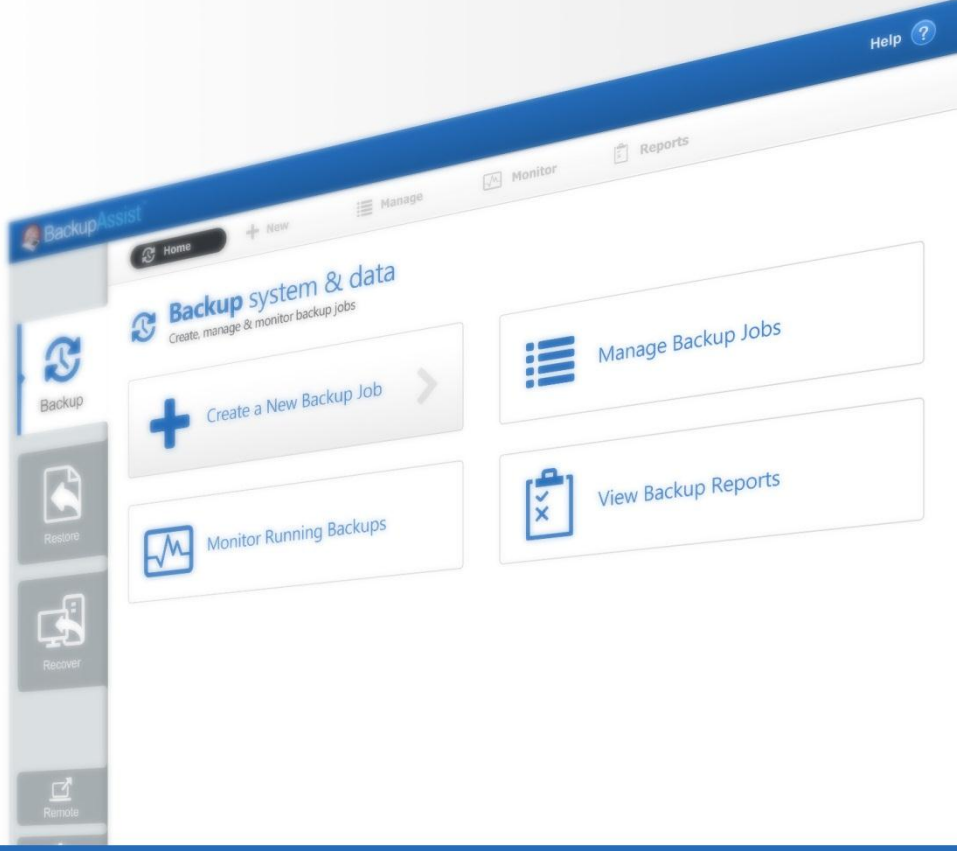

BackupAssist™

BackupAssist v7

File Protection – Using Rsync

User Guide

BackupAssist User Guides explain how to create and modify backup jobs, create backups and perform restores. These steps are explained in more detail in a guide's respective whitepaper.

Whitepapers should be used as the main reference documents when planning your backups and your data protection strategy. Whitepapers include important considerations, configuration explanations and the implementation information needed to use BackupAssist effectively.

Cortex I.T.

Contents

1. Overview

BackupAssist File Protection includes a powerful tool called Rsync that can back up data across the internet to any Rsync host. Critical files can be copied to a secure, offsite location, away from your office, and backing up across the internet overcomes the need to swap tapes or hard drives.

Once you have selected the host where your data will be stored, no further equipment or maintenance is required. Additional storage space can be easily added to the data host as your data requirements grow, and your data is available whenever you need it using BackupAssist.

Documentation

More information on Rsync can be found in the File Protection using Rsync whitepaper.

The whitepaper contains comprehensive information and should be referred to when planning backup strategy using Rsync.

Other BackupAssist documentation can be accessed through the Documentation webpage

Recommended reading:

- To learn more about the BackupAssist Recover tab, see the BackupAssist Recover Tab User Guide
- To learn more about the BackupAssist Settings tab, see the BackupAssist Settings Tab User Guide

Licensing

To back up data across the internet with Rsync, requires the *Rsync Add-on* license, once the initial trial period has expired. Please contact your local BackupAssist reseller or distributor for pricing information, or visit www.BackupAssist.com.

For instructions on how to activate / deactivate license keys, visit our Licensing BackupAssist page.

Backup settings

BackupAssist's settings can be entered and modified using the selections available in the **Settings tab.** Clicking on the *Settings tab* will display the selections as icons.

Backup user identity

Backup jobs require an administrator account with read access to the data source, and full read-write access to the backup's destination. It is recommended that a dedicated backup account is created for this purpose. The account's details are entered using the *Backup user identity* option on the *Settings* tab, and your backup jobs will be launched using these credentials.

Email server settings

This menu item is used to enter the details of the SMTP server used by BackupAssist to send email notifications. The SMTP server must be configured if you want to have an email *Notifications* step enabled when you create a backup job.

System State

Rsync cannot be used to create a System State backup. A System State backup is explained here.

2. Rsync Considerations

As Rsync is an open protocol, you have the option of either storing your data on a third party destination server, or supporting a destination Rsync server yourself.

For more information on how to get the most out of Rsync, visit our Video Presentations page.

Third Party data host

Third party data centers, ISPs and cloud providers can support Rsync backup destinations. These solutions have the advantage of high availability networks, and some datacenters also offer geo-redundant storage.

Third Party data host: setting up S3Rsync

BackupAssist includes a dedicated configuration screen for backups to Amazon S3 via the s3rsync.com service. To backup to Amazon S3 with Rsync you will need:

- **Amazon account**
  In your Amazon Web Services account, you will need to obtain your Access Key ID and generate a Secret Access Key. Then you will need to create an S3 bucket to use for your backups. See this article for a guide to the Amazon S3 Simple Storage Service.
- **S3Rsync account**
  When you sign up for an s3rsync.com account, you will be given a username and a private SSH key file. Save the SSH key file somewhere on the machine on which you wish to run BackupAssist.
- **BackupAssist**.
  Once you have performed these steps, you can set up your job in BackupAssist using the S3Rsync *Destination* selection.

Do-it-yourself host

Any Rsync Server such as an Rsync-enabled NAS device, Windows or Unix machine can be used to store backups using Rsync. The do-it-yourself approach has the advantage of keeping data in your control, and a lack of monthly hosting fees or limits to the amount of data backed up.

To set up your own Rsync host, please see the File Protection with Rsync Whitepaper.

Rsync as a BackupAssist add-on

Rsync will be shown as a destination option when you create a File Protection backup, but it requires the *Rsync Add-on* license (once the trial period has expired) or the Rsync standalone license (for offsite-only backups). It is important to maintain your Rsync license because your Rsync backups will not be accessible if your license has expired.

When your backup job is first set up, you should *seed* your data on the data host. Use removable media to physically transport the data, or if you are using a NAS host, run the job once over a local network. For more information, see the Rsync backup management section of this whitepaper.

3. Backup considerations

Before creating a backup job, it is important to understand what backup and restore options are available. This section provides guidance on some key considerations.

Exchange VM Detection

If your backup job contains a Hyper-V guest with an Exchange Server, the authentication information for that guest should be entered into the **Exchange VM Detection** tab on the **Selection** screen when you create the backup job. With these credentials, BackupAssist can detect what guests have an Exchange Server, and list the EDB file available for each guest when you perform a restore using the Exchange Granular Restore console

The Exchange VM Detection tab will appear when the Hyper-V role is installed and running on the server. If you are backing up multiple Exchange guests, each one should have the same username and password.

The Hyper-V process is automated but the restore requires both the *Exchange Granular Restore Add-on and the Hyper-V Granular Restore Add-on* licenses.

VSS Application backups

The Volume Shadow Copy Service (VSS) is a Microsoft Windows Service that creates a copy of an application's data so the data can be backed up while the application is running. This means the data will not change or be locked while a backup is taking place. BackupAssist is VSS-aware, so File Protection, File Archiving and System Protection backups can detect VSS applications such as Exchange, SQL, Hyper-V and SharePoint. BackupAssist will display a VSS application as an application container during the *Destination* step of the backup job's creation. You can select the container or individual components and BackupAssist will select the files that need to be backed up.

In some cases, only applications that are running will be detected. If an application is not listed, try re-starting the application and the VSS service and then click the *Refresh* button in BackupAssist.

For Windows Small Business Server 2003, a registry entry modification is needed to see an Exchange Server. See our online blog post, Backing up Exchange with SBS 2003, for more information.

Restore vs. Recovery

A restore is the process of accessing a backup and restoring it to the original (or a new) location, if your data is lost, corrupted or if you want an earlier version of that data. However, if your computer cannot start itself, you may need to perform a recovery.

A recovery is the process by which a computer is recovered after hardware has been replaced or an operating system failure has occurred, and your computer can no longer start itself. To perform a recovery you need a bootable media to start your computer, and an image backup that the bootable media can use to recover your operating system, data and applications.

For more information on data recovery, see the Recover tab & RecoverAssist Whitepaper.

4. Creating a File Protection backup using Rsync

The following instructions describe how to create a File Protection backup job using Rsync. The configurations include Rsync host connections for both windows and Linux systems.

Launch BackupAssist and follow the steps outlined below:

1. Select the **Backup** tab, and click **Create a new backup Job**
2. Select **File Protection**

   If this is the first time you have created a backup job, you will be asked to provide a *Backup user identity* if one has not been defined. See the BackupAssist settings, whitepaper for guidance.

3. **Selections:** The selections screen is used to select the data and applications that you would like to back up. Any VSS applications detected will be displayed here as application directory containers.

   An Exchange VM Detection tab will be available if you are backing up an Exchange VM guest.

   Select the volumes, folders, files and applications that you want to back up, and click **Next.**

4. **Destination media**: The destination screen is used to select the type of media that you want to back your data up to. This step's name will change to "Rsync", when you click next.

   Select **Rsync** or **S3Rsync** for your backup destination, and click **Next.**

   The *S3Rsync* option is for backups to Amazon S3 via the s3rsync.com service.

   Select **Enable Rsync file based encryption** if you want the backup data to be encrypted before being transmitted.

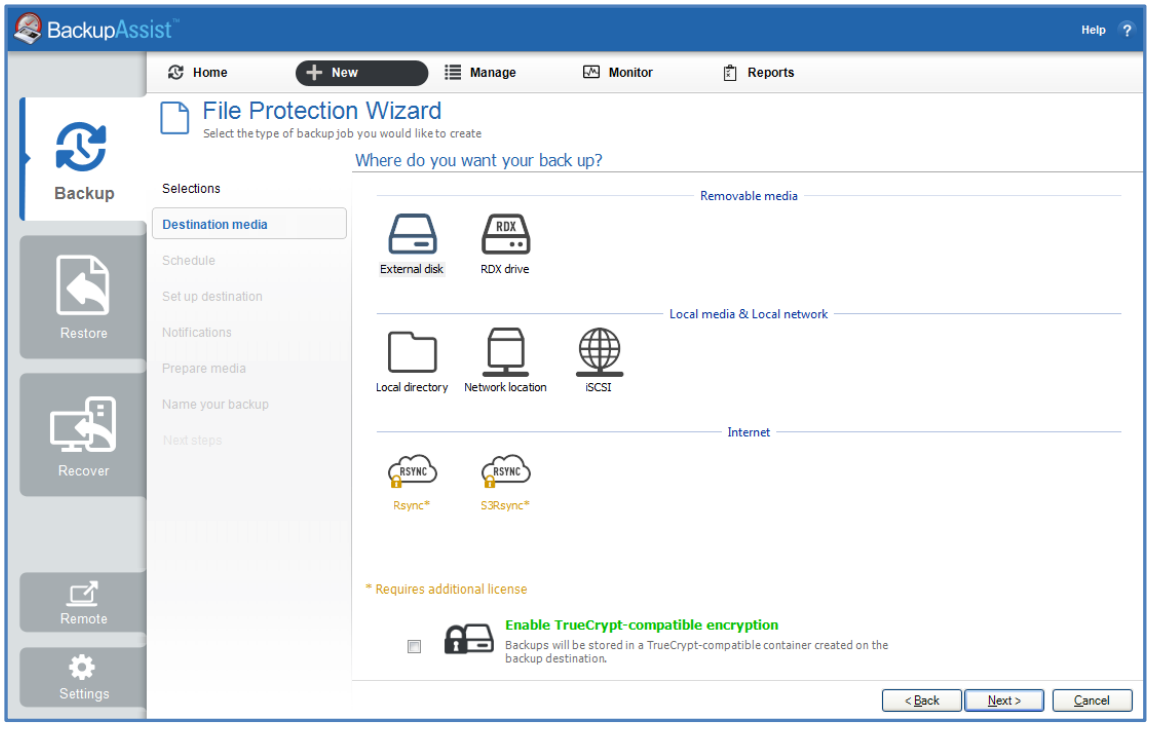

Figure 1: File Protection backup – Rsync destination selection

5. **Schedule:** This screen is used to select when and how you would like the backup job to run, and how long you would like the backup to be retained for. A selection of pre-configured schedules, called schemes, will be displayed.

   - The schemes available will depend on the type of destination media selected in step 4.
   - Clicking on a scheme will display information about the schedule used.
   - The schedule can be customized after the backup job has been created.

   For more information about creating custom schedules, refer to the Backup tab whitepaper.

6. **Set up destination**. The screen is used to configure your Rsync destination. The configuration screen displayed will depend on whether *Rsync* or *S3Rsync* was selected.

   IF the standard **Rsync** *Destination* was selected, follow the guidelines below:

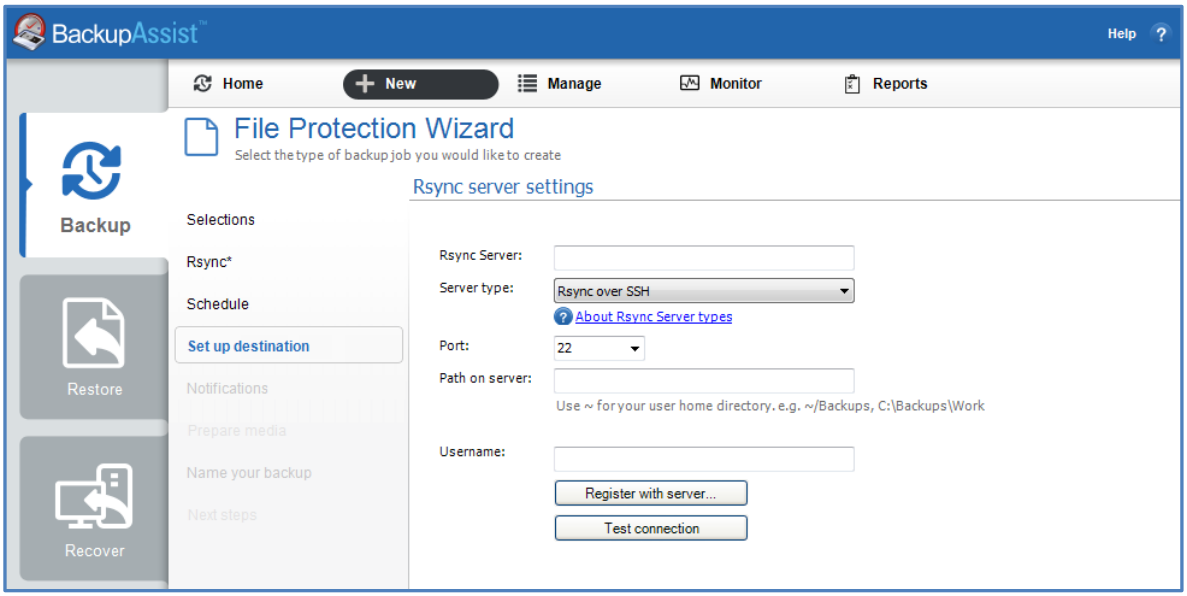

Figure 2: BackAssist File Protection – Rsync destination selection

a. **Rsync Server:** Enter your Rsync server name (or IP address).

b. **Server Type:** Select *Rsync over SSH*, *Rsync Daemon* or *Rsync Daemon over SSH tunnel*.

c. **Port:** The default port will display for the server type selected.

d. **Path on server:** It is best to use a new, empty directory for this path. The parent directory must exist although the sub directories will be created when the job is first run: /parent/sub_directory/.

   - If your data host is running **Windows**, you can enter a normal Windows path here, such as "C:\Backups". You can also enter a path relative to the user's home directory by starting with a tilde (e.g. "~/Backups").
   - If your data host is running **Linux**, you can use an absolute path by starting with a slash or a path relative to the user's home directory by starting with a tilde (e.g. "~/Backups").

e. **Username:** Enter the username that was activated while setting up your Rsync host.

f. **Register with server:** Select this option and you will be prompted to enter the password. BackupAssist will then create a public/private key pair to authenticate you to the data host.

g. **Test connection:** Click to test your connection to the Rsync server. If this step fails but the registration succeeded, it is probably that the *Path on server* cannot be accessed.

IF the **S3Rsync** *Destination* was selected, follow the guidelines below:

Once you have set up your **Rsync destination**, click **Next**

a. **Rsync Server**: This should be farm.s3rsync.com (the default setting) unless you have been advised otherwise by s3rsync.com.

b. **Port**: This should be 22.

c. **Amazon S3 bucket**: You can leave this blank unless you want to set up multiple backup jobs using the same bucket (not recommended).

d. **Set Path**: Specify any folders you have created in the bucket.

e. **Access Key ID**: Your S3 Access Key ID.

f. **Secret Access Key**: Your S3 Secret Access Key.

g. **S3rsync username**: Your username supplied by s3rsync.com (note: this is different to your Amazon username).

h. **S3Rsync SSH key path**: The location of the saved SSH key file provided by S3rsync.com.

i. If you selected *Enable Rsync file based encryption*, you will be prompted to create a password.

> **Note:** It is important that you keep a copy of your password in a safe place, as we cannot retrieve passwords if they are lost or forgotten.

➢ **Mail Server**: If you have not configured an SMTP mail server for BackupAssist, you will be prompted to provide those details after the backup destination step has been completed.

7. **Notifications**: Once a backup job has completed, BackupAssist can send an email to inform selected recipients of the result. This email notification can be enabled during the creation of a backup job, if the mail server has been configured. To enable email notifications:

   a. Select, **Add an email report notification.**

   b. Enter recipients into the **Send reports to this email address** field.

   c. Enter recipients into the **Also send reports to this email address** field. You can then select the condition under which the email should be sent, using the drop-down box.

   After the backup job has been created, you can modify the notifications by adding and removing recipients, setting additional notification conditions and including print and file notification types.

8. **Prepare media:** This step will be skipped because Rsync backups do not use removable media.

9. **Name your backup:** Provide a name for your backup. Click **Finish**.

▶ **The File Protection with Rsync backup job has now been created.**

Important: Once a backup job has been created, it should be reviewed and run using the *Manage* menu. This menu provides additional options to configure your backup. See the section, File Protection using Rsync backup management, for more information.

Important: Once a backup job has been run and a backup created, a MANUAL test restore should be performed to ensure the backup is working as intended. To perform a test restore, refer to the section, Restoring from a File Protection backup.

5. Restoring from a File Protection backup

This section provides instructions on how to restore data that was backed up using BackupAssist's File Protection using Rsync.

To restore data from a **File Protection** backup, start BackupAssist and follow these steps:

1. Select the **Restore tab**

   The *Restore tab* has a *Home page* and a *Tools menu*. The *Home page* is the default screen and the recommended starting point for performing a restore. The *Tools menu* should only be used by experienced administrators or users being assisted by technical support.

2. From the **Home page**, select the type of restore you want to perform. When you select one of the restore categories provided, BackupAssist will locate the corresponding backups for you.

   - *Files and folders* will display all data backups and all VSS application backups.
   - *Applications* will display backups that contain VSS applications, and exclude data only backups.
   - *Exchange*, *SQL* or *Hyper-V*, will display all backups that contain the selected application. Selecting an application type will display application specific restore tools (e.g. Hyper-V Granular Restore and SQL Restore) as well as the Restore Console.

3. Once you have selected the type of restore you want to perform, the *Home page* will display all backups catalogued by BackupAssist that match your selection. The backups will be grouped by the backup's source location, and by the restore tool that can be used.

   - If a backup can be used by two restore tools, it will appear in two groupings.
   - If a backup contains data from multiple locations, it will appear in a grouping for each location.

   If your backup included both data and VSS applications, both will be available to restore once the backup has been loaded in step 4, regardless of the restore type selected.

   Select the **Restore Console.**

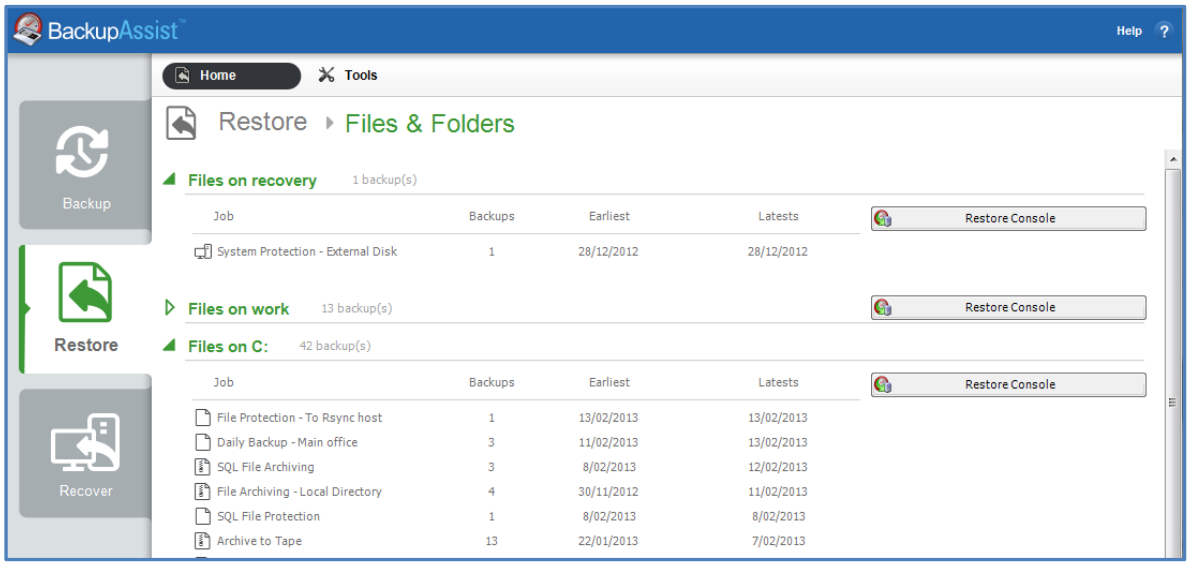

Figure 3: BackupAssist Restore Home page – selection results

4. **Restore Console – backup and data selection**

   The BackupAssist *Restore Console* will open and load all of the backups that were listed on the *Home page*. The next step is to locate the data you want to restore, from the loaded backups. The Restore Console provides two tools to locate your data:

   - The **Browse** tab. Select this tab if you know the backup and date you wish to restore from, or if you need to restore an entire backup set.

     a. Use the drop-down menu to choose the backup that you want to restore from.
     b. Use the calendar to select the date you want to restore from.
     c. Use the middle panes to expand the backup set.
     d. Select the data to restore.
     e. Click **Restore to** at the bottom right of the window.

   - The **Search** tab. Select this tab to search all of the loaded backups for the data you want to restore. You can display data filtered by name, date, size and type, for all backups. The results can be compared (e.g. the dates of two files) to identify the correct data selection.

     a. Enter your search term (The search accepts wild card searches, such as **.log* or **.doc*).
     b. Select a filter/s if required.
     c. Click the *Search* button.
     d. Select the data to restore.
     e. Click **Restore to** at the bottom right of the window.

   If the backup is not present, or if you wish to load additional backups, select the **Load backups** option. Click **Load all known backups** to load all backup catalogues.

5. **Restore Console – restore destination selection**

   When you select *Restore to,* a window will open showing the *Backup location,* the *Restore to* destination and the *Restore options*.

   a. Review **Backup location**.

   b. Review **Restore to:** Leave the *Original location* selected or chose an *Alternative path*.

      Restoring to an alternate location will use a minimal path. For example, restoring a single file to an alternate location will copy the file to the location without re-creating the original folder structure.

   c. Review the **Restore options:**

      - Select one of the following: *Overwrite all existing files*, *Do not overwrite existing files* or *Only overwrite older files*.
      - The option, *Restore NTFS security attributes* will be selected by default.

   d. Selecting *Create a log file listing all processed files*, will create a file that lists the success or failure of each file. The log is opened by selecting the log file's link in the backup report.

   e. *Queue all backup jobs when a restore is running*, is selected by default.

   f. Click the **Restore** button.

      - The Restore Console will connect to your Rsync host and restore the selected files.
      - The restore will run from the destination window and a **Report** link will appear once the restore has finished.

   g. Select **Done**.

▶ **Your File Protection using Rsync restore has now been completed.**

6. File Protection using Rsync backup management

Once you have created a backup job, you can modify the settings and access advanced configuration options using the *Manage* menu..

To access the backup management screen:

1. Select the BackupAssist, **Backup tab.**
2. Select **Manage** from the top menu. A list of all backup jobs will be displayed.
3. Select the backup job you want to modify, and select **Edit.**
4. Select the required configuration item on the left. Key configurations are described below.

To learn more about the backup management options, see the Backup tab whitepaper.

Manually running a backup job

All new and modified backup jobs should be manually run to ensure they work as intended.

1. Select the backup job, and select *Run*.
2. You will be prompted to *Rerun a past backup* or to *Run a future backup now*.
3. When the backup job starts, the screen will change to the *Monitor* view.
4. Once the backup has been completed, select the *Report* button and review the results.

Rsync Data Seeding

Rsync backups are incremental backups. The first time you perform your backup, no data will exist on your data host so a full backup is required. If you enable or disable encryption for an Rsync job, BackupAssist will need to *re-seed* the backup to the Rsync backup destination with a full set of data.

Seeding your backup via a slow internet connection may not be practical, so two methods are provided here to seed your data host. Once the initial seed to the data host is complete, each successive backup will be an incremental backup of data that has changed.

Option 1 – Seeding a permanently offsite data host

You can use BackupAssist to automatically seed data offsite using a removable media, which can be physically transported to the data host so the data uploaded locally. Seeding your data using this method is simple:

1. Connect a removable media device to the machine running BackupAssist.
2. Select the **Destination** menu item.
3. Click the **Seed backup** button and select the location of an empty folder on your portable media.

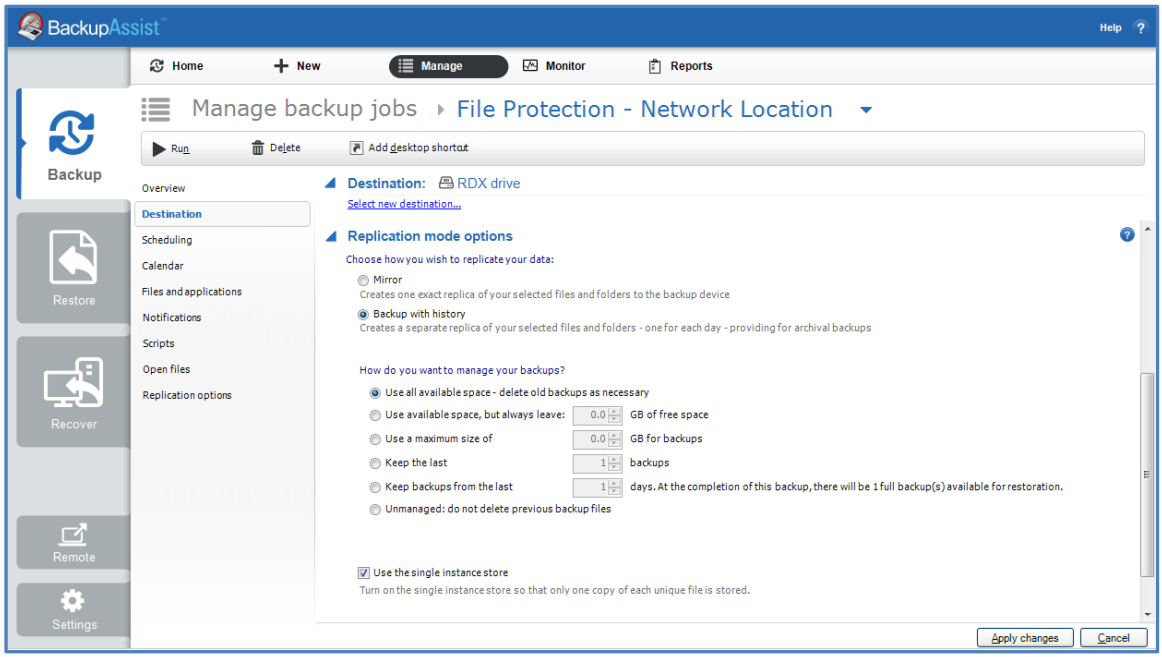

Figure 4: Manage backup jobs screen – Destination seeding

4. Once the seed is complete, your portable media should contain:
   - A **directory** with the seeded data
   - A **README.txt** file containing instructions on how to copy the seed to your Rsync server
   - An .**sh script file**, which is used to copy your data to your Rsync server.
5. Transport the portable media containing the seed to the site where your Rsync server is located.
6. Connect the device to the Rsync server and copy the seed to it:

   For a Windows server (assuming the seed is located on E:\SeedFolder)

   a. Go to the *Start* menu > CopSSH > Start a Unix BASH shell.

   b. Enter the following command: bash "/cygdrive/e/SeedFolder/seed.sh".

   For a Linux or Unix server (assuming the seed is located in /mnt/usbdrive/SeedFolder).

   a. Run the following command in your shell: bash "/mnt/usbdrive/SeedFolder/seed.sh".

   A complete seed of your data should now be copied to your Rsync server. Each successive backup from now on will be an in-file delta incremental backup of data that has changed.

Option 2 – Bringing your data host onsite to perform the seed

This method is suitable for "standalone" data hosts (where a data host is not shared among multiple clients) that can be physically transported onsite – such as NAS devices.

Seeding your data is easy – simply follow these instructions:

1. Connect your data host to the LAN and make a note of its IP address / Hostname.
2. Create your BackupAssist Rsync job, run it at convenient time and wait for it to complete.
3. Move your NAS to its permanent location.
4. Update the job settings in BackupAssist to reflect the new IP address / Hostname.