VTrak D5000 / Vess R3000

Data Service Guide

Contents

License Types

There are three types of license available including two classifications of trial licenses. The initial Trial License is available immediately by default. With the Extended Trial License an evaluation period of 60 days enables all the premium features. After the 60 day trial period, it is necessary to upgrade to a Perpetual License in order to continue using the premium services. The purpose of the trial license is to allow customers to evaluate the available premium services available.

Trail License

- This is used for customer to evaluate the feature before purchase the license.

- The license's life time is 60 days

- Enable all available features

Extended Trail License

- This license shall be used after Trial License was installed.

- The license's life time is 60 days

- Enable all available features

Perpetual License

Perpetual License is the license which shall not expired

Installing the License

To install acquired license, click on **Administration**, then the License menu link. Click on the **Choose File** button and locate the license file (format for license is .dat).

Click on **Upload** to begin importation of the license file. The Confirmation dialog menu appears, to continue importing the license file, you must confirm that you want to install the file by typing "Confirm" and clicking on the **Confirm** button. When the license is installed, you can go to the License menu to view the status of the license.

Please note the licensed services are not guaranteed to work normally when trial license or extended trial licensed expired. The delete operation is allowed, create operation is not.

To continue using the advanced features after the trial period, you can upgrade to a perpetual license. For more information on perpetual license and features included, please contact Promise Sales, channel partner or distributor.

License Features

The special features described in this document available with an active license include:

- Thick Clone

- Deduplication

- Compression

- Cloud Backup

- WORM

Hardware Limitations

The table below describes limitations for the available SKUs of Vess R3600 and VTrak D5000.

Feature Group	Feature	D5000 version	R3000 version	DRAM /ctrl	License
PerfectData	Snapshot/Clone/SSD Cache	FCS	FCS	>= 32G	Yes (Default included for D5000
	Local Replication/ Migration	SR1.2	N/A	>=32G	
	Thick Clone	SR1.3	SR1.0	>= 32G	
	Compression	SR1.3	N/A	>= 32G	
	De-dup	SR1.3	N/A	>= 128G	
PerfectMetro	Pair Management	SR1.3	SR1.0	>= 32G	Yes
	Remote Thickclone	SR1.3	SR1.0	>= 32G	
PerfectCloud	Cloud Backup	SR1.3	SR1.0	>= 32G	Yes
PerfectLock	WORM	SR1.3	SR1.0	>= 32G	Yes

Snapshots Overview

In order to use the Snapshot feature, a storage pool must first be created using the ***Advanced Type*** option in the **Create New Pool** menu. A volume snapshot is used to create a read-only,point-in-time copy of the volume that does not use extra storage space in the pool. The snapshot uses only space when the block references are changed. Snapshots preserve disk space by recording only the differences between the current dataset and a previous version. The snapshot is saved in case it is necessary to revert back to the volume status at the time of the snapshot for disaster recovery. This is called a rollback. A typical example use for a snapshot is to have a quick way of backing up the current state of the file system when a risky action like a software installation or a system upgrade is performed.

A snapshot of a volume cannot be directly accessed, but a clone can created for the purpose of rollback. Rolling back to a previous snapshot will discard all data changes that have occurred between the time of the snapshot and the current time. Snapshots and clones are a quick and low cost (in terms of capacity used) means of backing up a volume for the purpose of recovery.

To create a volume snapshot, click on the volume name in the left panel, click the **Snapshot & Clone** button, *you will see the Snapshot & Clone list*, then click the **Create Snapshot** button.

Create Snapshot

In the Create Snapshot menu, enter a name for the snapshot and click the **Submit** button. The snapshot will appear listed Snapshot and Clone list.

Clone options

Note that clones are a feature of the *Advanced Type* pool. This option is available when you first create a pool in the **Create New Pool** menu. Additionally, there are two methods used for creating clones. The standard clone is based on snapshots only; the snapshots are created the same pool. The 'Thick' clone uses snapshots and replication; this methods stores the snapshot and the volume or NAS share on a different pool (Destination Pool), or even on a remote system.

A standard clone is created from a snapshot as a means of backing up the snapshot. If you intend to delete a snapshot that has a clone, you must first delete the clone. Thick clone supports full backup and incremental backup, must be scheduled, and will continue with the back up until instructed to stop.

Creating a (thin) clone

To create a snapshot clone, first create the snapshot, select it in the Snapshot & Clone list, click on the **Create Clone** button.

Create Clone

In the Create Clone menu, enter a name for the clone and click the **Submit** button.

Rolling back a snapshot

To rollback using a snapshot, select the snapshot, click on the gear icon, and choose the *Rollback Snapshot* option. You are required to confirm that you want to rollback using the snapshot in a pop-up menu. Type "confirm" and click on the **Confirm** button to proceed with the rollback. Remember, any changes in the volume that have occurred since the snapshot will be lost.

Rollback Snapshot

Thick Clone

Thick Clone is used for data backup and recovery, a safeguard against data loss.
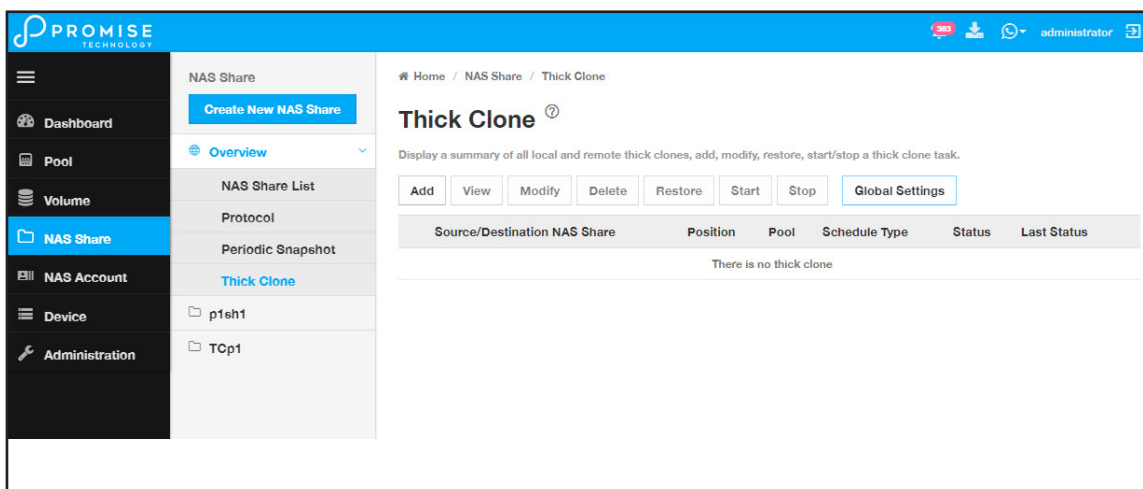
Key features of Thick Clone:

- Supports cloning for NAS Share and Volume

- Asynchronous backup and scheduled backup

- Supports Full Clone and Incremental Clone

- Supports Local Clone and Remote Clone

- Supports encrypted data transmission

- Minimal impact on ordinary IO

Adding a Thick Clone task

The Thick Clone back up feature should be used with Periodic Snapshot. Since this is backup process is asynchronous, it will take more time than the standard clone, therefore Thick Clone should be scheduled for off-peak hours, especially is there is a large volume to clone.

To setup Thick Clone, go to **NAS Share**, click on the **Thick Clone** menu link, click on the **Add** button.

Thick Clone menu

Add Thick Clone menu

In the Add menu:

- Choose **Location**, *Local* or *Remote*

- Choose **Source Pool** and **Source NAS Share**

- Choose **Destination Pool** and **Destination NAS Share**

- Choose *Manual*, *Daily* or *Weekly* **Schedule Type**

- For Weekly and Daily options, you must configure a schedule. Thick Clone can require a long time to complete, so it is best to choose a schedule during off peak hours.

Thick Clone schedule

Select **Start Time** and **Weekday** (for weekly Thick Clone).

Click on the **Submit** button to apply the Thick Clone settings.

Deduplication

Deduplication is a method used to avoid automatically writing the same data twice by detecting duplicate data blocks and keeping track of the multiple places where the same block is needed. This can save space and unnecessary IO operations which can also improve performance.

Deduplication is sychronous (happens instantly during writes, without any need for background deduplication processes), safe (no chance that two data blocks are mistakenly treated as equal) and efficient (designed to scale with the filesystem).

If stored data has a high probability there will be duplicate data block, deduplication is worthwhile and effective for both performance and storage efficiency. Typical applications where deduplication makes sense include Virtualization Storage, File Servers, Mail Servers, and Backups, especially if many users are backing up to disk.

When properly used and setup, deduplication duplicate data blocks are removed as they are written to disk. The result is that only unique data is stored on disk and common components are shared between files, as shown in the figure below.

Unique data blocks are stored on disk
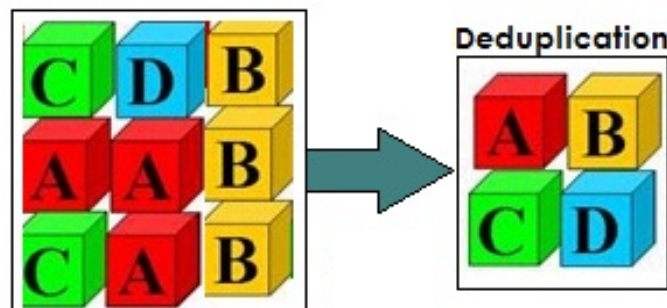
In some cases, deduplication can result in savings in disk space usage and cost. However, you must consider the memory requirements before enabling deduplication. Also, consider whether enabling compression on your file systems would provide an excellent way to reduce disk space consumption.

Deduplication is performed using checksums. If a block has the same checksum as a block that is already written to the pool, it is considered to be a duplicate and, thus, just a pointer to the already-stored block is written to disk. Therefore, the process of trying to deduplicate data that cannot be deduplicated wastes CPU resources. Deduplication is in-band. This means that deduplication occurs when you write data to disk and impacts both CPU and memory resources.

For example, if the estimated deduplication ratio is greater than 2, you might see deduplication space savings. In the example shown in Listing 1, the deduplication ratio is less than 2, so enabling dedup is not recommended.

Deduplication ratio

The deduplication ratio shows the size of archives in a deduplicating vault in relation to the size they would occupy in a non-deduplicating vault.

For example, suppose that you are backing up two files with identical content from two machines. If the size of each file is one gigabyte, then the size of the backups in a non-deduplicating vault will be approximately 2 GB, but this size will be just about 1 GB in a deduplicating vault. This gives a deduplication ratio of 2:1, or 50%.

Conversely, if the two files had different content, the backup sizes in non-deduplicating and duplicating vaults would be the same (2 GB), and the deduplication ratio would be 1:1, or 100%.

The more deduplication saves in terms of space, the more the benefits will outweigh the costs. But if your data is unique all the time, there won't be a benefit from deduplication and the cost will become more prevalent in terms of performance.

Memory requirements

For deduplication about 1GB of RAM is required for each TB of data. For Pools the rule is about 5 GB of Ram for each TB of data.

For testing purposes, switch deduplication on, then test for a week or so, and the deduplication ratio will tell you how much space you saved, while you'll be able to observe how performance is impacted. You can determine if deduplication is worth it, if not, you can simply switch it off again.

Deduplication Settings

Deduplication controls whether duplicate copies of data are eliminated. Deduplication is synchronous, pool-wide, block-based, and can be enabled or disabled to experiment and observe if it is worthwhile.

Deduplication and Compression is enabled when creating a volume.
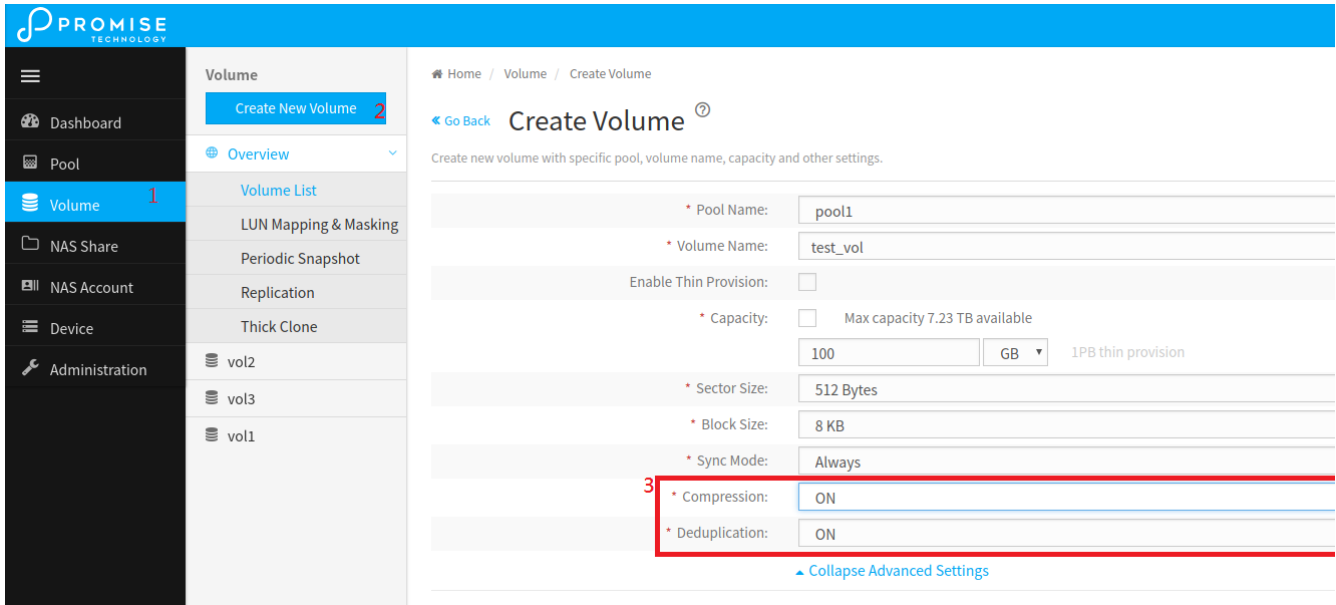
Compression Settings

Data compression is another feature used to improve storage efficiency and possibly performance. Compression is applied only to the new and modified data. The degree of compression varies depending of which algorithm is selected. There is also a performance cost that varies with algorithm, generally more compression will use more CPU resources. Like deduplication, compression is synchronous, i.e. it happens when new or modified data is written to disk.

Deduplication and Compression is enabled when creating a volume.

WORM

The Write Once Read Many (WORM) feature is used to store data that is not allowed to get modified, a method of immutable data storage. WORM is well suited for certain types of data, such as log files, financial records or medical records where data integrity is essential. Government, medical, financial institutions, IT departments often maintain records that must remain unchanged. The WORM feature is applied to a NAS Share and can be applied indefinitely or for a specified period.

To setup WORM on a NAS Share, choose a NAS Share, expand the **More** menu and choose *WORM*.

NAS Share List
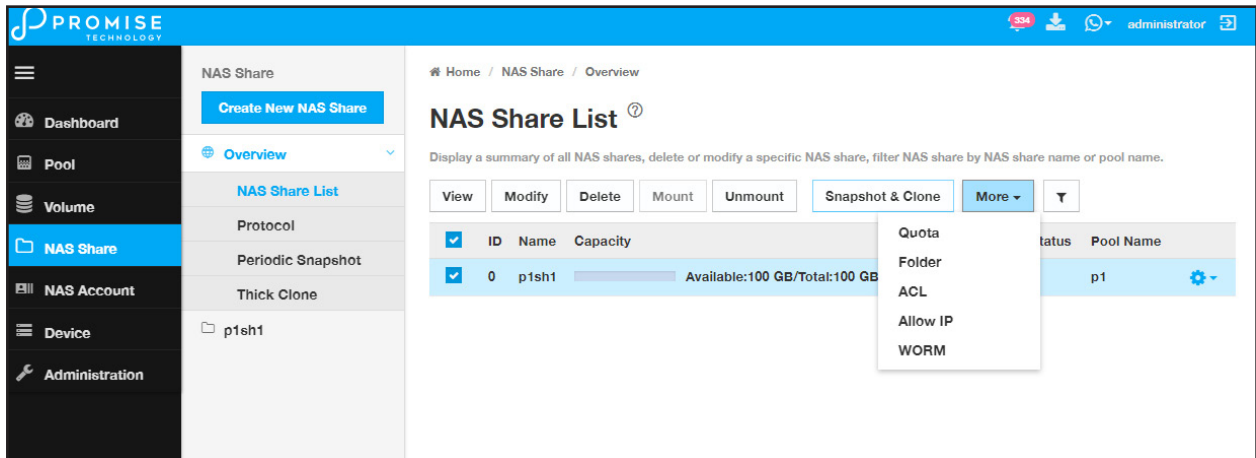
WORM settings
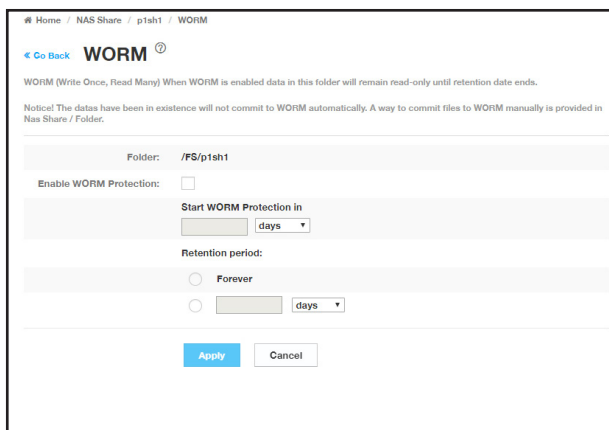
In the new menu, click on the *Enable WORM Protection* to enable it, and configure settings for start time and duration of the protection.

Click on **Apply** to make the settings active.

Note that the WORM cannot be disabled once it is enabled.

Cloud Backup

The Cloud Backup enables syncing NAS user files shared between the Vess R3000 or VTrak D5000 and one of the supported cloud services.

Supported cloud services:

- Aliyun
- Amazon S3

- Azure Blob
- Dropbox
- Google Drive

- OneDrive
- OpenStack Swift

Requirements for use

In addition to a valid connection to the Internet, the following are required to use Cloud Backup:

- A cloud drive with one of the supported vendors

- DNS settings correctly configured

- NAS user account

- System time consistency (consistent with world time, NTP server, etc.)

To setup Cloud Backup, the NAS user must first login,then click on Cloud Backup to view the task list. Click on **Add New Task**.

Cloud Backup task list
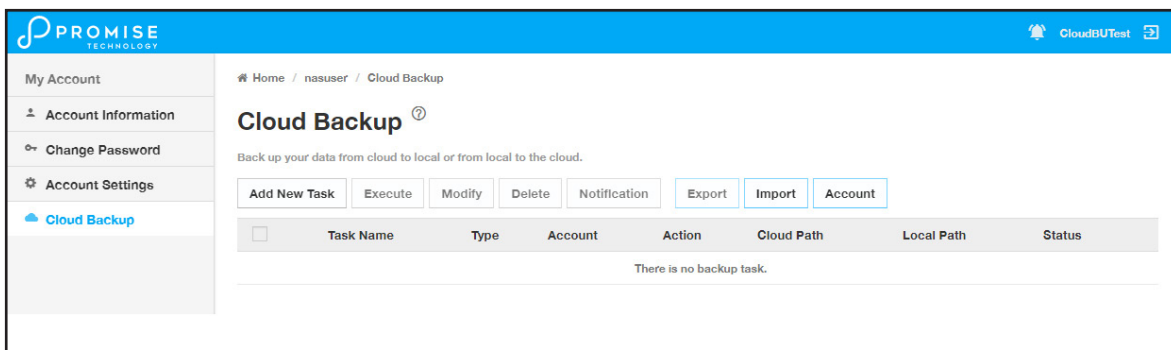
Add New Task menu

In the Add New Task menu, select the **Cloud Type** (Cloud Service) to be used for backup and click on the **Add New Account** button. A login menu for the selected cloud service appears. Choose the account to use for this task. The login takes you to a permission menu where it is necessary to first allow permission to accept backups from the Vess R3600 or VTrak D5000.
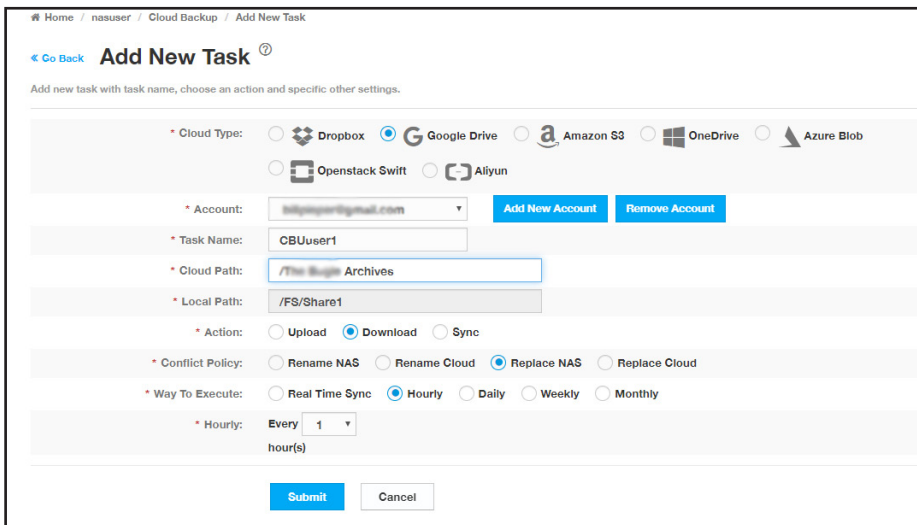
Select cloud account menu (Google Drive)

Allow access permission menu (Google Drive)
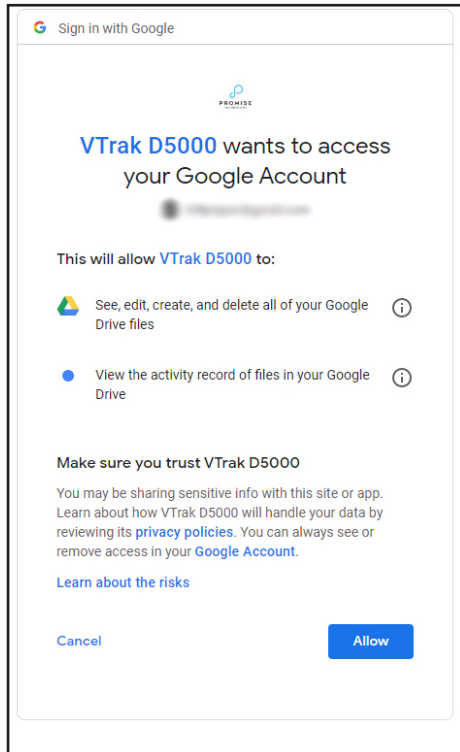
The menu that appears depends on the service used. In this example, a Google Drive account is used for the task. Follow the instructions for the service to allow permission (i.e. click **Allow** or **OK** button). You will return to the task settings menu.

Add New Task menu with account added

Enter settings for **Task Name**, **Cloud Path** (choose folder on cloud service storage), **Local Path** (choose location on system), **Action** (*Upload*, *Download, Sync*), **Conflict Policy** (action taken if there is a data conflict such as a file with the same name or path) and configure the schedule for the task, or choose Real Time Sync to update continuously.