

Hitachi NAS Platform

Docker Volume Plugin

This document details Docker Volume Plugin that allows Docker Volumes to be created and managed on the HNAS platform.

MK-92HNAS087-01

May 2023

© 2021 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively, "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at https://support.HitachiVantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

- 1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
- 2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

EXPORT CONTROLS - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, Flash Copy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, BCPI™ and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or web site are properties of their respective owners.

Table of Contents

Preface	4
About this document.....	4
Document conventions.....	4
Intended audience.....	4
Accessing product downloads.....	4
Comments	4
Getting Help.....	5
Chapter 1: Overview.....	6
Chapter 2: Docker Volumes	6
Chapter 3: Hitachi NAS Docker Volume Plugin	6
Plugin Installation and Configuration	7
Configuration Parameters	8
Volume Creation.....	9
Volume Creation Parameters.....	10
Additional Volume Creation Examples.....	11
Volume Deletion	11
Update HNAS Credentials	12
Persistent Approach	12
Alternate Approach.....	12
Chapter 4: Uninstall	13
Chapter 5: Troubleshooting	13
HNAS Docker Volume Plugin can't connect to the HNAS server:.....	13
Docker host unable to NFS mount the HNAS server:	14

Preface

About this document

This document details Docker Volume Plugin that allows Docker Volumes to be created and managed on the HNAS platform.

Document conventions

This document uses the following typographic convention:

Convention	Description
Bold	<ul style="list-style-type: none">Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK.Indicates emphasized words in list items.
<i>Italic</i>	Indicates a document title or emphasized words in text.
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>

Intended audience

This document is intended for customers and Hitachi Vantara partners who license and use NAS Platform.

Accessing product downloads

Product software, drivers, and firmware downloads are available on Hitachi Vantara Support Connect: <https://support.hitachivantara.com/>.

Log in and select Product Downloads to access the most current downloads, including updates that may have been made after the release of the product.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara.

Getting Help

[Hitachi Vantara Support Connect](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Chapter 1: Overview

The NAS server platform (Hitachi NAS platform and NAS module) can be used to provide remote storage for container images running within Docker.

All instructions in this document assume that Docker is already installed and working.

As the plugin relies on the ability for containers to mount HNAS NFS exports, it can only be used on Linux based systems.

The plugin is supported on Docker version 18 and newer, and currently only on stand-alone systems, rather than clusters/docker swarm.

Chapter 2: Docker Volumes

Docker volumes are the preferred mechanism for persisting data generated by and used by Docker containers. Volumes are easy to back up or migrate and can be managed using the Docker CLI commands or the Docker API. Volumes work on both Linux and Windows containers and can be safely shared among multiple containers. Volume drivers let you store volumes on remote hosts or cloud providers, to encrypt the contents of volumes, or to add other functionality. New volumes can have their content pre-populated and provide a way to easily provision datasets to container images.

Docker volumes are generally implemented and managed via plugins. They provide a means to manage the volumes and can be used to effectively share part of the Docker hosts filesystem. Plugins provide a mechanism to allow remote storage to be attached to a Docker container. NFS is commonly used to provide remote storage.

Chapter 3: Hitachi NAS Docker Volume Plugin

The Hitachi NAS Docker Volume Plugin allows storage to be allocated on HNAS systems from the docker command line. It is able to manage the creation and deletion of NFS exports, virtual volumes and allows existing data already present on the HNAS to be cloned and presented to multiple different containers. It can also delete the volume contents when it's no longer needed, if configured to do so. NFS must be licensed on the HNAS system for the plugin to work, and a FileClone license must be present on the HNAS system for the clone operation to work.

Backup and restore of important data contained within container images can be an issue under normal circumstances. By providing NFS mounted remote storage, it's possible to use normal backup/restore tools to ensure that the important data associated with the containers can be properly backed up and restored should it be necessary.

The plugin does not get involved with the normal operation of the containers, and NFS mounted exports behave in exactly the same way as if they were created and mounted manually. The plugin makes it easier to create external storage on the HNAS system, without the need to directly interact with the HNAS management interfaces.

The plugin communicates with the HNAS system via the management REST API, which must be enabled on the HNAS system, and requires HNAS software 13.3 or newer.

The plugin is also able to instruct the Docker host to mount the NFS exports once they have been created. Specific mount options can be supplied as part of the Docker command line – for example, which NFS version to use. NFSv3 is recommended when mounting HNAS volumes, as it is generally faster than NFSv4, and has less overhead.

One restriction on the plugin is that it does not work on a Windows Docker host, only Linux based systems.

Plugin Installation and Configuration

As a first step, it is advisable to create a separate SUPERVISOR level user on the HNAS system for use with the plugin. To create the new user, run the following command from the HNAS command line interface:

```
hnas-1:$ user add
```

The command will prompt for the new username, the user level, which must be “SUPERVISOR” and then for the password.

Before installing the plugin, check that the docker host can successfully communicate with the HNAS REST API. Refer to the “HNAS Docker Volume Plugin can’t connect to the HNAS server” topic in the Troubleshooting section of this document for details on how to prove connectivity.

The plugin is supplied as a compressed tar file (.tgz file). To install the plugin copy the tgz file to the /tmp folder on the Docker host, and follow the steps below, refer to the “Install and Configuration Parameters” section for details on the command line parameters that need to be supplied as part of the `docker plugin set` command:

```
# cd /tmp
# tar -xvf hnas-volume-v1.0.tgz
# docker plugin create hnas-volume:latest ./plugin
# rm -rf plugin
# docker plugin set hnas-volume:latest apiurl=<apiurl> username=<username>
password=<password>
# docker plugin enable hnas-volume
```

Check plugin is installed:

```
# docker plugin list
ID                NAME                DESCRIPTION                ENABLED
398cb2a1519d     hnas-volume:latest  Hitachi NAS volume plugin  true
```

The plugin is only able to manage volumes on a single HNAS system/cluster, but it’s possible to install multiple copies of the plugin on the same Docker host. To install multiple copies of the plugin managing different HNAS systems, repeat the install steps above, but replace `hnas-volume` with a different name. *It’s best to use a unique name for each install that references each HNAS system if more than a single instance is to be installed.*

Configuration Parameters

It is also possible to modify any of the parameters after installing the plugin (if it's disabled) by using the following command:

```
docker plugin set hnas-volume KEY=VALUE [KEY=VALUE]
```

where KEY is one of the parameters listed below and VALUE is what the parameter is to be set to. As a minimum, the `apiurl`, `username` and `password` parameters need to be set for the plugin to work.

```
apiurl=<required>
```

The base URL used to connect to the HNAS REST API from the plugin. This parameter is required, and there is no default value. The address should be a publicly accessible admin vnode address, and must be accessible to the Docker host. The format should be `https://10.1.2.3:8444` – https and port 8444 are the only options currently supported.

```
apiversion=4
```

The version of the HNAS REST API to use - currently only versions 4 and 7 are supported. The default value is 4.

```
username=<required>
```

The REST API username that the Docker plugin uses to connect to the HNAS server – this needs to be a SUPERVISOR level user.

```
password=<required>
```

The REST API password associated with the `username` parameter.

```
managementIP=127.0.0.1
```

Only applicable if `apiversion 4` is used, and shouldn't need to be changed – default is 127.0.0.1. This address is used internally within the HNAS system.

```
verifyssl=false
```

Whether the plugin should verify the SSL certificate when connecting to the HNAS API. For initial install and testing, it is best to leave this parameter as the default. If this parameter is set to `true`, then an appropriate valid certificate will need to be installed on the HNAS server. Installation of a valid certificate is beyond the scope of this document.

Volume Creation

The simplest way to create a new volume using the HNAS Volume Plugin is using the following format:

```
docker volume create --driver <driver name> \  
  -o filesystem=<HNAS filesystem> \  
  -o <NFS mount options> \  
  <name of new volume>
```

This will create a new NFS export, matching the `<name of new volume>` on the HNAS associated with the `<driver name>` parameter, on the specified filesystem. The export created on the HNAS will point to a directory with the same name in the root of the specified filesystem. The export will then be mounted using any (optional) additional `<NFS mount options>` supplied on the command line.

The following shows a specific example that creates a new NFS export called `hnavoll` on filesystem `fs1`, and it's mounted using NFS version 3 (`vers=3` option).

```
# docker volume create --driver hnas-volume -o filesystem=fs1 -o vers=3 hnavoll
```

To check the volume has been successfully created, run the following command:

```
# docker volume ls  
DRIVER          VOLUME NAME  
hnav-volume:latest hnavoll
```

Once the new volume has been created, it can be mounted by a container using the `-v` command line option, which is used to specify the new volume and the mount point within the container. As with any Docker volume, it is not possible to mount sub-folders within a volume, only the root may be mounted.

The following example instructs Docker to run a `Debian:stretch-slim` image, and mount the `hnavoll` volume to the local container path `/mountpath`. It then starts the container in interactive mode, and starts a bash prompt. From the bash prompt, a directory listing of the remote storage is shown – it is initially empty, except for the `.snapshot` folder, which could potentially hold HNAS snapshot data.

```
# docker run --rm --name debian-test -v hnavoll:/mountpath -it debian:stretch-slim  
/bin/bash  
root@0bda9bd2192a:/# ls -al /mountpath/  
total 12  
drwxrwxrwx  3 root  root   2048 Jun  7 12:23 .  
drwxr-xr-x 28 root  root   4096 Jun  7 13:46 ..  
drwxrwxrwx  2 root  root   2048 Jan  1  1970 .snapshot  
root@0bda9bd2192a:/#
```

Volume Creation Parameters

In addition to the standard Docker volume creation parameters, some additional parameters are provided by the HNAS Docker Volume Plugin which are available during volume creation – only the `filesystem` parameter is required, all others are optional. Each volume creation option must be preceded by `-o`

`-o delete`

This is a flag, and therefore takes no value. It indicates that when the volume is removed, the export's path should be recursively deleted from the filesystem. If the export's path is `/`, the flag will be ignored (and an error logged). Creating a volume with this flag will fail if the export already exists. If this option is omitted from the creation command, the files on the HNAS will be left intact when the Docker volume is deleted on the host.

`-o filesystem=<hnas-fs-name>`

The HNAS filesystem on which to create the docker volume.

`-o ip=<file-serving-IP-address>`

An available IP address of the EVS that owns the given filesystem. If none is provided, the volume will attempt to find an accessible IP address associated with the EVS hosting the filesystem. This parameter may be needed if the NFS mount fails once the HNAS export has been created, or if the EVS has multiple IP addresses on different subnets.

`-o config=<nfs-export-access-config>`

An access configuration for the export that will be created for the volume. Refer to the Hitachi NAS File Services Administration guide, SMU online help or man pages for the list of available options and format. A simple example would be: `"*(root_squash)"`

`-o clone=<volume>`

The name of a volume/export to clone for the new volume. Creating a volume by cloning will fail if the export already exists. This uses a tree clone job on the server.

`-o path=<base-path>`

The path to use for the export. If this isn't given, the path defaults to `/<name of new volume>` (i.e. a directory named after the volume in the root directory). If this option is provided and an export with the same name already exists, then the given path must match the export's path or an error will be returned.

`-o size=<size-in-bytes>`

If this parameter is specified, then an HNAS Virtual Volume, with a size- based quota is also created. This means that the amount of space reported to the operating system running within the Docker container will match the size specified here, and that the size restriction will be enforced via a hard quota. If this parameter is not used, the size reported to the operating system running within the Docker container will be the full size of the filesystem that hosts the NFS export, and no restrictions will be made on how much data can be written to that filesystem. The size can be specified in multiple formats – namely bytes, 1.2G or 500MiB.

Note: If a volume is created using the “size” parameter, the HNAS Virtual Volume size may later be adjusted by an administrator using one of the HNAS management interfaces, but it is not possible to adjust the size of an existing volume using Docker.

```
-o email=<email-list>
```

A comma separated list of email addresses that will be associated with the virtual volume that gets created if the size parameter is used. Emails are sent when the storage gets filled to certain levels – a warning message will be sent at 75% full and a severe message at 95% full.

Any extra options, which must also be prefixed by `-o`, are passed as mount options when mounting the export by the container. For example, creating a volume with:

```
# docker volume create --driver hnas-volume -o filesystem=<fs-name> -o vers=3 -o
timeo=600 myvolume
```

would result in the `vers` and `timeo` options being passed to the NFS mount command, when the first container attempts to use it:

```
mount -o vers=3 -o timeo=600 ...
```

Additional Volume Creation Examples

To make it easier to use and manage external NFS stored volumes, it is best to make use of a few of the additional options available when creating volumes.

Rather than placing all exports in the root or a filesystem, it is best to place them under a single subdirectory, and then export folders from within that subdirectory. To achieve this, use the `path` option to specify the location of the folder used to store the contents of the volume.

```
# docker volume create --driver hnas-volume -o path=/docker/voll -o filesystem=fs1 -o
vers=3 hnasvoll
```

As docker containers quite often run as the root user, it is important to make sure that unrestricted access is provided to the process running within the container. By default, the root user gets the user and group ID squashed to the unknown user, thus preventing a root user from being able to perform certain filesystem operations. To disable this behavior, it's necessary to specify the NFS export is created with the “norootsquash” option. Use the `config` option to specify this behaviour.

```
# docker volume create --driver hnas-volume -o config=(norootsquash) -o filesystem=fs1
-o vers=3 hnasvol2
```

Volume Deletion

To delete a volume, it must no longer be in use by any containers, this includes containers which are no longer running. If the container persists in a stopped state, the container must also be deleted before the volume can be deleted.

Get the name of the volume to be deleted.

```
# docker volume ls
DRIVER          VOLUME NAME
hnas-volume:latest  hnasvol1
```

Delete the volume, using the normal `docker volume rm` command.

```
# docker volume rm hnasvol1
```

When the volume is deleted, the HNAS is instructed to delete the NFS export. If the volume was created with the `delete` option, the HNAS will also be instructed to delete all files contained within the exports folder. This is done as a background process, using HNAS tree delete functionality.

Update HNAS Credentials

Due to limitations with the Docker plugin architecture, it is not possible to make changes to running plugins, and generally, they must be stopped to make changes to their settings – plugins can't be stopped while containers have volumes mounted either, so changing the credentials would require an outage of the running containers. Two approaches are available to change the HNAS Volume plugin credentials – persistent, and an alternative approach, which can update the running plugin, but will be lost on a restart.

Persistent Approach

To persistently change the HNAS credentials that the plugin uses requires the plugin to be disabled. Disabling the plugin can't be done while containers have volumes mounted.

To persistently change the credentials, stop all containers that are using HNAS volumes, and then stop the HNAS volume plugin:

```
# docker plugin disable -f hnas-volume
```

Change the credentials.

```
# docker plugin set hnas-volume username=<username> password=<password>
```

Re-enable the plugin.

```
# docker plugin enable hnas-volume
```

Then restart any containers that were using the HNAS volumes.

Alternate Approach

It is possible to update the credentials for a running instance of the HNAS Volume plugin without stopping all the containers using existing HNAS volumes, however they are not persistently saved, so would revert to the previous credentials should the HNAS plugin or Docker host get restarted.

The following command can be used to update the credentials – Note an error message will be returned, but it should indicate that the running credentials have been updated.

```
# docker volume create -d hnas-volume -o updatecredentials -o username=<username> -o
password=<password>
Error response from daemon: create
2b7fc405f5851811051b6abf976c1990c3faf455a94357a9717dc23660cd1972: VolumeDriver.Create:
The username and password have been updated in the running environment - this change
will be lost when the plugin is disabled or restarted.
```

Note: If this method is used, the next time the Docker host is restarted, the Persistent Approach should be used to permanently change the credentials.

Chapter 4: Uninstall

Before the HNAS Volume plugin can be removed from a Docker host, it is necessary to unmount and delete all volumes created with the plugin. Refer to the “Volume Deletion” section of this document for details on how to delete any volumes.

Delete the plugin with the following command:

```
# docker plugin rm hnas-volume
```

Chapter 5: Troubleshooting

HNAS Docker Volume Plugin can't connect to the HNAS server:

Ensure that the HNAS REST API is enabled. This can be checked via the server command line interface as follows:

```
hnas-1:$ rest-server-status

Node      Rest Server
----      -
1         running
2         running
```

If the status is “not running” for either node, then the REST server can be started with the following command:

```
hnas-1:$ rest-server-start
The RESTful API server has been successfully started.
```

The connection parameters and credentials can be checked using a simple curl command, run on the Docker host machine. Run the following command, and replace the <username>, <password> and <admin_IP> parameters with the ones to be used by the plugin. If the command works, it should return some information about the cluster nodes, and some additional debug info associated with the connection. If the credentials are wrong, a message should be returned with a “401 Unauthorized” error code. If the <admin_IP> address is wrong, the connection should either be refused, or timeout.

```
curl -v -k -H "X-Subsystem-User: <username>" -H "X-Subsystem-Password:
<password>" -H "X-Subsystem-Type: FILE" https://<admin_IP>:8444/v4/storage/nodes
```

For more details on troubleshooting connectivity issues to the REST API, refer to the “Hitachi NAS File Storage, REST API Reference Guide - MK-92HNAS088-00”.

Docker host unable to NFS mount the HNAS server:

If an error is returned when attempting to use the HNAS hosted volume with a container – see the example below:

```
# docker run -v hnasvoll:/mountpath -it debian:stretch-slim /bin/bash
docker: Error response from daemon: error while mounting volume '/mnt/volumes/18c5f9ba-10c6-4697-b0c9-c72f9205f9d5': VolumeDriver.Mount: exit status 32.
```

This indicates that the NFS mount failed. The most likely cause is that the HNAS EVS hosting the newly created NFS export has multiple IP addresses, which are not all contactable from the Docker host, and that one of the non-routable ones have been selected for use. To verify that this is the case, it is necessary to extract the debug message from the syslog of the Docker host – this can be done as follows. Use the UUID from the original error message to search within the syslog file for the debug error message:

```
# cat /var/log/syslog | grep 18c5f9ba-10c6-4697-b0c9-c72f9205f9d5
Nov 11 15:25:49 hnas-docker-test dockerd[589]: time="2019-11-11T15:25:49Z" level=error
msg="time=\"2019-11-11 15:25:49.707\" level=debug msg=\"[mount 172.16.19.250:/hnasvoll1
/mnt/volumes/18c5f9ba-10c6-4697-b0c9-c72f9205f9d5 -o vers=3]\"
plugin=173e11fd97d73fef5a49e6317f6b1634f41f4ef077be938512f15e9b6a37ad62
Nov 11 15:27:01 hnas-docker-test dockerd[589]: time="2019-11-11T15:27:01.416036747Z"
level=error msg="Handler for POST
/v1.40/containers/5c907e751a85dc917463a7c7c01da6b2adaa2408f0ea3bd66e58fcb1d4fe478e/start
returned error: error while mounting volume '/mnt/volumes/18c5f9ba-10c6-4697-b0c9-
c72f9205f9d5': VolumeDriver.Mount: exit status 32"
```

The complete mount command can be extracted from the debug message, and it shows the IP address and NFS export used to perform the mount request, along with any options specified – shown in bold above. If the address used in the mount command is not contactable by the Docker host, then delete the volume and recreate it specifying the `-o ip=<file-serving-IP-address>` parameter and set it to one of the EVS IP addresses that is contactable from the Docker host. A list of all file serving addresses can be obtained from the SMU.

Another potential NFS mount failure condition, which would also fail with the same error message, is that by default HNAS systems are not configured to support NFSv4, and if no mount options are specified to the volume creation command, the Docker host may attempt to use NFSv4 to mount the HNAS export. There are two solutions to this issue - enable NFSv4 on the HNAS or specify NFSv3 as an option when creating the Docker volume. Specifying NFSv3 as a mount option should be the preferred solution, as enabling NFSv4 on an HNAS system that does not already have it enabled could potentially cause other NFS clients to mount via NFSv4 when next mounted, which may have unintentional behavior changes.

Hitachi Vantara

Corporate Headquarters 2535 Augustine Drive
Santa Clara, CA 95054 USA www.HitachiVantara.com community.HitachiVantara.com

Regional Contact Information

Americas: +1 866 374 5822 or info@hitachivantara.com

Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com

Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com

