AutoCDP Plugin User Guide

Table of Contents

Setting up remote mirroring – CDP

NexentaStor provides CDP (Continuous Data Protection) capability via pluggable module.

For information on NexentaStor plugins, please visit NexentaStor F.A.Q. pages, or see Section "Frequently Asked Questions" in this document.

Introduction

Auto-CDP (Automatic Continuous Data Protection a. k. a. **auto-cdp**) service provides remote block-level mirroring capability and is distributed as a NexentaStor pluggable extension (plugin).

The service allows to replicate disks between two different NexentaStor appliances in real time, at a block level.

Conceptually, the service performs a function similar to local disk mirroring scheme of RAID1, except that in the case of Auto-CDP this is done over IP network.

Auto-CDP services operates in synchronous mode. Synchronous mode is the Remote Mirror mode where the I/O operation is not confirmed as complete until the remote volume has been updated.

Similar to all NexentaStor extension, Auto-CDP plugin has its own page on the webiste:

- Auto-CDP plugin web page

Please visit this page for the most recent information, updates, and download instructions.

The NexentaStor Auto-CDP service plugin provides the following features:

- Fully integrated into the NexentaStor Appliance and provide the wizard UI for NMC;
- Utilizing enterpise-grade robust Sun AVS (Sun Availability Suite) framework, SNDR service;
- ZFS integrated with major pool operations such as disk removal and pool grow;
- Provides autocdp-check NexentaStor trigger to monitor states and statuses;
- Socket-level transfers monitoring functionality.

Installation

Before installing the Auto-CDP plugin, please make sure that data volume (that you intend to replicate) exists, both appliances are SSH-bound and networking connectivity is properly setup. Please refer to NexentaStor User Guide for details on how to create data volume, setup SSH and networking connectivity.

Even if your data volume is unused, the initial syncing will take significant amount of time necessary for block-level sector-by-sector transfer of all its (unused) blocks over the IP network.

Please see the following F.A.Q.:

- I'm using auto-cdp plugin to block-mirror my storage. Initial replication is very slow..

The Auto-CDP plugin requires NexentaStor Trial or Enterprise Edition. The plugin can be remotely uploaded into your installation-private APT repository. For the most recent information on the plugin, please visit its page on the website: Auto-CDP plugin web page.

To verify that Auto-CDP plugin is available for installation, run the following command:

```
nmc$ show plugin remotely-available
```

To install Auto-CDP plugin, please run the following command:

```
nmc$ setup plugin install autocdp
```

Or, you can view, install and uninstall the NexentaStor extension module using appliance's web GUI, as shown below:

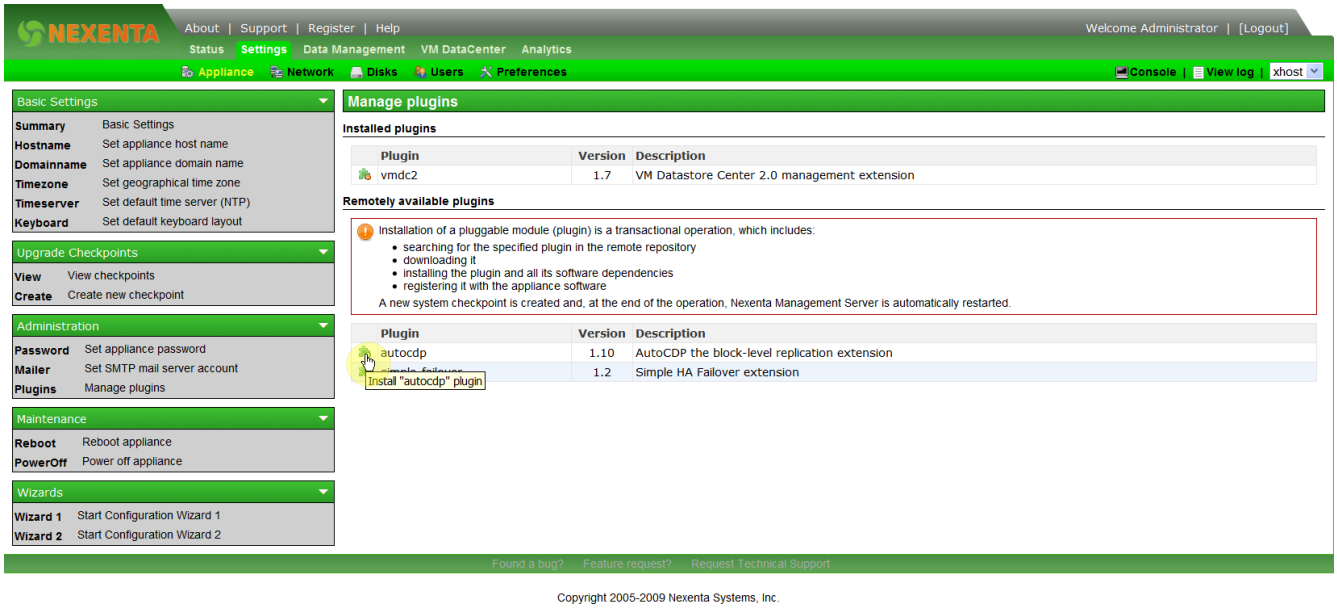

If you are using the appliance's web GUI to install plugin, the operation will need to be confirmed:

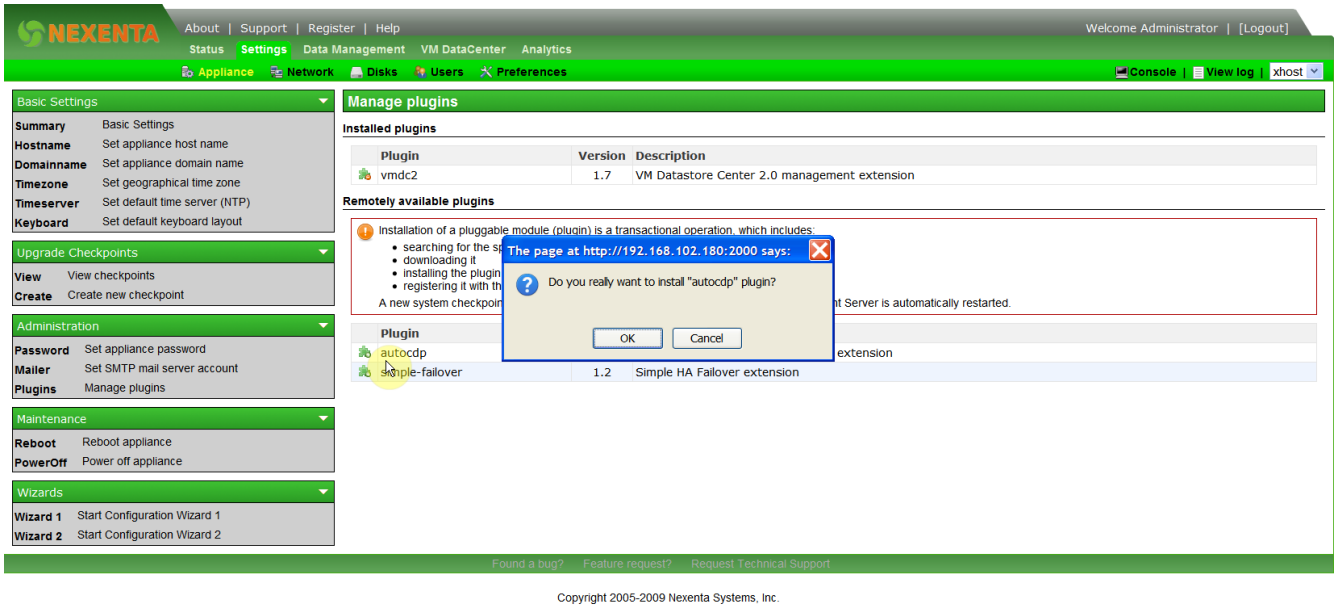

The installation will require NMS restart and NMC re-login. After installation, use NMV or NMC to verify that the plugin is successfully installed. In NMC, that corresponding operation would be:

```
nmc$ show plugin autocdp
```

The command will display plugin version, as well as other useful information. For more details on plugin management in NexentaStor please refer to User Guide.

Getting Started

Creating of Auto-CDP service instance is easy, and takes the following steps:

1.Select local (**primary**) data volume to replicate. The name of the service instance is in form "**:volname**".

Document Convention

Assuming there is a volume named '**vol1**', the corresponding Auto-CDP service will be named **:vol1**.

Here, and throughout the rest of this document, **:volname** indicates the name of the volume to block-mirror using Auto-CDP, and simultaneously, the name of the corresponding Auto-CDP service.

Note, Auto-CDP service cannot be created for **syspool** – the appliance's system volume.

2.Select remote appliance. Specifying existing SSH-bound appliance registered on local appliance-creator;

3.Select disks on the remote appliance to serve as block-level replicas of the disks of the local (primary) volume.

As always, to carry out the 1-2-3 steps, NMC provides a guided multiple-choice interactive environment. As always, the same steps can be executed via command line, using the options specified above.

Once initiated, Auto-CDP will transfer the local (primary) volume's metadata, which will effectively create a secondary (remote) volume out of the corresponding remote disks.

The appliance's Auto-CDP will keep both data and ZFS metadata on the replicated disks in-sync, at all times.

Note: Auto-CDP requires using either DNS hostname for the local and remote appliances, or their "replacement" via local host tables. See 'setup appliance hosts -h' for more information.

The following NMC wizard command can be used for service instance creation:

```
nmc$ create auto-cdp
```

```
nmc$ create auto-cdp -h
```

```
nmc$ setup auto-cdp :volname
```

where "**:volname**" is a service instance, and simultaneously, the name of the volume that is being block-mirrored via (or by) Auto-CDP.

The alternative hostname

The alternative local hostname can be specified instead of the one automatically selected by the Auto-CDP service.

Background:

> Auto-CDP identifies the primary <=> secondary nexus by the specified IP addresses and their corresponding fully qualified hostnames. In a simple case there is a single networking interface and appliance's fully qualified domain name (FQDN). If the appliance has multiple IP interfaces, an attempt is made to find the matching primary IP to be used for the specified secondary IP address. The logic to find the best matching pair of interfaces for the CDP data transfers is built-in.

To override the default logic and specify an alternative hostname, use the -H option as shown below during service creation command:

```
nmc$ create auto-cdp -H althost
```

where "althost" is the alternative hostname to be used for the Auto-CDP service instance created.

Enabling Auto-CDP service instance

After Auto-CDP service instance was successfully created the wizard will ask permission on service enable. Service can be enabled/disabled at any given time using the following command:

```
nmc$ setup auto-cdp :volname enable
```

This enables the remote mirror replication for the primary volume and also uses the remote

mirror scoreboard logs to start the resynchronization process so that the corresponding secondary data volume becomes a full replica of the primary volume.

Sizes of the remote (**secondary**) disks (a. k. a. LUNs) must be equal or greater than the corresponding primary disks that are being replicated.

Once enabled, NexentaStor Auto-CDP service instance will update a remotely mirrored data volume. Only the blocks logged as changed in the remote mirror scoreboards are updated.

Use '-f' (force) option when the primary and the secondary volumes/luns might be different and no logging information exists to incrementally resynchronize the volumes/luns.

Reverse synchronization and DR (disaster recovery)

At some point in time the secondary setup will be used as a disaster recovery (DR) site. There are two scenarios which needs to be considered while failing over to secondary setup:

1.Primary site is still active and you just want to manually switch to the secondary for primary maintenance operations. The assumption is that Auto-CDP service instance was in replication mode. In this case, the following command needs to be executed on primary:

```
nmc$ setup volume volname export
```

The command above will gracefully disable instance ":volname" after export completes. After command is complete, the following command needs to be executed on secondary:

```
nmc$ setup volume import volname
```

The "**:volname**" instance will stay in logging mode and once primary site is up the export/import operations can be repeated. After import command is complete on primary, execute the following command to enable reverse synchronization back from secondary to primary:

```
nmc$ setup auto-cdp :volname enable -r
```

The '-r' option (reverse) used to reverse the direction of synchronization: that is, synchronize from the secondary volume to the primary. With the '-r' option the primary volume becomes a passive destination while the secondary volume is considered active source (of the changes).

2.Primary site is down and you need to forcefully switch to secondary. The graceful Auto-CDP service instance disable is not possible in this case and loss of data may occur. However, the filesystem on a secondary setup is going to be always in consistent state due to transactional nature of ZFS and synchronous mode of SNDR operation. On the secondary setup, execute the following command to forcefully import data volume:

```
nmc$ setup volume import volname -f
```

The '-f' option will disregard host checking and forcefully import data volume on a secondary setup. The rest of operations could be done similar to (1).

The reverse operation '-r' then resumes Remote Mirror replication of new updates from the primary volume to the secondary volume automatically so that the volume sets remain synchronized. We recommend to quiesce the workload to the volume sets during the restore/ refresh operation. This action ensures that the primary and secondary volumes match before

replication of new updates resumes.

Volume operations and Auto-CDP

Auto-CDP service tightly integrated with main NMC/NMS management commands which simplifies service maintenance. The following data volume operations supported:

1.Volume grow. Add new LUN or group of LUNs to an existing data volume will interactively invoke Auto-CDP NMC wizard for disk pair setup;

2.LUN removal or replacement will ensure that appropriate Auto-CDP NMC wizard is invoked

3.Volume export will ensure that corresponding Auto-CDP instance is disabled if desired;

4.Volume import will ensure that corresponding Auto-CDP instance is enabled if desired;

5.Volume destroy will ensure that corresponding Auto-CDP instance is removed;

6.Simple-Failover service ensures that Auto-CDP configuration is securely transferred to all machines with-in the simple-failover group and automatically activates Auto-CDP service on failover machine;

Service monitoring

There are number of ways to monitor Auto-CDP service health and statuses:

1.Automatic NexentaStor trigger “autocdp-check”. As any other NexentaStor trigger it will send notification events (e-mails) if some of services instances suddenly enters logging mode or when initial syncing enters into normal replication mode;

2.AVS level monitoring of service status and I/O progress. The following command can be used:

```
nmc$ show auto-cdp :volname iostat

name              t s   pct   role  rkps  rtps wkps wtps
v/rdsk/c2t0d0s0   P L  0.05  net     8     3    0    0
v/rdsk/c2t1d0s0   P L  0.40  net     0     0  345    0
```

The output includes the following information:

t	Volume type
s	Volume status
pct	Percentage of volume requiring sync
role	Role of the item being reported
rtps	Number of reads
wtps	Number of writes
rkps	Kilobytes read
wkps	Kilobytes written

3.Service-generated socket-level network traffic. The following command can be used:

```
nmc$ show auto-cdp :volname stats
```

4.Service logs can be viewed by executing the following command:

```
nmc$ show auto-cdp :volname log
```

5.To monitor logging activity interactively, execute the following command:

```
nmc$ show auto-cdp :volname logtail -f
```

Auto-CDP configuration properties

Auto-CDP service provides the following configurable properties:

1.**cdproot** - AutoCDP internal database root folder. Location where NMS and Auto-CDP service will store its internal data (AVS bitmaps devices). Default is “syspool/.cdp”. This option allows user to create dedicated volume to store zvol-emulated bitmap devices and isolate system volume;

2.**cdp_scoreboard_log_size** - specifies Auto-CDP scoreboard log size. Default is '1G' (one gigabyte). The AVS bitmap devices. Zvol-emulated and thin-provisioned.

Service States

The following are the states that can be set to the operational service instance:

- **volume failed (VF)** - an I/O operation to the local data volume has failed ;
- **bitmap failed (BF)** - an I/O operation to the local bitmap volume has failed ;
- **need sync (SN)** - a sync to this volume has been interrupted. It needs to be completed. The direction of the data flow must not be changed until one or the other is done;
- **need reverse sync (RN)** - a reverse sync to this volume has been interrupted. It needs to be completed (or restored via Point-in-Time Copy). The direction of the data flow must not be changed until one or the other is done;
- **logging (L)** - incoming writes are logged in the bitmap devices only. Data is not replicated to the remote site. **need sync** and **need reverse sync** are all sub-states of **logging** such that writes are logged in the bitmap, but not replicated;
- **reverse syncing (RS)** - a secondary to primary copy is in progress;
- **syncing** – a primary to secondary copy is in progress.

Troubleshooting

To troubleshoot, execute either one of the following commands:

1.To re-enable the entire service corresponding to instance “**:volname**”:

```
nmc$ setup auto-cdp :volname enable
```

2.To re-enable the entire service and fully resynchronize the associated primary volume “volname” to secondary volume. Beware, this operation may take a long time to complete:

```
nmc$ setup auto-cdp :volname enable -f
```

3.To re-enable specific <LUN> pair for service instance “:volname”:

```
nmc$ setup auto-cdp :volname lun <LUN> enable
```

4.To re-enable and fully resynchronize the specific <LUN> pair for service instance "**:volname**":

```
nmc$ setup auto-cdp :volname lun <LUN> enable -f
```

5.To reset all Auto-CDP services and re-initialize AVS databases on both sides active and passive:

```
nmc$ setup auto-cdp reset
```

WARNING! This operation will reset Auto-CDP service to its initial (post-creation) state.

6.As a troubleshooting example, to replace failed drive just run normal volume command:

```
nmc$ setup volume volname replace-lun
```

The major difference between all of the commands listed above is: **granularity**.

The first two commands (1,2) execute on a level of the entire service instance, with the corresponding action applied to all associated disk pairs. Use command (1) if you want to move service from "logging" mode back to "replication".

The second pair (3,4) of troubleshooting actions is LUN specific. These two commands (3,4) are especially useful when a single or a few specific pairs of syncing LUNs appear to be stuck in a so called "logging" mode and will not change states. Another relevant scenario is related to importing of the secondary volume. If the newly imported mirrored volume shows faulted drive(s), use the LUN specific re-synchronization to troubleshoot.

The reset command (5) is plugin/service wide and will affect all instance on both active and passive sides.

The disk replacement, as well as all other disk management operations, tightly integrated with the service and the right action will be taken if corresponding Auto-CDP instance is present. Simply execute normal volume management operation and do not worry about the complexity associated with AVS disk set management.

Creating auto-cdp – example

In the following example appliance '**testbox1**' is a primary, '**testbox2**' is a secondary. This example includes all NMC prompts – it is a complete demonstration of **auto-cdp** creation:

```
nmc@testbox1:/$ create auto-cdp
Remote appliance  : 192.168.37.128
Remote for c2t1d0 : c2t1d0
Remote for c2t0d0 : c2t0d0
Creating new Auto CDP service 'auto-cdp:vol1', please wait...
Successfully created service 'auto-cdp:vol1'
Enable it now?  Yes
Enabling service, please wait...
PROPERTY                VALUE
name                    :vol1
max_q_fbas              16384
autosync                off
max_q_writes            4096
async_threads           2
state                   syncing
to_host                 testbox2
from_host               testbox1
type                    active
TESTBOX1          TESTBOX2
c2t1d0      =>      c2t1d0
c2t0d0      =>      c2t0d0
The local host is 'active' auto-cdp node.
```

Once the initial synchronization between a pair of active (primary) and passive (secondary) volumes commences, you can monitor it either via '**show auto-cdp <name> stats**' or '**show auto-cdp <name> iostat**' NMC commands.

In fact, these two commands are always useful, in terms of monitoring the data replication traffic, whether this is **auto-cdp**, **auto-sync** or **auto-tier** service. However, **auto-cdp** traffic monitoring is particularly useful at the time of the initial block-level syncing:

It is recommended not to use the primary (active) volume during the very first (the initial) CDP synchronization. Any updates on the primary during this period of time may considerably delay the initial synchronization. Note also that during this phase a major part of the available I/O bandwidth is used by the auto-cdp service, which is yet another reason to let it run through as soon as possible.

See '**show auto-cdp <name> stats**' for more information.

Following is an example of '**show auto-cdp <name> stats**' output:

```
nmc@testbox1:/$ show auto-cdp :vol1 stats -i 1
TCP CONNECTIONS                                SNEXT      RNEXT      TRANSFER
192.168.37.128.1022-192.168.37.134.121         1313611534 3140553278 1.60 MB
192.168.37.128.1022-192.168.37.134.121         1314180374 3140554114 569.68 KB
192.168.37.128.1022-192.168.37.134.121         1314838374 3140554994 658.88 KB
192.168.37.128.1022-192.168.37.134.121         1316976874 3140557854 2.14 MB
192.168.37.128.1022-192.168.37.134.121         1321352574 3140563706 1.25 MB
192.168.37.128.1022-192.168.37.134.121         1327471974 3140571890 955.38 KB
192.168.37.128.1022-192.168.37.134.121         1328722174 3140573562 1.25 MB
...
```

Once the traffic stops, you'll be able to see the block-level replicated volume on the remote side:

```
nmc@testbox2:/$ show auto-cdp :vol1 -v
PROPERTY                VALUE
name                    :vol1
max_q_fbas              16384
autosync                off
max_q_writes            4096
async_threads           2
state                   logging
to_host                 testbox2
from_host               testbox1
type                    passive
TESTBOX1          TESTBOX2
c2t1d0      =>      c2t1d0
c2t0d0      =>      c2t0d0
The local host is 'passive' auto-cdp node.
```

Reverse mirroring – example

In the following 6-steps example appliance '**testbox1**' is again a primary, and '**testbox2**' is a secondary. The reverse mirroring starts from exporting a volume from the primary appliance (Step #1)...

One critically important guideline in re CDP:

It is recommended NOT to have the primary and secondary volume imported simultaneously. In fact, NexentaStor software will prevent this from happening.

Still, note: the remotely mirrored volume may be imported only at one side, primary or secondary, at any given moment.

In short, several preparation steps need to be performed before actually enabling reverse mirroring from '**testbox2**' to '**testbox1**' (Step #5 below):

```
Step #1. testbox1 (primary): first, export vol1
nmc@testbox1:/$ setup volume vol1 export
Export volume 'vol1' and destroy all associated shares ?  Yes

Step #2. testbox2 (secondary): import vol1

nmc@testbox2:/$ setup volume import vol1
volume: vol1
 state: ONLINE
 scrub: none requested
config:
        NAME        STATE     READ WRITE CKSUM
        vol1        ONLINE       0     0     0
          mirror    ONLINE       0     0     0
            c2t0d0  ONLINE       0     0     0
            c2t1d0  ONLINE       0     0     0

Step #3. ...using secondary volume until (and if) the problem with
primary is resolved...

Step #4. testbox2 (secondary): export vol1
nmc@testbox2:/$ setup volume vol1 export
Export volume 'vol1' and destroy all associated shares ?  Yes

Step #5. testbox1 (primary): reverse syncing
nmc@testbox1:/$ setup auto-cdp :vol1 enable -r
Enable reverse synchronization for auto CDP service 'vol1'?  Yes
Enabling service, please wait...
Auto CDP service ':vol1' enabled.

Step #6. testbox1 (primary): import vol1
nmc@testbox1:/$ setup volume import vol1
volume: vol1
 state: ONLINE
 scrub: none requested
config:
        NAME        STATE     READ WRITE CKSUM
        vol1        ONLINE       0     0     0
          mirror    ONLINE       0     0     0
            c2t0d0  ONLINE       0     0     0
            c2t1d0  ONLINE       0     0     0
```