storage Support |
| AS-96371 | Data Service | Data service may restart repeatedly due to faulty SSD | If the array encounters a faulty SSD with multiple bad segments, the Log-structured File System (LFS) attempts to resolve the errors. If the array is under high load and LFS has a high number of block errors to resolve it may exhaust the IORW buffers causing the Data Service to restart repeatedly. | Contact HPE Nimble Storage Support to identify and replace the faulty SSD. |
| AS-106093 | Data Service | Data service may restart due to a race condition | While committing internal transactions, Data Service may hit a rare race condition. To recover from this Data Service might restart | Not applicable |
| AS-94834 | Data Service | Data Service may unexpectedly restart | A disruption in network connections can cause Data Service to restart unexpectedly. | Not applicable |
| AS-108094 | Data Service | Replication interruptions and Data Service restarts may occur due to network errors | In environments with frequent network checksum errors, interruptions of Replication and Data Service restarts may occur. | Review network devices to identify and reduce network checksum errors. |
| AS-105607 | Data Service | Snapshot replication of deduplication-enabled volumes may lead to File System restart | During snapshot replication of a dedupe-enabled volume, the downstream array file system may restart due to an out-of-memory condition. | Not applicable |

| Known Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|---|-----------|---|--|--|
| ID | Component | Title | Description | Workaround |
| AS-94961 | Platform | Performance affecting firmware defect in a subset of 6TB drives. | HPE Nimble Storage has identified a rare firmware defect in a subset of drives which can, under certain write intensive workloads, cause the array to under-perform. | Contact HPE Nimble Storage Support. |
| AS-86099 | Platform | Data service may restart during when file operation timeout is exceeded | During internal file operations, processes may be waiting for a lock to be released. If the wait time exceeds 30 seconds, a service health check may restart the Data service to recover. | Not applicable |
| AS-104517 | Platform | Data Service may restart due to health check failure | The Data Service on the array may restart when timeout for internal communication between array controllers has been exceeded. The service restarts to restore the communication. | Not applicable |
| AS-108793 | Platform | Data Service may restart if the array has multiple bad drives | In rare instances, the Data Service may restart if the array has multiple bad drives which make IO handling very slow. | Contact HPE Nimble Storage Support to review disks for replacement. |
| AS-96053 | Platform | NDER process may lead to host reconnects | Nimble Drive Error Recovery (NDER) is activated for drives failing I/O in an attempt to recover the drive. In rare instances, the process may surpass iSCSI host timeout values, causing host I/O inaccessibility. | Not applicable |
| AS-107489 | Platform | Limited thermal monitoring policy for PCIe components | The current thermal policy has a limitation where individual PCIe components are not monitored on card-by-card basis. As a result, the high level temperature policy that is current implemented, is sometimes incapable of regulating the temperatures of individual cards. | Not applicable |
| AS-51053 | Platform | Array remains in solo/stale after controller reset. | In rare instances the interconnect between controllers may be down following a controller reset. | Contact HPE Nimble Support for assistance in restoring array to Active/Standby status. |

| Known Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|---|-------------------|---|--|---|
| ID | Component | Title | Description | Workaround |
| AS-93456 | Platform | Incorrect software update error message /tmp is out of space. | During software update from NimbleOS 3.x versions to later version, a failure to extract the software update package may incorrectly return the error message /tmp is out of space. even if there is still space in the directory. | Verify that the software update package has been downloaded to the system and then attempt the software update again using the software --resume_update command from the CLI. |
| AS-110030 | System Management | Group Management service may restart due to race condition | Group Management service may restart due to a race condition between threads when a volume access control list is removed. The service restarts to recover from the condition. | Not applicable |
| AS-94737 | System Management | No Automatic Failover in the event the host loses all FC connectivity to an array | An Automatic Failover (AFO) of the Group Management Services will not be initiated if all Fibre Channel (FC) interfaces on the Group Leader array fail on both controllers. | A Manual Group Leader Failover will be required to restore Fibre Channel connectivity to the hosts. |
| AS-89701 | System Management | Automatic Switchover Service restarts due to thread limitations | The Automatic Switchover Service internally creates and closes threads each time during Automatic Failover (AFO) quorum setup and tear down. This may cause the service to eventually crash after reaching the maximum thread limit. The system recovers automatically when the Automatic Switchover Service restarts. | Not applicable |
| AS-107367 | System Management | Array resetup fails due to previous complications with Array Switchover Service | The Automatic Switchover Service may restart unexpectedly when the witness is removed or Automatic Switchover is disabled on an array group. In rare instances, this may lead to an issue with a database entry within the array is not cleared successfully. If this array resetup is attempted on this array, the operation will fail when the Array Management Service encounters this stale entry. | Please contact HPE Nimble Storage Support |

| Known Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|---|-------------------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-65615 | System Management | Group Management Service must be restarted to unlock additional volume limits after controller upgrade | When performing a controller upgrade to a high-end model, the object limits will still show the lower limits if the Group Management Service is not restarted. | A failover can be initiated in order to restart the Group Management Service. You may also contact HPE Nimble Storage Support to restart the service manually. |
| AS-106276 | System Management | Array group remains out-of-sync following network recovery. | When there are network communication issues between the Group Leader and Backup Group Leader, the system goes into an out-of-sync condition. In rare circumstances, even after network connection is restored, the array group may still remain out-of-sync. | Please contact HPE Nimble Storage Support. |
| AS-61614 | System Management | Group Management service may restart during array shutdown | The Group Management service may restart during an array shutdown while processing REST request. No user operations are impacted because the array is already in the middle of a shutdown. The shutdown proceeds normally. | Not applicable |
| AS-94113 | System Management | Group Management service may restart due to duplicate snapshot collection name | If an attempt is made to create a snapshot collection with same the name of an existing snapshot collection, an error indicating object already exists will be reported and the Group Management service may restart. The service will recover upon restart. | Use unique name when creating snapshot collection. |
| AS-100254 | System Management | Group Management Service restarts under heavy load | A system management process can restart when the system is under heavy load. The system recovers automatically. The Data service is not affected. | Not applicable |
| AS-98124 | System Management | Array Management service restart during service shutdown | The Array Management service may restart due to a race condition encountered during service shutdown. The service will recover after the restart. | Not applicable |
| AS-87736 | System Management | Software precheck failures return generic error message | If a software update precheck fails, in some cases it will return only the failure status without providing additional information about the cause of the failure. | Contact HPE Nimble Storage Support for assistance in determining the cause of the failure. |

| Known Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|---|-------------------|---|--|---|
| ID | Component | Title | Description | Workaround |
| AS-98194 | System Management | Group Management Services unavailable temporarily when performing volume restore | Group Management may become unavailable temporarily when a large number of restores are performed on a volume. | Not applicable |
| AS-95169 | System Management | Graceful shutdown takes longer than expected | In rare occurrences, a customer-initiated reboot may cause a kernel reboot on the active controller. This will cause a longer reboot cycle. | Not applicable |
| AS-105612 | System Management | Controller failover may be triggered by delayed shutdown of an internal array process | The arrays Data Service may restart under certain conditions in order to maintain data availability. In some cases, the restart of this process takes longer than expected, resulting in an unexpected timeout and a controller failover. The controller experiencing the timeout will reboot and become the standby controller. This event is non-disruptive. | There is no workaround for this issue. Controller failover will maintain data availability. |
| AS-84499 | System Management | Controller may unexpectedly restart due to high memory utilization | Controller may reboot unexpectedly if there is high memory utilization for the java and/or jetty processes on the array. | Not applicable |
| AS-92379 | System Management | Unable to Filter volumes using Synchronous Replication | There is currently no way to filter volumes using Synchronous Replication within the array GUI. | Use volume collections to check syncRep volumes, or use other filters to meet the needs |
| AS-97968 | System Management | Page footer in GUI may fail to update after bulk update operation | After performing a bulk update operation from the Manage > Data Storage > Volumes > volumename > Data Protection tab in the GUI, the page footer may not update and previous button may be unavailable. | Refresh the page to restore button functionality. |

Resolved Critical Issues

| Resolved Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|--|--------------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-105741 | Data Service | Data Service may restart during fingerprint lookup | When a blocks fingerprint is requested and the corresponding domains Fingerprint Index is being merged in parallel, Data Service may restart due to a race between Fingerprint Index lookup and merge operations. | Not applicable |
| AS-105458 | Data Service | Data Service may restart unexpectedly due to health check failure. | Under certain conditions, the Data service may restart during array internal index processing within a short time span. Transactions during the processing may take too long to complete within the defined time span, which causes the service to restart. | Not applicable |
| AS-106021 | Data Service | Index verification fails if a 16 TiB volume is completely unmapped causing Data Service to go down | In rare cases during index creation, when a 16 TiB volume is fully unmapped, the resulting index structure fails verification and brings down the Data Service leading to an outage. | Contact HPE Nimble Storage Support. |
| AS-110870 | Data Service | File system service may restart unexpectedly during data block allocation status API execution | During data block allocation status application programming interface (API) execution the file system service may restart. This happens when data block allocation start offset is chunk-size-bytes aligned but data block allocation end offset is not aligned with chunk-size-bytes. | Contact HPE Nimble Storage Support |
| AS-96300 | Data Service | Data Service may restart due to volume manager health check failure | Generation delete operations and NVRAM to disk data flush operation can cause Data Service to restart due to health check failure, as it can hold checkpoint for a long time. | Not applicable, on restart Data Service would behave normally. |

| Resolved Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|--|
| ID | Component | Title | Description | Workaround |
| AS-104732 | Host Integration | VASA Sessions created and cached even though Group Management service is unavailable. | When the Group Management service restarts for unrelated reasons, the host sends set context (creates a new session) because its VASA sessions have been invalidated with the Group Management service going down. However, when the new session call is made, we create a new session, then realize Group Management is down, and send back an error to the host. Unfortunately, the new session is created (object in memory), but is not truly a valid session. These sessions could accumulate to a large number over time, which may cause an out of memory condition for Jetty. Once the service runs out of memory, it can also restart unexpectedly. | Not applicable |
| AS-106391 | Platform | Services may fail to start on a replacement controller | Under specific circumstances, services may fail to start on a replacement controller. | Please contact HPE Nimble Storage Support if you encounter this issue. |
| AS-103129 | Platform | Data Service may restart while committing large internal transactions | In rare cases while committing large internal transactions, the process may timeout. As a result, the Data Service may restart to recover the condition. | Contact HPE Nimble Storage Support if there are multiple restarts to work around the issue. |
| AS-104924 | dHCI | Plugin: Cannot add 4 or more servers in the dHCI deployment | Currently, if customer plans to add 4 or more Proliant servers in their dHCI deployment via the plugin, the operation fails. | In order to add more 4 or more Proliant Servers, the workaround is to add up to 3 servers at a time from the plugin. |
| AS-103976 | System Management | Group management service restart during shelf activation | Group management service may restart during shelf activation on Backup Group Leader (BGL) array. This occurs if a user tries to activate a shelf for BGL array which is not associated with any pool. The Group management service restarts because of an empty pool. | Create a pool on the Backup Group Leader array, then activate the shelf. |

| Resolved Critical Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|---|--|
| ID | Component | Title | Description | Workaround |
| AS-105878 | System Management | Unsafe Automatic Switchover (ASO) in the event of connectivity loss. | In rare circumstances, especially in deployments without network redundancy, if the group leader array loses connectivity to both the backup group leader and the witness, this will result in Automatic Switchover (ASO). There is a small risk that not all writes will have been mirrored to the partner array making ASO an unsafe operation. | Disable Automatic Switchover using the steps below: 1. Login to the array GUI 2. Select Administration > Availability 3. Uncheck the Enable checkbox for Automatic Switchover 4. Click Save Use Manual Failover for Synchronous Replication Volumes in place of ASO. |
| AS-92465 | System Management | Intermittent login failures due to Active directory lookups failures | There is a possibility of sporadic, transient, active directory authentication failures. In these cases, the system will recover on its own, requiring no user interaction. | Not applicable |

Resolved Issues

| Resolved Issues in NimbleOS version 5.3.0.0 | | | | |
|---|--------------|--|---|--|
| ID | Component | Title | Description | Workaround |
| AS-106778 | Data Service | Data Service restart during deduplication domain deletion and software update | When a dedupe domain deletion is in progress during software update, internal metadata needs to be cleaned up in order for the deletion to make progress. During this process, the Data Service may restart in an attempt to automatically recover. | Not applicable |
| AS-78101 | Data Service | Services that rely on TCP connections might be interrupted if IPs are configured but link status is down | This scenario is rare to happen as timeout values for most of services would prevent this scenario to be triggered. Eventually an IP with link up status will be picked up within the timeout period of the service. | Remove the IP with link status down from configuration |
| AS-77921 | Data Service | Data Service may restart when gathering garbage collection telemetry | Data Service may restart when gathering garbage collection telemetry if cache subsystem is not fully initialized. | None |

| Resolved Issues in NimbleOS version 5.3.0.0 | | | | |
|---|------------------|---|---|---|
| ID | Component | Title | Description | Workaround |
| AS-103070 | Host Integration | vCenter plugin not getting deployed properly because of version mismatch | When array software is updated using the GUI, the new vCenter plugin is not getting installed properly this is because there is a mismatch of plugin version specified in the plugin-package.xml verses the actual plugin version. | Contact HPE Nimble Storage Support for assistance with plugin registration. |
| AS-90455 | Platform | IPMI software may not handle command exchange correctly with BMC leading to unexpected reboots of AFx/HFx controllers | In rare cases, out of order commands being sent to the Baseboard Management Controller (BMC) may return out of order responses that are not handled in the correct order by Intelligent Platform Management Interface (IPMI) software. In this instance, the IPMI message queue loses track of message order. The IPMI message queue not being able to return IPMI Watchdog messages to the watchdog thread causes the watchdog thread to timeout leading to an automatic reboot. While this BMC Watchdog timeout issue is specific to the AFx/HFx systems, this is not a hardware issue. Therefore, hardware replacement is unnecessary. | After the controller reboots, BMC firmware is restarted and is functional again automatically. |
| AS-91379 | Platform | Interfaces on Quad Port 10GbE BaseT and 10GbE SFP cards using VLAN tagging may become unresponsive | When a subnet on the array is configured to use VLAN tagging, Quad Port interfaces in that subnet may become unresponsive. As a result, hosts connected to an unresponsive interface, on that VLAN, may experience communication interruption. Over time, this condition may occur on every interface within that subnet, causing instability with that subnet. | Before an interface is in state, gracefully remove VLAN tagging from the Quad Port interface. After an interface is in state, perform controller failover to recover the condition. However, with VLAN tagging still enabled, subsequent events can occur. Therefore, consider the aforementioned workaround to minimize chance of a re-occurrence. |

| Resolved Issues in NimbleOS version 5.3.0.0 | | | | |
|---|-------------------|--|--|---|
| ID | Component | Title | Description | Workaround |
| AS-94761 | SAN | File System service may restart when an invalid write request is received | This scenario is rare to happen but presently not handled gracefully leading to a File System restart for recovery. The issue is triggered when a write request with valid length is received in SCSI Command Descriptor Block but invalid (zero) value in Data-Out Buffer. | Not Applicable |
| AS-101273 | System Management | Event Management service restart when two folders in different pools have the same name | When an alert is raised on one of two folders having the same name but are located in two separate pools, the Event Management service may restart repeatedly. | Rename the folders to be unique. If the Event Management service continues to restart or remains unavailable, contact HPE Nimble Storage Support. |
| AS-96143 | System Management | Group management service may restart due to assertion failure | NimbleOS uses a defined state machine for the replication workflow. At the end of the execution of each step defined in the state machine, it moves to the next step. If it leads to any unexpected step throughout the workflow then it will lead to assertion failure which results in Group Management service restart. | Not applicable |
| AS-105064 | System Management | Group management service may restart unexpectedly | Internal workflow processing related to the replication partner object may cause the process to deadlock, resulting in Group management service restart. | Not applicable |
| AS-105432 | System Management | Deletion of a volume is not completed due to the presence of stale ACLs associated with it | In certain scenarios, a volume deletion will not complete due to the existence of a stale ACL associated with it. This stale ACL is associated with a snapshot of the volume that was previously deleted. These volumes will not show up in the CLI/GUI as they are in a hidden state. | Contact HPE Nimble Storage Support to identify ACLs in forced-delete/create-retry state, verify these ACLs are associated with snapshots that no longer exist, and delete these ACLs from the internal NimbleOS database. |
| AS-106168 | System Internals | Data Service may restart due to health check failure | In rare cases the Data Service may restart due to a system health check failure caused by locked index processing threads. The service will recover after the restart. | Not applicable |

Known Issues

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|--------------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-108086 | Data Service | Data Service restart while converting internal data-structures in NimbleOS | After updating the array to NimbleOS 4.x.x.x or later, a process runs to convert an on-disk data-structure specific to NimbleOS by initiating an operation that walks the existing tree of this data-structure. In some cases, this operation keeps running even after the conversion has completed successfully, leading to health check timeout. This results in a Data Service restart. | Contact HPE Nimble Storage Support to increase the health check timeout on the array |
| AS-81863 | Data Service | Data Service may restart unexpectedly when RAID is in degraded mode | When RAID is degraded, IO needs to be reconstructed by reading from multiple disks, and an internal buffer may exhaust its allocated resources. In rare cases when multiple disks are degraded, the Data Service may restart unexpectedly. | If RAID is degraded for an extended period and Data Service restarts occur, contact HPE Nimble Storage to assess adjusting allocated buffer resources. |
| AS-62942 | Data Service | Data Service may restart unexpectedly for CS2XX and CS3XX arrays when under a heavy load | On CS2XX or CS3XX, the array may run out of data pages when under a heavy load. This will result in an unexpected Data Service restart. | Reducing the load on the array is the only known workaround for this issue. |
| AS-102001 | Data Service | Data Service may restart unexpectedly due to internal database communication | In rare instances, the Data Service may restart when internal database communication between services is not available. The service restarts to restore the communication between services. | Not applicable |
| AS-111378 | Data Service | Data Service may restart due to health check failure | A lost wake-up to an operation may cause the operation blocking the checkpoint which results in Data Service restart. | Not applicable |
| AS-110889 | Data Service | Data Service may restart when array runs out of scratch pages | This can occur when the memory footprint used by internal indexes exhaust the default allocation of scratch pages in the system after several volumes have been deleted. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|--------------|--|---|---|
| ID | Component | Title | Description | Workaround |
| AS-102881 | Data Service | Data Service Restart due to a race condition during metadata sync | During NimbleOS metadata sync, in rare instances, the data service may restart unexpectedly. The metadata sync operation itself wont be affected and the restart will reset the race condition; the data service will stabilize after the restart. | Data service will be available after restart. |
| AS-85848 | Data Service | Data Service may restart unexpectedly with health check failure when internal index data structures take too long to merge | As part of the file system checkpoint process, internal indexes are required to merge to disk. The health check process uses a heartbeat mechanism to determine if the merge is making progress. Under some circumstances, the process performing the merge of a batch of updates to the index fails to heartbeat within a deadline. When that happens, the Data Service may restart with a health check failure. | Not applicable |
| AS-87108 | Data Service | Data Service may restart unexpectedly with VM health check failure due to lock contention. | Under certain conditions, the Data Service may restart when a large number of internal indexes merge within a short timespan and when some of these indexes need to merge multiple times in the same checkpoint. In this case, there can be lock contention between operations responsible for picking what to merge and operations processing the merge, leading to health check timeout. | Not applicable |
| AS-109266 | Data Service | Data Service may restart during metadata invalidation | The Data Service may restart unexpectedly on an array group with Storage Class Memory and Synchronous Replication configured, during an automatic failover (AFO) event. This occurs when invalidating index data in memory exceeds the timeout of 30 seconds. | The array group will recover itself following the Data Service restart. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|--------------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-111235 | Data Service | Data Service restarts when one of the processes incorrectly sets its return status. | In one of the code paths for the process that handles fingerprint mapping, the merge status is not properly set. This may lead to an unexpected Data Service restart. | Not applicable |
| AS-96703 | Data Service | Data Service may restart due to volume manager health check failure during generation deletion | Generation delete loads a large number of ondisk metadata blocks which may prevent block index operation checkpoint from finishing. This causes the volume manager health check to fail which results in Data Service restart. | Not Applicable, the Data Service will resume normal operation after restart. |
| AS-96779 | Data Service | Data Service may restart unexpectedly due to race condition | When a read op finds partial data in-core, it issues a media read to get remaining data. By the time, media read returns, the in-core data is synced, and tree is reopened for deletion. The read does not expect tree to be in delete state and causes the Data Service to restart. | Not applicable |
| AS-94473 | Data Service | Data Service may restart when running out of buffers | When flash cache Garbage Collection copies forward live data of a fragmented segment, it could consume more buffers than provisioned and cause the Data Service to restart to recover. | Contact Nimble Storage Support. |
| AS-108519 | Data Service | File system restart to recover from stalled replication | Due to issues in communicating with the partner array during replication, there are few cases where the operation is not able to make progress. As a result, the file system may restart to correct this condition. | Not applicable. |
| AS-111347 | Data Service | Data service may restart due to a race condition | While committing internal transactions, Data Service may hit a rare race condition. To recover from this Data Service may restart. | Not applicable |
| AS-81739 | Data Service | Data Service restart due to slow disk IO or disk IO failure | If IO to disk are slow or fail, the Data Service may restart to try to recover the condition. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|--------------|--|---|--|
| ID | Component | Title | Description | Workaround |
| AS-106924 | Data Service | Data Service may restart due to network errors | In rare cases, Data Service may restart during snapshot replication due to the failure of checksum algorithms to detect all network errors. | Not applicable |
| AS-92170 | Data Service | Data Service can restart unexpectedly during shutdown process | Due to a race condition, the Data Service may restart during a graceful shutdown causing unexpected Data Services restart messages to be generated. This should not cause any I/O impact because the Data Service is already in the process of shutting down. | Not applicable |
| AS-97038 | Data Service | Disabling encryption may cause Synchronous Replicated volumes to remain out of sync | Disabling encryption might cause NVRAM data to fill up with data waiting for master passphrase, which can cause Synchronous Replicated volumes to go out of sync. | Enabling encryption will resolve the issue |
| AS-105714 | Data Service | Data Service may restart if network issue is encountered between Group Leader and Backup Group Leader arrays | A network issue between upstream and downstream could abruptly stop operations running downstream causing them to exit prematurely, resulting in Data Service restart. | Contact HPE Nimble Storage Support |
| AS-94545 | Data Service | Very rare race between Vol claim (with all snapshots marked for deletion) and space recalculation on replica downstream volume | The service may restart when removing the downstream replica using the steps below. 1. Deletion of all snapshots for the replica volume. 2. Claim the replica volume 3. Delete the replica volume | Remove downstream replica using the correct steps ordered below. 1. Claim the replica volume. 2. Delete all of the snapshots for the replica volume. 3. Delete the replica volume. |
| AS-98217 | Data Service | Data service may restart during array shutdown | Volume manager does not reset internal callbacks during the shutdown phase causing the Data service to restart. | The array will continue to shutdown after the Data service restart. |
| AS-79265 | Data Service | Data Service may restart on downstream array due to race condition during volume deletion | The volume manager will decrement bin and child volume folder entries during volume deletion. In rare occurrences, another thread may decrement the child volume entry and not the bin entry, causing the Data service to restart. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|------------------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-86720 | Data Service | Unassigning and reassigning array to a pool within 5 minutes will fail | Assigning an array to a pool immediately after unassigning it from the same pool will fail with the following error - Failed to assign arrays to the pool: A service is not running or is not reachable | Retry operation after a few minutes to reassign array to pool. |
| AS-111454 | Data Service | Aggressive Volume creation and deletion may result in Data Service restart | Aggressive volume creation and deletion may cause one of the tree data structures to hit maximum children it can have. This limit can occur on volume creation because volume deletion destroys the children asynchronously. The Data Service will restart when this limit has been reached. | Not applicable |
| AS-108647 | Host Integration | Cimserver stops when memory limit is exceeded. | In some instances the cimserver service will exceed its memory limit and no longer be able to start. | Contact HPE Nimble Storage Support to increase the memory limit for the service. |
| AS-50033 | Platform | Log partition may fill up | Log files generated by processes in the Nimble Operating systems are rotated and archived using certain parameters. Under some conditions, if processes log more than normal, this can result in the log partitions being filled up, leading to other issues with critical processes. | Contact Nimble Storage Support to identify the reason for the partition filling up and to clean up the log partition. |
| AS-53621 | Platform | Both power supplies showing up as missing | Under certain circumstances, internal commands may cause the integrated circuit to hang which causes both power supplies of the head shelf to appear as missing. The problem does not impact the power supply's ability to deliver power to the array. | <ol style="list-style-type: none"> 1. Reseat one of the power supplies to clear the hang condition. 2. Or update NimbleOS to version 3.6.0.0 or later to minimize the number of internal commands that can trigger this issue. |
| AS-100088 | Platform | Controller does not power on following a power cycle. | In rare incidents, controllers do not power on following power cycle. | Please contact HPE Nimble Storage Support |
| AS-67242 | Platform | Incorrect shelf cabling can cause a series of unexpected Data Service restarts | If a shelf is cabled incorrectly and added to an array, the raw capacity calculation could fail leading to multiple restarts of the Data Service. | Referencing the corresponding Expansion Shelf Quick Start Guide, correct the hardware cabling issue. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-----------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-33725 | Platform | Unexpected controller takeover due to incorrect state of the SAS HBA | When the SAS HBA detects faulty states, to recover, the array needs to reset the SAS HBA's firmware. The SAS HBA firmware reset can block disk I/Os significantly longer than our High Availability monitoring timeouts allow. Instead, a controller reboot is triggered immediately if this state is detected, resulting in an unexpected takeover event. | The controller reboot should restore SAS HBA to normal state. HPE Nimble Support may contact customer to collect additional diagnostics if required. |
| AS-95294 | Platform | Enclosure Management service may unexpectedly restart due to drive status race condition | The Enclosure Management service may unexpectedly restart due to a race condition which is resulting from drive status which has already been marked removed by the Data Service. The restart of the service clears the condition. | Not applicable |
| AS-90850 | Platform | Data Service may unexpectedly restart | The Data Service may restart unexpectedly when it detects an internal check error. The restart of the service clears the condition. | Not applicable |
| AS-103802 | Platform | Data Service restart due to resource allocation failure. | The Data Service may restart due to a transient resource allocation failure. This happens when the service cannot complete a disk IO due to transient memory allocation failure. This does not cause a service outage as Data service continues normally after a restart. | Not applicable |
| AS-99567 | Platform | Data Service may restart if a controller is low on memory | Data Service may restart in the rare case when a controller is low on memory. | The restart of the service will clear the low memory condition, no further action is needed. |
| AS-110036 | Platform | Data Service may restart if the array has an faulty drive with a high number of medium read errors | Data Service may become latent and restart when the array has a faulty drive with a lot of medium read and IO timeout errors. | Contact HPE Nimble Storage Support to identify the faulty drive for replacement. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-----------|---|--|--|
| ID | Component | Title | Description | Workaround |
| AS-46629 | Platform | Kernel panics while processing host bus adapter errors | On very rare occasions of HBA error handling, interrupts could be turned off long enough for the NVRAM driver to timeout while trying to complete direct memory access operations. The controller recovers after a re-boot is initiated by the kernel. | Not applicable |
| AS-101570 | Platform | Delay with Data Service starting during shelf state change | On rare occasions, the array groups Data Service may fail to initialize if a shelf state change occurs simultaneously. | No workaround is required. The array will recover itself automatically by restarting the Data Service. |
| AS-99428 | Platform | Replacement disk reports foreign for disk state in GUI/CLI | Disk will report foreign for disk state in the GUI Hardware page or in the output of disk --list from the CLI. This typically occurs if diagnostic data may not have been removed after testing. | Add the disk from the CLI using the disk add command and output from disk list: 1. Run disk --list 2. Note the slot number, and shelf location for the disk labeled foreign. 3. Add the disk: disk --add <slot number>; --array <arrayname>; --shelf_location <shelf location>; Note: the --force option may be required Contact HPE Nimble Storage Support if the disk does not move to resynchronizing state after completing the commands. |
| AS-93296 | Platform | Data service may restart if maximum cache exceeded for CS215, CS235, CS300, CS500, CS700 arrays | ES2 and AFS2 expansion shelves contain additional slots for upgrading cache capacity of the array. Older array models have a maximum cache limit that can be handled by the array. If ES2 or AFS2 expansion shelves are added to an array and the cache exceeds the max cache limit for the array type, the data service may restart due to running out of data pages. | Review the array configuration matrix for the array model: https://infosight.hpe.com/InfoSight/media/local/active/34/CSxxx%20Config%20Matrix.pdf Remove any additional cache from the expansion shelf that exceeds the max cache limit based on array model. |
| AS-105053 | Platform | Enclosure Management Service may restart unexpectedly | The Enclosure Management service may unexpectedly restart when it detects an internal check error. The service restart clears the condition. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-----------|--|--|---|
| ID | Component | Title | Description | Workaround |
| AS-91522 | Platform | SSD has reached its endurance limit (wear leveling) but the disk is not marked failed. | In rare cases, an SSD may reach its endurance limit but continues to pass Nimble Drive Error Recovery algorithm. This causes a never-ending process of off-lining and on-lining the drive. This may occur for Intel and Toshiba SSDs. | Not applicable |
| AS-98885 | Platform | Unexpected Group Management Service restart due to receive buffer exhaustion | The Group Management Service may restart unexpectedly due to receive buffer exhaustion on the management network interface. No visible impact has been reported because of this issue. | Not applicable |
| AS-76174 | Platform | VM creation fails in VMWare VVol environment after update to NimbleOS 4.x.x.x | An attempt to create new VMs in a VMWare vVol environment fails after updating to NimbleOS 4.x.x.x. This occurs when a system partition within the array is running out of space. | Please contact HPE Nimble Storage Support to resolve this issue. This issue is resolved in NimbleOS version 5.0.2.0 and later. |
| AS-103315 | dHCI | dHCI update workflow can fail if /var mountpoint on the host is full. | As part of the dHCI update workflow which involves the ESXi server update, some commands need to be run on the server. Running these commands fails with error Error: A general system error occurred: Internal error. if the /var is full. This is a known issue with 6.7 builds running with Emulex driver which has been fixed as part of ESX 6.7 U2. https://docs.vmware.com/en/VMware-vSphere/6.7/rn/vsphere-esxi-67u2-release-notes.html#resolvedissues | Free up /var disk space by deleting unwanted log files (/var/log/EMU/mili/mili2d.log) and rebooting the host. |
| AS-103247 | dHCI | An in-progress dHCI update fails if the group leader fails over | When the dHCI unified update feature is used to update the dHCI stack, the update will fail if a Group Leader Failover occurs during the process. | The only way to resume the update would be to failover to the original group leader array and then resuming the update through the dHCI vCenter plugin. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-----------|--|--|---|
| ID | Component | Title | Description | Workaround |
| AS-97685 | dHCI | Nimble Add ProStack server task may hang if array failover occurs | If an array failover occurs while Nimble Add ProStack server task is running from vCenter, the process may hang and not complete. | After confirming the array has returned to Active/Standby status, stop the hung task and run the Nimble Add ProStack server task again. |
| AS-103769 | dHCI | The update page on the dHCI plugin takes 4 hours to refresh | When an ESXi server is added to dHCI cluster, the update page does not get updated to include the newly added ESXi server. This refresh happens every 4 hours. After the next refresh, the new ESXi servers version will be included and accounted for on the update page. | Not applicable |
| AS-95054 | dHCI | Addition of a server with expired ESXi license fails | When adding a server with an expired ESXi license to the dHCI cluster, through the vCenter plugin, you may see an error saying - Failed to submit a task to add server. | A valid ESXi license must be assigned to the server. |
| AS-101915 | dHCI | dHCI update fails when Admission Control is enabled | For ESXi server update, DRS is used to migrate VMs running on the server. If admission control is enabled on the dHCI cluster, DRS is not able to migrate VMs off a server. | Admission Control should be disabled on a dHCI cluster for the update to proceed. |
| AS-108946 | SAN | Data Service may restart when host does not issue abort to timed out command | The Data service will wait up to 300 seconds for host response to iSCSI commands. If host does not abort commands that take more than 300 seconds to complete, the Data service will restart. | Not applicable |
| AS-100197 | SAN | Data Service restart during shutting down FC service on the standby controller | During a controller reboot, due to resource contention between new Fibre Channel (FC) connection attempts and shutdown of the FC module, the Data Service on the array may restart unexpectedly. | Not applicable. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-----------|--|---|---|
| ID | Component | Title | Description | Workaround |
| AS-107345 | SAN | NimbleOS services may restart unexpectedly due to memory leak in login path | Logins to CHAP authentication enabled volumes could leak a small amount of memory by repeated failed login attempts to offline volumes or stale targets. Over a period of days and weeks, this leak can result in one or more of the NimbleOS processes running out of memory. As a result, NimbleOS services may restart unexpectedly. | Identify the offline volumes or stale targets and initiate a host side cleanup/rescan to avoid repeated login attempts and failure to the offline or stale targets. |
| AS-109412 | SAN | NimbleOS services may restart unexpectedly due to slow leak with CHAP logins | Logins to CHAP authentication enabled volumes could leak a small amount of memory the size of CHAP username. Over a period of days and weeks, this leak can result in one or more of the NimbleOS processes running out of memory. As a result, NimbleOS services may restart unexpectedly. The issue is exacerbated by repeated failed login attempts to offline volumes or stale targets. | Identify the offline volumes or stale targets and initiate a host side cleanup/rescan to avoid repeated login attempts and failure to the offline or stale targets. |
| AS-64790 | SAN | Data Service may restart due to race condition | When the Data Service is in process of shutdown, the service may restart due to race condition when shutdown threads do not wait or abort pending operation threads. | Not applicable |
| AS-98042 | SAN | The Data Service restarts unexpectedly during shutdown | When the active controller is being shutdown, the Data Service runs into an internal error condition that causes the service to restart unexpectedly. Since the process is already being shutdown, there is no impact to user data availability. | Not applicable |
| AS-101325 | SAN | Data Service may restart unexpectedly while removing member array | Under certain conditions, the Data Service on the Group Leader array may restart unexpectedly while removing member array. This is due to a race condition when processing SCSI RTPG (REPORT TARGET PORT GROUPS) commands. The service should stabilize on its own shortly following the restart. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|----------------|
| ID | Component | Title | Description | Workaround |
| AS-89753 | SAN | Service may restart due to race condition | The Data Service or SCSI High Availability Service may restart due to race condition encountered during process shutdown. | Not applicable |
| AS-110367 | System Management | Group Management service may restart due to race condition | A race condition may occur while updating records in the Scale Out database when two threads attempt to update the same record. | Not applicable |
| AS-103766 | System Management | Group Management Service may restart due to race condition | Group Management Service may restart unexpectedly as one thread has taken a ReadWrite lock which has another writer thread, which is waiting for Scale-Out Database (SODB) transaction to be completed. The service restarts due to the SODB transaction exceeding the expected timeout. | Not applicable |
| AS-54302 | System Management | REST API Clone of volume does not support setting ACL at time of clone creation | When created with the REST API, by default the cloned volume inherits the ACL from the parent volume. There is no way to set or change the ACL for the clone at the time of creation through the REST API. Clients should make two REST calls, one to create the clone and then one to change the ACL to the desired value. In the future, we could change the REST API to not set the ACL by default and allow the client to set the ACL at the time of the clone creation to the desired value. This only applies to volumes that have agent_type set to none | Not applicable |
| AS-105929 | System Management | Group Management Service restart due to race condition | Group Management service may restart due to Volume Management thread and API thread for snapshot creation/deletion entered race condition. The service recovers after restart. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|--|
| ID | Component | Title | Description | Workaround |
| AS-105431 | System Management | Alarm IDs in alarm list may appear out of order. | Due to the multi-threaded nature of the Alerts and Alarms Service, it is possible for an event with a later id to be posted prior to an event with an earlier id. Therefore it will have an earlier timestamp even though its ID is higher. This can cause the IDs to appear out of order. However, the alarms in the list are ordered correctly by timestamp. | Not applicable |
| AS-106124 | System Management | Member array alarms are still visible after array is removed from group | Alarms raised by a member array are visible when issuing alarm --list even after the member array is removed from the group. | Run the following command via the array CLI: alarm --delete alarm_id |
| AS-94398 | System Management | Alarm may not clear as expected | Due to a race condition it is possible for an alarm to remain uncleared on the system even when the alarm condition it is reporting is no longer the case. This can happen if the onset alert and recovery alert were generated at close to the same time. | The problematic alarm can be manually deleted through the CLI command: alarm --delete <alarm_id> |
| AS-83604 | System Management | Event service may restart unexpectedly | API calls occur between the Event service and Group Management service when an alarm is cleared. The calls will lookup the alarm ID and onset event in the Scale Out Database (SODB). If the alarm is cleared before the onset event lookup completes and is not found during these calls, the service may restart unexpectedly. | Not applicable |
| AS-50821 | System Management | Alerts and Alarms processing service may restart unexpectedly | Alerts and Alarms processing service may restart unexpectedly when certain operation surpass the designated health check timeout. The process will stabilize following the restart. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|---|--|
| ID | Component | Title | Description | Workaround |
| AS-72902 | System Management | Alerts sent via SMTP may be rejected | When the array sends an alert via SMTP, the message may be rejected by the server indicating 550 5.6.11 SMTPSEND.Bare-LinefeedsAreIllegal. This can occur if the receiving server does not support BDAT command for SMTP chunking | The recipient can create an inbound transport rule to append a disclaimer to the messages from the problematic sender. The disclaimer will append the expected CR-LF combination to the message so that it can be delivered. (This disclaimer may consist of a single character such as a period or a dash.) |
| AS-85608 | System Management | The Event service may restart unexpectedly | The Event service may restart unexpectedly due to a memory access issue. The restart will not impact data connections to the array and the process will recover after the restart. | Not applicable |
| AS-103567 | System Management | Eventd process may restart unexpectedly | The Eventd process may restart unexpectedly due to a rare deadlock condition between its threads. The restart of the process will clear the condition. | Not applicable |
| AS-46024 | System Management | Eventd process may restart intermittently | The Eventd process may infrequently crash, due to a bug in an external library used by the DNS resolver. The restart will not impact data connections to the array and the process will recover after the restart. | Not applicable |
| AS-68651 | System Management | Flood of timeouts causing Event Management service restart | We create multiple threads to deliver emails, but we use a non-threadsafe libcurl call to dispatch them. Therefore, the lock needs to be around libcurl call. If there is a misconfigured DNS or SMTP server, the curl call will timeout. If there are greater than 7 emails waiting to be delivered and all are suffering a timeout, we will starve the health checking for more than the 300-second health check timeout causing the Event Management service to restart. | Correct the DNS or SMTP configuration to a valid address by ensuring that a ping to the defined address succeeds. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|--|
| ID | Component | Title | Description | Workaround |
| AS-82919 | System Management | Group Management Service restarts due to excessively long username | An excessively long username of over 255 characters will exceed a character limit within the array groups auditing framework. This would cause the Group Management service to restart unexpectedly. | The Group Management service will recover after the crash. The workaround would be to no longer use an excessively long username, and reduce the username to under 255 characters. |
| AS-96966 | System Management | Purge inactive encrypted keys appears in audit log. | Every night, when inactive keys for deleted encrypted volume are deleted by the array, it creates an audit log entry with root as the user performing the action. | Not applicable |
| AS-71090 | System Management | No Audit Log entry is created if user does not have the privilege to create user | If a user tries to create a new user account, but the user doesn't have the privilege to do so, the user creation will fail. However, an audit log entry is not created. | Not applicable |
| AS-104640 | System Management | Group Management Service restarts following volume creation timeout | The Group Management Service on the array may restart unexpectedly when a SOAP timeout is encountered after trying to create a volume. This occurs due to a race condition where the array attempts to delete the volume after the creation attempt fails. | Not applicable |
| AS-109127 | System Management | Group Management service may restart when connections to Scale-Out Database exceeds threshold value | Connection to the Scale-Out Database can not be reinitialized when a transaction is in progress. The Group Management Service will restart to restore connection. | Not applicable |
| AS-104185 | System Management | Group Management service may restart when performing Autosupport Validation | When autosupport configuration validation is performed, internal process tracking may abort causing the Group Management Service to restart. | Perform the Autosupport validation process again. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|---|----------------|
| ID | Component | Title | Description | Workaround |
| AS-108146 | System Management | Group Management service may restart due to key value metadata handling | During key value metadata handling, the process may delete primary keys in one table of the Scale-Out Database (SODB) and then try to reference the deleted primary key as foreign key in another table in SODB. This causes foreign key violations in Postgres and eventually leads the Group Management service to restart. | Not applicable |
| AS-105454 | System Management | Group Management service may restart due to snapshot key value metadata handling | During key value metadata handling, the process may delete primary keys in one table of the Scale-Out Database (SODB) and then try to reference the deleted primary key as foreign key in another table in SODB. This causes foreign key violations in Postgres and eventually leads the Group Management service to restart. | Not applicable |
| AS-99300 | System Management | Group Management service may restart during startup | During the startup process for Group Management, key value metadata is loaded from the Scale Out Database (SODB) into memory. The Group Management Service may restart due to exceeding the expected time to complete startup if there is a large amount of metadata to load. | Not applicable |
| AS-99704 | System Management | group --status CLI output shows incorrect Failover Mode during network connectivity issues | If there is a network connectivity issue between the witness and Group Leader array, the group status CLI output will update the Failover Mode from Automatic to Manual until the connection is reestablished. It also displays the Witness Status as N/A as opposed to Unreachable. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|--|
| ID | Component | Title | Description | Workaround |
| AS-94649 | System Management | Peer Persistence Automatic Switchover (ASO) is disabled during software update | During software update of array group with Peer Persistence configuration, Automatic Switchover (ASO) is disabled. If an array goes down (both controllers down) during the software update process, due to a power failure or other unexpected event, hosts could lose access to data until the failed array recovers, or a manual switchover of the affected volumes is done. | Not applicable |
| AS-98953 | System Management | Array Management Service restarts during Backup Group Leader discovery | The arrays database system may become unavailable for a limited time when there is a failure in setting up the Backup Group Leader. When attempting to discover a new Backup Group Leader, the Array Management Service may restart due to a race condition. | This restart is non-disruptive to the data on the array, and the Array Management recovers after the restart occurs. |
| AS-94683 | System Management | Network isolation of the Group Leader and Backup Group Leader array may lead to Automatic Switchover service restarts | In Automatic Switchover environments, in rare instances, network isolation of the Group Leader and Backup Group Leader may cause the service that handles the automatic switchovers to restart unexpectedly. | No workaround is needed. The service recovers on its own. |
| AS-101342 | System Management | Group limits command lists internal identifiers | The array group CLI command with limits option (group --list-limits) displays numeric internal identifiers as part of the information listed for the volume information. These numeric identifiers are used by the array only and can be ignored. | Not applicable |
| AS-99702 | System Management | Backup Group Leader is not assigned due to power outage | Following a power outage, it is possible that the Backup Group Leader is not assigned to the group. This may occur if the SODB database does not start due to an SSH key issue. | Please contact HPE Nimble Support. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|--|
| ID | Component | Title | Description | Workaround |
| AS-104812 | System Management | Array Management service restarts due to memory allocation issue | The Array Management service may restart unexpectedly due to an a memory allocation failure when attempting to synchronize configuration with a member array. The restart of the Array Management service clears the situation. | Not applicable |
| AS-93469 | System Management | Group Management service may restart while collecting member array statistics | The Group Management service may restart while collecting statistics from member array. This can occur when the request from the group leader to member array exceeds timeout, causing the service to restart to recover. | Not applicable |
| AS-98297 | System Management | Array Management service restart due to memory allocation exception | There is a memory leak in the stat (statistics) component of the Array Management service. Magnitude of the memory leak is proportional to the frequency of stats query. | If array is queried for statistics, then its frequency should be reduced. This is specially relevant for the setup where script or some monitoring tool does stats query frequently in a loop. |
| AS-87749 | System Management | Max limit of 120 nics in netconfig alarm does not get cleared | When alarm for number of nics in array net config reaching 120 is triggered, it doesnt get cleared even when the number of nics goes down. | The workaround would be to delete the alarm using the alarm --delete CLI. |
| AS-73595 | System Management | A lossy network causes Group Management service restart | Service threads within the group management API handler are restricted to run for no more than 5 minutes. In a very lossy network, TCP throughput can be throttled to almost nothing due to retransmission timeout back-off. In one observed case, where the API response was about 190 KB, the connection managed to send only about 90 KB before the 5 minute timeout occurred and caused the Group Management service restart. | The only workaround is to improve the quality of the network connection so that the loss of TCP packets is minimized. |
| AS-103275 | System Management | Event Service may restart due to invalid database entry | The Event service may restart unexpectedly if an invalid network interface entry is found in the Scale Out Database (SODB). | Contact HPE Nimble Storage Support |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|---|
| ID | Component | Title | Description | Workaround |
| AS-74242 | System Management | Force deletion of user defined performance policy should not be supported | There is a --force switch available when deleting a performance policy via the HPE Nimble Storage Array CLI. This --force switch does not work and will fail with the following: ERROR: Failed to delete performance policy. Resource busy. The --force command is not supported since the specified performance policy should not be removed without first checking its volume or folder associations. | Not applicable |
| AS-101832 | System Management | Volume Migration estimates may be inaccurate when multiple volumes are being migrated | Volume move operations copy both data and associated metadata from the source array to destination array. For groups with multi-array pools, copying the metadata can take a significant amount of time, and the estimate calculation may be inaccurate initially. These estimates will auto-correct themselves by using feedback mechanisms. | Not applicable |
| AS-97327 | System Management | Group Management service may restart due to communication timeout exceeded | If communication between Group Management and Postgres services does not complete within expected timeout, the Group Management service may restart. The restart will resume the communication and try the transaction again. | Not applicable |
| AS-94835 | System Management | Array Management process may restart during automatic failover | Array Management services may be unavailable for a short time due to restart during automatic failover. | Not applicable |
| AS-101420 | System Management | Array Management Service restarts unexpectedly under high load | The Array Management service may restart unexpectedly when the array is under high workload. | The service will stabilize on its own following the restarting. |
| AS-97697 | System Management | Group Management Service may restart unexpectedly | Group Management service may restart due to health check timeout exceeded. The service will recover after the restart. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|---|
| ID | Component | Title | Description | Workaround |
| AS-98434 | System Management | After group leader migration, quorum is not removed after witness disconnection | After Group migration is performed, backup group leader becomes group leader. If witness is disconnected for more than five minutes, the new leader will not remove quorum. | Disable automatic failover and remove witness from configuration. |
| AS-96241 | System Management | Group Management service may restart due to high memory usage | When the system has a high number of objects, the Group Management service may restart while running REST query or CLI command. | Not applicable |
| AS-99520 | System Management | Both upstream and downstream may claim the volume collection ownership when excessive handovers are performed | If a user performs multiple volume collection handovers between two arrays during a short time span, this may cause a situation where both upstream and downstream array may claim volume collection ownership. This is due to a race condition in the workflow. | Not applicable |
| AS-57574 | System Management | Replication of dedupe and encrypted volumes will stall between 3.4 and previous 3.x releases | Encrypted and deduped volumes cannot be replicated from 3.4.x and later to previous 3.x releases because they do not support encryption and dedupe at the same time. Note that this does not affect replication from 3.4.x and later to 2.3.x where the volumes will not be deduped on the downstream running 2.3.x. | Update the downstream array to NimbleOS 3.4.x or later. |
| AS-98650 | System Management | Alert for aborted handover does not specify reason | In the case where the downstream array is reaching its snapshot rate limit and the user performs the volume collection handover, the handover will be abort if the limit is surpassed. An alert will be raised but the alert message may be missing the reason for aborting handover. | Not applicable |
| AS-106490 | System Management | Group management service may restart due to race condition | While syncing the config changes on the downstream array, group management service may restart due to a race condition between protection policy deletion and adding the volume to the protection policy. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|--|---|
| ID | Component | Title | Description | Workaround |
| AS-98378 | System Management | Error No message received after issuing CLI command to disassociate volume from collection | Under system busy conditions, when an excessive amount of operations are being issued in parallel or too many internal retries are occurring to perform tasks, you may receive a No message received error after issuing a CLI command. | Please reissue the command. If the operation was already performed by the earlier command, an appropriate message will be returned. |
| AS-98155 | System Management | Group management service may restart unexpectedly | Arrays with volumes that have large branch structures may cause internal command processing timeout to be exceeded. This will cause the Group Management service to restart due to health check failure. | Contact HPE Nimble Storage Support. |
| AS-90633 | System Management | Error No message received after issuing CLI command to associate volume to volume collection | Under system busy conditions, when an excessive amount of operations are being issued in parallel or too many internal retries are occurring to perform tasks, you may receive a No message received error after issuing a CLI command. | Please reissue the command. If the operation was already performed by the earlier command, an appropriate message will be returned. |
| AS-90649 | System Management | Configuration of deduplication volumes for sync replication might fail | If the Default Deduplication setting differs for upstream and downstream pools, the configuration of deduplication volumes for replication might fail with the following error Deduplication not allowed since no application category is assigned to the performance policy | Update the downstream pools Deduplication setting to match the upstream pool. |
| AS-90286 | System Management | volcoll --info output lacks pool/folder qualifications for associated volumes | For the volcoll --info output for sync replication volume collections, the Associated volumes: and Associated pinned volumes: fields lack pool/folder qualification for the associated volumes. | vol --list can be used to determine pool/folder attributes of these volumes. |
| AS-89124 | System Management | Synchronous Replication Volume Count Limit | The group --list_limits CLI command does not list the Synchronous Replication volume count Limit. Synchronous Replication on 5.1.0.0 and later can protect up to 128 volumes. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|---|--|
| ID | Component | Title | Description | Workaround |
| AS-108765 | System Management | Group Management service restarts during the array shutdown | During a planned array shutdown, the Group Management service may restart due a race condition. There is no data or management interruption since the service is already shutting down. | Not applicable |
| AS-95610 | System Management | Group Management Service restarts during bulk volume update | Due to a rare race condition, the Group Management Service may restart unexpectedly during a bulk volume update operation. | Retry the command for the failed volumes. |
| AS-108868 | System Management | Group Management service may restart while collecting user information from Active Directory | In environments with an array integrated with Active Directory, the create/validate session code holds a lock on an underlying sessions table and then goes to Active Directory to collect more information about the user. If this operation takes too long, the Group Management service may timeout to free this lock. | Confirm all Domain Controllers in the Active Directory environment that is integrated with the array are reachable. |
| AS-91638 | System Management | Group Management Service restarts due to packet loss in network | If the network response to a REST request takes more than 5 minutes, a thread performing the REST request times out and as a result Group Management Service restarts. The service stabilizes itself and as long as the network is serving the requests faster. A single instance of the Group Management service restart should not cause any disruptions. | Please review the network and see if there is a consistent packet loss and fix any network glitches. If you need any assistance, please reach out to HPE Nimble Storage Support. |
| AS-66182 | System Management | Discovering volumes after array resetup may cause Group Data service restart | If a user runs --resetup followed by setup on an array that was previously configured and setup does not complete successfully, data IPs may be unconfigured until the setup actually completes successfully. As a result, GDD may restart unexpectedly if any attempts are made to discover the volumes on those unconfigured data IPs. | Ensure setup completes successfully before attempting to discover volumes. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|---|
| ID | Component | Title | Description | Workaround |
| AS-94517 | System Management | Group Management service may restart due to memory exhaustion | Group Management service may restart due to memory exhaustion in configurations that approach 10,000 volumes and 300,000 snapshots. | Not applicable |
| AS-68782 | System Management | System limitation of writable snapshots is not reported by NimbleOS | The group --list_limits CLI command does not list system limitation for writable snapshots. Also, no alerts or alarms are generated as the array group approaches the limit. | Please refer to the System Limits and Timeout Values section in the NimbleOS Administration Guide. |
| AS-97899 | System Management | Group Management service may restart due to communication timeout exceeded | If communication between Group Management and Postgres services does not complete within expected timeout, the Group Management service may restart. The restart will resume the communication and try the transaction again. | Not applicable |
| AS-105944 | System Management | Time to Live (TTL) expiry date on last replicated snapshots can be negative | NimbleOS protects the last replicated collection, in some cases, the TTL expiry date on those snapshots can become negative when the snapshots exist beyond TTL. | The TTL can be updated on the snapshots which have a negative value to a current value. The snapshot may also be removed if it has been confirmed it is no longer needed. |
| AS-105291 | System Management | Group Management Service may restart due to a race condition | Due to a race condition, the Group Management Service on a downstream group may restart while updating volume collections from the upstream group. | Not applicable |
| AS-93113 | System Management | Unmanaged snapshots remain after cleanup is enabled | If clones are created using an unmanaged snapshot, then this unmanaged snapshot will not be deleted even if cleanup is enabled. | Not applicable |
| AS-95868 | System Management | Group Management service may restart due to internal database timeout | The Group Management performs queries on the information stored in the Scale Out Database (SODB) for the array. If the query does not complete within the expected timeframe, the service may restart to recover from the issue. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|---|--|
| ID | Component | Title | Description | Workaround |
| AS-104965 | System Management | Group Management service may restart during array shutdown | During planned service stop, such as array shutdown, the Group Management service may restart due to a race condition. There is no data or management interruption since the service is already shutting down. | Not applicable |
| AS-103982 | System Management | Group Management Service may restart unexpectedly due to network connectivity | In rare instances, the Group Management Service may restart unexpectedly when the Group Leader and Member array have lost connectivity due to network outage. The service restart recovers GUI and CLI access, data services are not impacted by the restart. | Not applicable |
| AS-99615 | System Management | Array Management Service restarts unexpectedly following automatic Group Leader Failover | The Array Management Service restarts unexpectedly following automatic Group Leader Failover (AFO). The restart is non-disruptive. | Not applicable |
| AS-95132 | System Management | Process Management service may restart during software upgrade | In rare instances, the Process Management service restart may occur during software update. The system recovers after the restart of the service. | Not applicable |
| AS-98504 | System Management | Group Management service may restart unexpectedly | If internal database processing for array statistics exceeds the expected timeout, the Group Management service will restart due to health check failure to recover. | Not applicable |
| AS-105804 | System Management | Group Management service may restart unexpectedly when performing high snapshot activity | Group Management service may restart on the array when there is a high amount of snapshot activity being performed. The service restart will recover from the condition and the snapshot operations will resume. | Scheduling snapshots to occur at different times instead of all at once may help alleviate this issue. |
| AS-101392 | System Management | Services may not start on the array after it is powered on and off several times | When the array is powered on and off excessively, services may fail to start on the array. | Please contact HPE Nimble Storage Support |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|---|---|
| ID | Component | Title | Description | Workaround |
| AS-105453 | System Management | Group Management service may restart unexpectedly | The Group Management service may restart when service communication for internal database processing is terminated. The service restarts to restore connections between the services. | Not applicable |
| AS-110374 | System Management | Group Data Service may restart when the array is under heavy load during software update | The underlying scale-out database competes with CASL and other system processes for IOPS. During software update, a migration script runs against the database. Under heavy file system load, the migration steps may not complete within the expected amount of time. As a result, the migration may timeout leading to a restart of the Group Data Service. After the restart, the migration should complete as normal without any user impact or intervention. | No workaround available. To avoid encountering this issue, reduce IO load when performing software update. Software Update will succeed after one or more System Management service restarts. |
| AS-98694 | System Management | Snapshot limit warning alarms persist after update to 5.1.x.x or later | After the update to 5.1.x.x or later, the Snapshot limit warning alarm is no longer used. This presents a situation where stale alarms are present on the array and they will not be cleared even if the space situation is rectified. The alarm follows the following format: WARNING Mon DD YYYY HH:MM:SS Acknowledged - Volume <volume name>; snapshot space usage is over the configured warning limit. | The alarms can be deleted manually either in the GUI or on the CLI. |
| AS-77045 | System Management | Alarm not cleared after volume or pool drops below warning threshold | Alarms are generated on the array when volume or pools exceed thresholds. In some instances the recovery event to clear the alarm when the condition is cleared does not clear the alarm. | Confirm the usage level for volume or pool indicated by the alarm is below threshold. Once confirmed usage level is below threshold, clear the alarm from the Command Line Interface (CLI): 1. List the alarms to find the alert ID alarm --list 2. Delete the alarm that is no longer valid alarm --delete |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|---|
| ID | Component | Title | Description | Workaround |
| AS-110123 | System Management | Group Management service may become unavailable when shut down | Infrequently the Group Management service may encounter an error while shutting down. If the service is being restarted it may take a few more seconds to start. GUI and CLI will be unavailable for a few seconds. | Not applicable |
| AS-108432 | System Management | Group management service may restart when there are many REST requests for volume statistics values | High concurrent REST volume reads with statistics may cause Group Management service to restart. Current REST requests will fail, GUI and CLI will be unavailable. The system will recover after the restart of the service. | Not applicable |
| AS-106848 | System Management | Arrays with Automatic Switchover enabled fail software update with generic message | Software updates to 5.1.4.200 are not allowed when Automatic Switchover (ASO) is configured. If a software update to 5.1.4.200 fails for this reason, a generic software update failure message is returned in the GUI. The cause of the failure would need to be determined by looking at the system configuration and determining if ASO is configured. | The ASO checkbox is enabled by default, however ASO is not enabled until a witness has been configured. In the GUI, navigate to Administration > Availability. If witness is configured and the ASO check box is checked, disable ASO by unchecking the box and clicking save. Perform the array software update again. If the update continues to fail with generic messaging, contact HPE Nimble Support. |
| AS-57173 | System Management | Updating array groups from earlier NimbleOS 3.x releases to 3.3.x or later releases could timeout | Due to the number of firmware updates involved, software update going from pre-3.3.x versions to 3.3.x or later could take too long and eventually timeout. These update times may take longer in larger multi-array group configurations. Software update in smaller array groups should complete without intervention. | If a timeout occurs, the update can be completed using the software <code>--resume_update</code> command. |
| AS-72559 | System Management | Group management service may restart during software update | Group management service may restart during software update due to race condition involving unlocking the download lock file. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|---|
| ID | Component | Title | Description | Workaround |
| AS-66997 | System Management | Health check timeout may cause software update failure | The timing is close enough that it is possible for the individual array precheck during software update to take long enough that the health check timeout is triggered, causing the group management process to restart and the software update to fail. | This is an intermittent issue, so if the software update fails in this manner it should pass if the software update is resumed. |
| AS-54519 | System Management | Software update inactivity timeout messaging | When a software update is initiated on the array and the prompt for EULA acceptance is not answered, the software update session will timeout after several hours with a message indicating Requires Authentication or Contact HPE Nimble Storage Support. | Initiate the software update again and answer the EULA prompt and the software update will proceed as expected. |
| AS-40516 | System Management | Timeouts during software update | Under rare conditions, a software update may report an error even though the actual update has completed successfully. This occurs when software update takes longer than 4 hours. | Running the software --resume_update command from the console will clear this condition. |
| AS-91962 | System Management | Volume management operations may fail with unrelated error messages while handover is in-progress | While the volume ownership is being changed, any management operation involving that volume could fail. This is expected behavior. Depending on the progress of that handover operation, management operation is failing with different errors. | Ignore the error message and retry the management operation |
| AS-92209 | System Management | Group Management Service may restart unexpectedly during creation of Peer Persistent Snapshot | Group Management Service might restart unexpectedly while creating Peer Persistent snapshot during Daylight Saving Time adjustment window. | Not applicable |
| AS-109805 | System Management | Group Management service may restart due to health check timeout | The Group Management service may restart when internal database processing exceeds the expected timeout value. The service restarts to recover from the condition. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|------------------------------------|
| ID | Component | Title | Description | Workaround |
| AS-95212 | System Management | HPE Nimble Storage array compatibility issues with MIT Kerberos trust types | When the HPE Nimble Storage array is configured to use Active Directory integration, the array is joined to one specific domain, as a domain member. Under normal circumstances, users in trusted domains will also be able to authenticate to the array. If one or more trusted domains are joined to the forest using an MIT Kerberos type trust relationship, users and groups in any trusted domain (e.g. not the domain the array is joined to) will be unable to authenticate to the array. | Not applicable |
| AS-104512 | System Management | Active Directory authentication in some cases, may lead to a Group Management service restart | Active Directory Authentication causes the arrays Group management service to wait for a response from the Active Directory. If this response is delayed, the Group Management Service may restart unexpectedly. | Not applicable |
| AS-105035 | System Management | Group Management service may crash while trying to create users while array is under heavy load | In rare instances, if the array is under heavy load and Group Management service restarts while new user is being created, the service can fail to start. | Contact HPE Nimble Storage Support |
| AS-71137 | System Management | Group Management Service restarts while authenticating AD users | The Group Management Service may restart unexpectedly if it takes longer than 300 seconds to authenticate an Active Directory (AD) user. | Not applicable |
| AS-69084 | System Management | Group Management service may restart while collecting user information from Active Directory | If collection of user group information from Active Directory takes longer than expected, the Group Management service timeout may be exceeded. The service will restart to resume information collection. | Not applicable |
| AS-66437 | System Management | Command to join Active Directory (AD) may fail causing Group Management service Restart | Command to join AD may fail because of latency in getting back a response from AD server. This may cause a health check failure for Group Management causing the restart of the service to recover. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|---|
| ID | Component | Title | Description | Workaround |
| AS-74830 | System Management | Group Management service may restart when joining array to Active Directory | When joining Active Directory via GUI or userauth command in CLI, the command is expected to complete within 300 seconds. If this timeout is exceeded during the process, the Group Management service will restart to recover. | Not applicable |
| AS-78946 | System Management | Latency in communicating with Active directory during may cause Group Management service to restart | AD communication may be slow leading to a delay in responses that may cause Group Management service to crash. Leaving the AD domain requires interaction with AD wherein this delay can cause a restart. | Not applicable |
| AS-65654 | System Management | Active Directory connectivity issue may result in Group Management service restart. | If there are connectivity issues between an array with Active Directory integration enabled and an Active Directory Domain Controller, then the responses may take more than 300 seconds and timeout. As a result, the Group Management service may restart. | Resolve any connectivity issues between the array and Active Directory Domain Controller(s) to avoid unexpected service restarts. |
| AS-109549 | System Management | Group Management service may restart due to Active directory communication taking too long. | When logging in with a Active directory (AD) user, delayed responses from AD may lead to Group Management service restart. | Not applicable |
| AS-102893 | System Management | Enabling synchronous replication fails upon reaching volume limit | Following operations will fail upon reaching the volume limit: -adding Synchronous replication schedule to a volume collection -associating a volume to a volume collection with Synchronous Replication enabled - editing a Volume Collection schedule to add Synchronous Replication partner | Delete unused clones or volumes to bring down the volume count. |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|---|---|
| ID | Component | Title | Description | Workaround |
| AS-86545 | System Management | Unable to create dedupe enabled volumes on a new install | After a CSx000 array is installed, it takes one minute for the array to determine its deduplication capability. If a volume is created prior to this, it will not have dedupe enabled even if the array is dedupe capable. | Once the array is able to determine its deduplication capability, all newly created volumes will have dedupe enabled, if specified. In order to enable dedupe on the previously created volumes, you may run the following command via the HPE Nimble Storage Array CLI: <code>vol --edit &lt;vol_name>; --dedupe_enabled yes</code> |
| AS-101535 | System Management | Group Management Service is temporarily unavailable after deleting volumes | Enabling and disabling the dedupe setting on volumes and concurrently deleting volumes can cause the Group Management Services to become temporarily unavailable on the array. | The Group Management Service will eventually restart itself |
| AS-92157 | System Management | No CLI support for changing the Witness Port | Currently, there is no CLI support for changing the witness port. The <code>nimble-witnessd.service</code> file needs to be edited manually. | Not applicable |
| AS-100067 | System Management | Member array might not be displayed under Add Array to Group option | A member array might not be listed under the Add Array to Group option within the GUI if the member is configured with a different protocol (iSCSI vs Fibre Channel). Also when there are multiple arrays in the subnet, arrays which can't be discovered within the stipulated time may not be listed in Add Array to Group. | Not applicable |
| AS-107015 | System Management | Group Data service may restart on startup | In rare instances, the Group Data service may restart during startup due to a misconfiguration in NimbleOS pertaining to a file descriptor limit. | Not applicable |
| AS-99431 | System Management | Array Management Service restarts or Takeover occurs unexpectedly following automatic Group Leader Failover | In rare circumstances, following an Automatic Failover (AFO) a race condition may cause the Array Management Service to restart or an unexpected controller takeover. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|--|---|
| ID | Component | Title | Description | Workaround |
| AS-99343 | System Management | Custom SSL certificate import not supported on older versions of Google Chrome | Within the HPE Nimble Storage array GUI, custom SSL certificate import is only supported on Google Chrome version 71 or later. | Not applicable |
| AS-98177 | System Management | Setting alarm reminder frequency to the same value from GUI does not change next reminder time | When updating an alarm from the Events & Alarms page in GUI, selecting an alarm and clicking CHANGE REMINDER button, without changing the reminder frequency time, and clicking SAVE button, does not change next reminder time. This behavior is different from CLI. Setting alarm reminder frequency to the same value from CLI resets the next reminder time based on the current time. | To keep the same reminder frequency and reset the next reminder time based on the current time, change the reminder frequency to a different value, save it, and change it back and save it, or use CLI to make the change. |
| AS-87701 | System Management | Incorrect information on hardware page displayed when controller is down | When a controller is down, the user may see incorrect representation of physical ports within the Hardware Page of the array GUI. This is due to the lack of information from the missing controller. | When the controller is back up, all the information is displayed correctly on hardware page. |
| AS-102299 | System Management | GUI Error when entering a valid folder overdraft limit value | The Array GUI incorrectly returns an error when a valid value for the folder overdraft limit has been entered. This happens only in Internet Explorer and Microsoft Edge browsers. | Use Google Chrome or Mozilla Firefox browser. |
| AS-77372 | System Management | Group Merge via GUI unable to process large amount of conflicts | Currently within the HPE Nimble Storage Array GUI, when performing a group merge, if there is a large amount of group merge conflicts (1000 or more), the GUI is unable to process and resolve all of them. | To work around this issue: 1. Log into the destination array to resolve the conflicts. 2. Attempt the group merge again. |
| AS-87886 | System Management | GUI may show Successful message when group merge fails | During group merge, the GUI might show Successful message even though the group merge backend processing fails. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|---|--|---|
| ID | Component | Title | Description | Workaround |
| AS-95591 | System Management | Incorrect ordering of pool merge error messages when Synchronous Replication and Witness are configured | Pool merge is not allowed if Synchronous Replication is enabled and pool merge is not allowed when a witness configured. If an array group has a witness configured for Automatic Switchover and has Synchronous Replication configured, when a user tries to perform a pool merge, the following error is generated: "pool merge is not allowed when witness is configured". In this case, if the user removes the witness and then re-attempts the pool merge, the following error is then generated: "Pool merge is not allowed when involved in sync replication". This error should supersede the previous error. | Not applicable |
| AS-94575 | System Management | Unable to edit a storage pool and assign an array at the same time | When attempting to edit a storage pool and assign an array at the same time, you receive the following error: Cannot update array list and name or description simultaneously. | Edit the pool name and assign / un-assign the array a in separate steps. |
| AS-104099 | System Management | Volume collection association for a volume can fail due to a name conflict on the downstream array | When associating multiple volumes to a volume collection from, the volume association for all volumes can fail due to a name conflict for one of the volumes on the downstream. | Fix the name conflict on the downstream array. |
| AS-48847 | System Management | Browser throws a web server communication error | When the GUI runs for a long time, the browser may encounter an out of memory issue. | If the page hangs or throws a web server communication error, try to clean the cache and refresh. If the problem persists, restart the browser. |
| AS-93157 | System Management | Array GUI does not specify which snapshots are unmanaged | The Array GUI does not specify which snapshots are unmanaged and no longer belong to a volume collection. | Run the following command via CLI: snap --list --all --unmanaged |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|-------------------|--|--|--|
| ID | Component | Title | Description | Workaround |
| AS-99024 | System Management | Browser becomes unstable upon certificate change | After changing a certificate, the GUI may present an error such as follows: The web service is very slow or unreachable... | After a new custom certificate has been imported or existing certificate is deleted, please close the browser where the action was performed and reopen a new one to guarantee a new connection request to the NimbleOS web interface. |
| AS-92634 | System Management | Volume performance numbers may report invalid values after software update | The volume performance numbers displayed in the GUI under Manage > Data Storage > Volumes > Performance Tab may display invalid values temporarily after an array software update. | The values should report correctly within 24 hours after the update has completed. |
| AS-86901 | System Internals | Group Data Service may restart when the array is under heavy load during software update | The underlying scale-out database competes with CASL and other system processes for IOPS. During software update, a migration script runs against the database. Under heavy file system load, the migration steps may not complete within the expected amount of time. As a result, the migration may timeout leading to a restart of the Group Data Service. After the restart, the migration should eventually complete as normal without any user impact or intervention. | There is no workaround. To avoid encountering this issue, reduce IO load when performing software update. |
| AS-104567 | System Internals | Array Management Service restarts when Group Leader cannot reach Backup Group Leader | When the Group Leader attempts to complete the Backup Group Leader promotion, if there is not a healthy data path, the Backup Group Leader promotion fails. Despite, the network error, the Backup Group Leader promotion goes into a loop and ultimately leads to an unexpected restart of the Array Management Service. | Not applicable |

| Known Issues in NimbleOS version 5.3.0.0 | | | | |
|--|------------------|---|---|---|
| ID | Component | Title | Description | Workaround |
| AS-69561 | System Internals | Data Service can restart unexpectedly during shutdown process | Due to a race condition, the Data Service can crash during a graceful shutdown causing unexpected Data Services restart messages to be generated. This should not cause any I/O impact because the Data Service is already in the process of shutting down. | Not applicable |
| AS-49046 | System Internals | Data Service may restart when the array runs out of scratch pages | Under unique customer workloads, the index footprint may exhaust default scratch page allocation in memory. As a result, the Data Service may restart unexpectedly. | Please contact Nimble Storage Support to increase the scratch page pool size. |