BackupAssist™ v9

File Protection using rsync

User guide

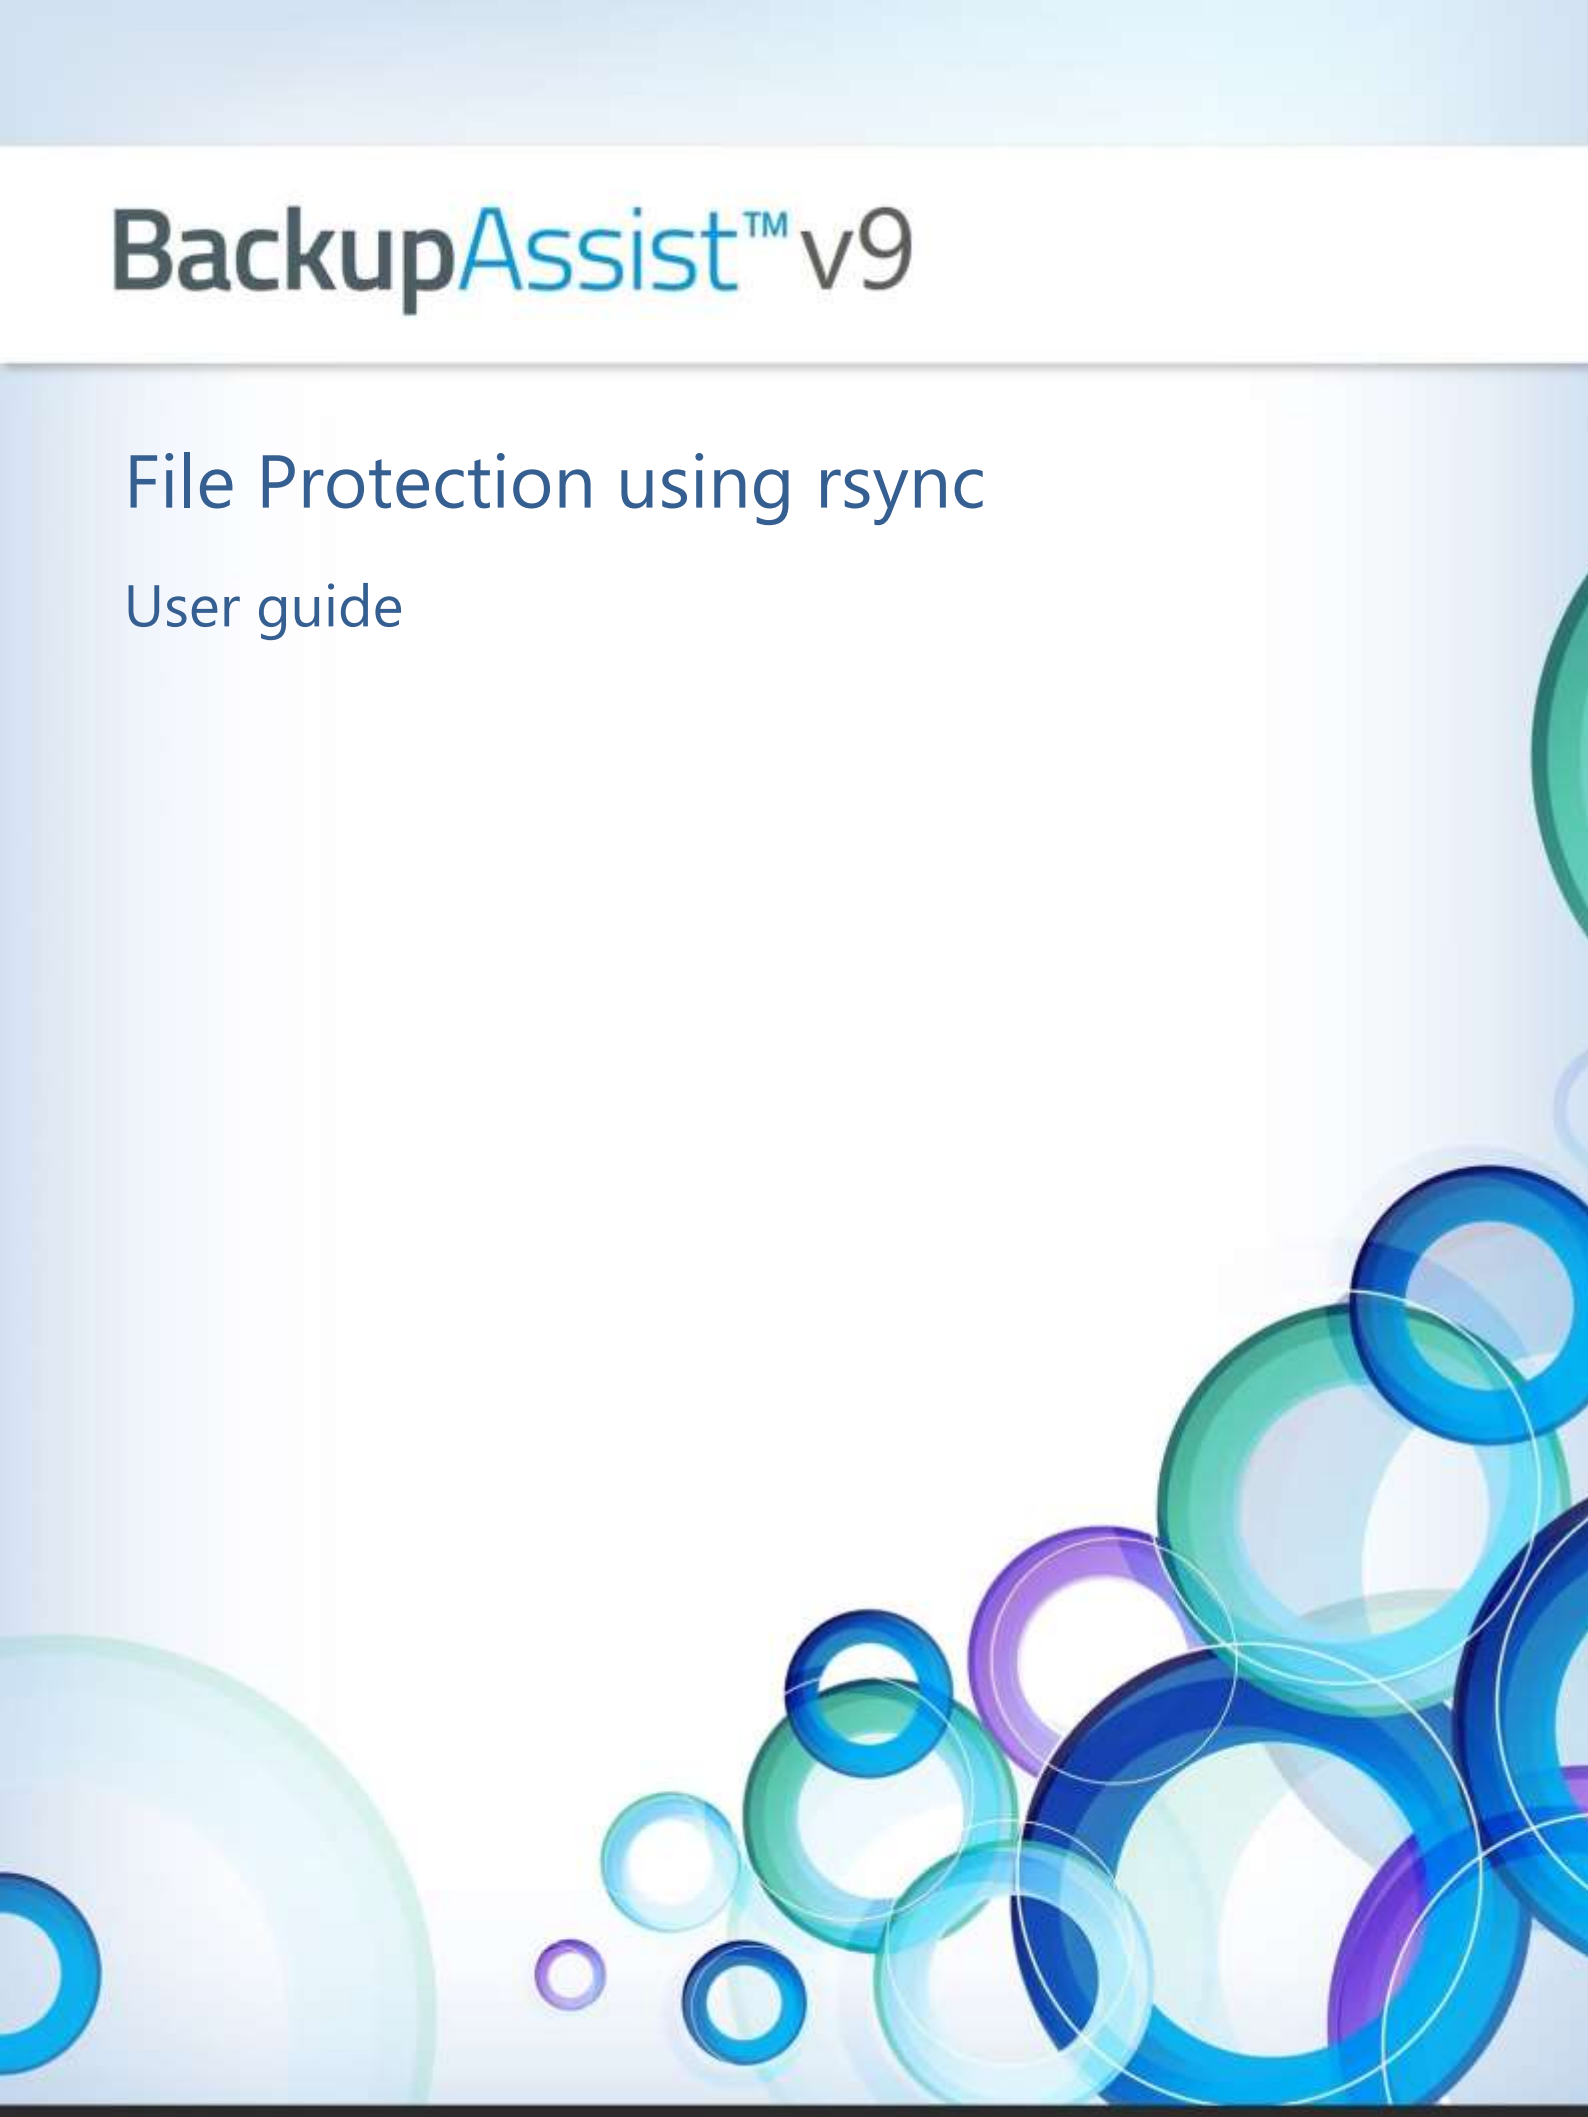

Contents

1. Introduction

BackupAssist File Protection includes a powerful tool called Rsync that can back up data across the internet to any rsync host. This guide outlines how to use Rsync to protect your data.

Adding rsync backups to your backup strategy is an excellent way of insure yourself against data loss. 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've 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, so you don't have to worry about purchasing replacement hardware. Best of all, your critical files are available whenever you need them and can be accessed from wherever you are, using BackupAssist.

Documentation

This guide explains how to create backups and perform restores using an rsync destination with BackupAssist File Protection. To learn how to set up your rsync destination, see our rsync setup guide.

File Protection– Rsync setup guide.

BackupAssist – Backup tab user guide

Rsync how-to video - Video Presentations page.

Licensing

File Protection is a standard feature included with the BackupAssist license. To back up data across the internet with rsync, requires the *Offsite Backups 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 out Licensing BackupAssist page.

Overview

Rsync is an open source application used to synchronize files and directories from one location to another. BackupAssist's implementation of this technology is in the form of an rsync destination option for File Protection backups, which allows you to back up data across the internet. The data transfer is minimized because only the data that has changed is transmitted and all data packets are compressed. You can also use built-in rsync encryption to protect the data on the rsync host.

The rsync destination that you use can be either an rsync server that you maintain yourself or a third party destination that supports rsync. Third party destinations include data centers, ISPs and cloud providers. These solutions have the advantage of high availability networks with saleable storage.

BackupAssist includes a dedicated configuration screen for backups to Amazon S3 via the s3rsync (www.s3rsync.com) service. To backup to Amazon S3 with rsync, you will need both an Amazon account and an S3Rsync account.

2. Rsync considerations

The performance and flexibility of backing up across the internet can depend on how rsync is implemented. Below are some key considerations when planning your rsync backup solution.

VSS applications

VSS applications including Exchange, SQL and Hyper-V, can be backed up to an rsync destination using File Protection. For Hyper-V however, we recommended System Protection backups, which do not support the rsync destination but do support granular Hyper-V gest restores.

Synchronizing drive images using rsync

Rsync is a destination for File Protection backups. It is possible for the data source to be a System Protection image backup, but this solution is not recommended because significant performance issues that can arise. If you want to back up important files to an rsync host, the best way is to back up those files using File Protection directly. Continue to create your image backups, but back up the important files independently using a File Protection rsync backup job.

We also advise against using File Protection's rsync, to transfer File Archiving backups to an rsync host. This is because rsync uses a checksum method to perform the bit level data transfer. Rsync checks whether any data has changed by looking at the file size and modification date. This is fast and simple on a regular file system, but if you have a very large single archive file (>10 GB) it will take much longer to complete, even if only a small element has changed.

Seeding

Rsync backups are incremental backups. The first time you perform your backup, no data will exist on the data host so a full backup will be required. Seeding your backup via an internet connection may not be practical, so two methods are provided to seed your data host. These are explained in the Rsync backup management section of this document.

Single-Instance store

File Protection backups cannot use single-instance store when the backup is saved on a ReFS formatted rsync destination. This means all the data will be backed up each time the backup job runs.

Backup source & frequency

Run your rsync job regularly. Regular daily backups will ensure that you keep your data transfer to a minimum and your data up-to-date.

Simultaneous backups

If you have a large number of backup jobs sending data to a host at the same time, the connections may become unreliable. It is recommended that you limit the host connections 5 at a time. Depending on storage requirements and the bandwidth available, you may increase this number with caution.

Backup user accounts

Rsync backup jobs require an BackupAssist administrator account with read access to the data source. This is set up using the *Backup user identity*, option in the *Settings* tab. The backup job will also need an rsync host account with read-write access to the rsync destination. This is enabled on the host server, and enter in the rsync destination screen.

3. Creating a File Protection backup using rsync

The following instructions describe how to create a File Protection backup job to back up your data to an rsync host.

Launch BackupAssist and follow the steps outlined below:

1. Select the **Backup** tab, and click **Create a new backup Job**
2. Select **File Protection**, and enter the details in the screen provided.

   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.
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.

   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 that use both Amazon S3 and the www.s3rsync.com service.

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

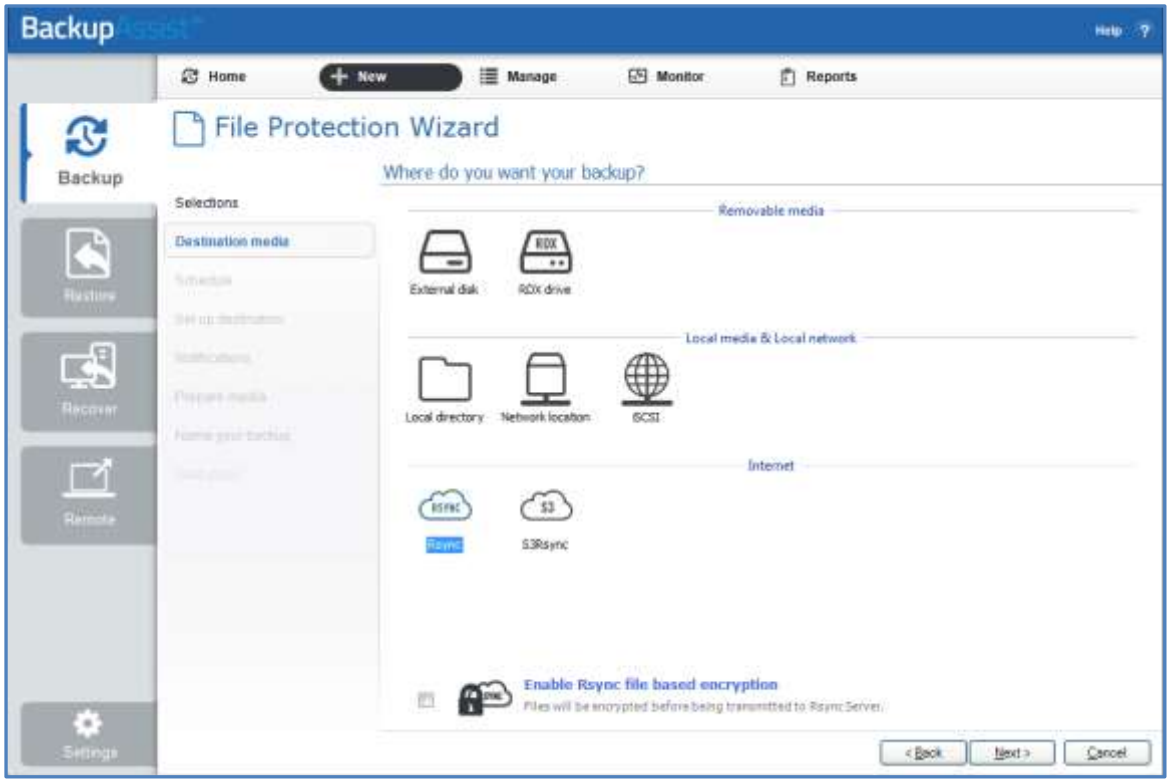

Figure 1: File Protection backup – Rsync destination selection

5. **Schedule**

This screen is used to select when the backup job is to run, and what the mix of daily and archive backups will be. A set of pre-configured schedules, called schemes, will be displayed.

Select an appropriate scheme and click **Next**.

- 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.

To learn more about File Protection schedules, refer to the Backup management section below.

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:

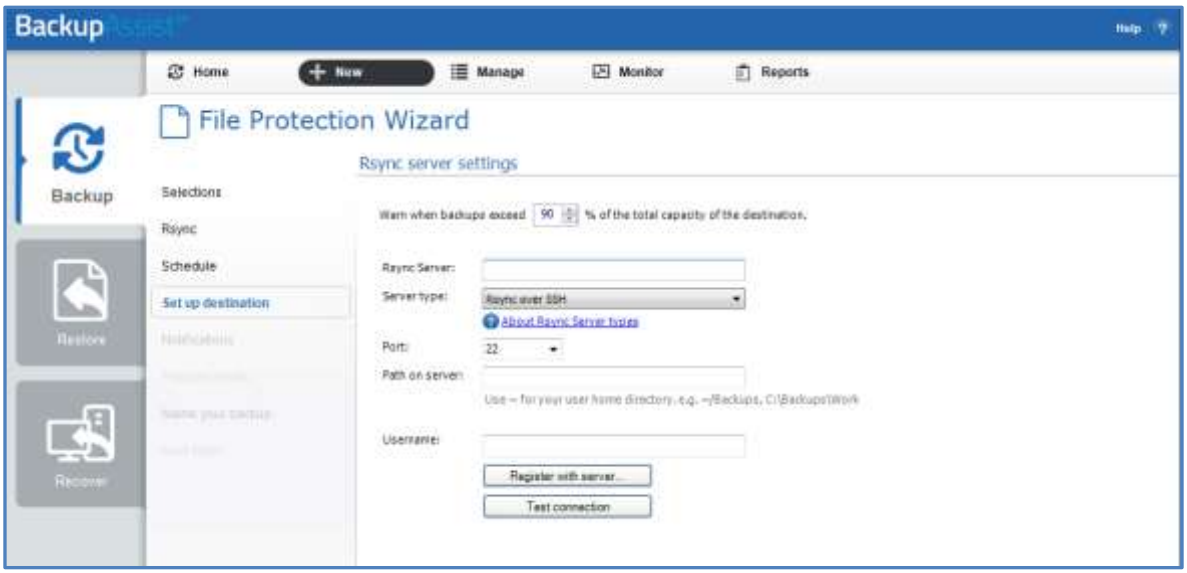

Figure 2: BackAssist File Protection – Rsync destination selection

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. The sub directories will be created when the job is first run: /parent/sub_directory/.

   - If your host is running *Windows*, enter a normal Windows path here, such as C:\Backups. Alternatively, enter a path relative to the user's home directory using a tilde (~/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 when the rsync host was set up.

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

g. **Test connection:** Use this button to test your connection to the rsync server. If this step fails but registration succeeded, the problem may be that the *Path on server* cannot be accessed.

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

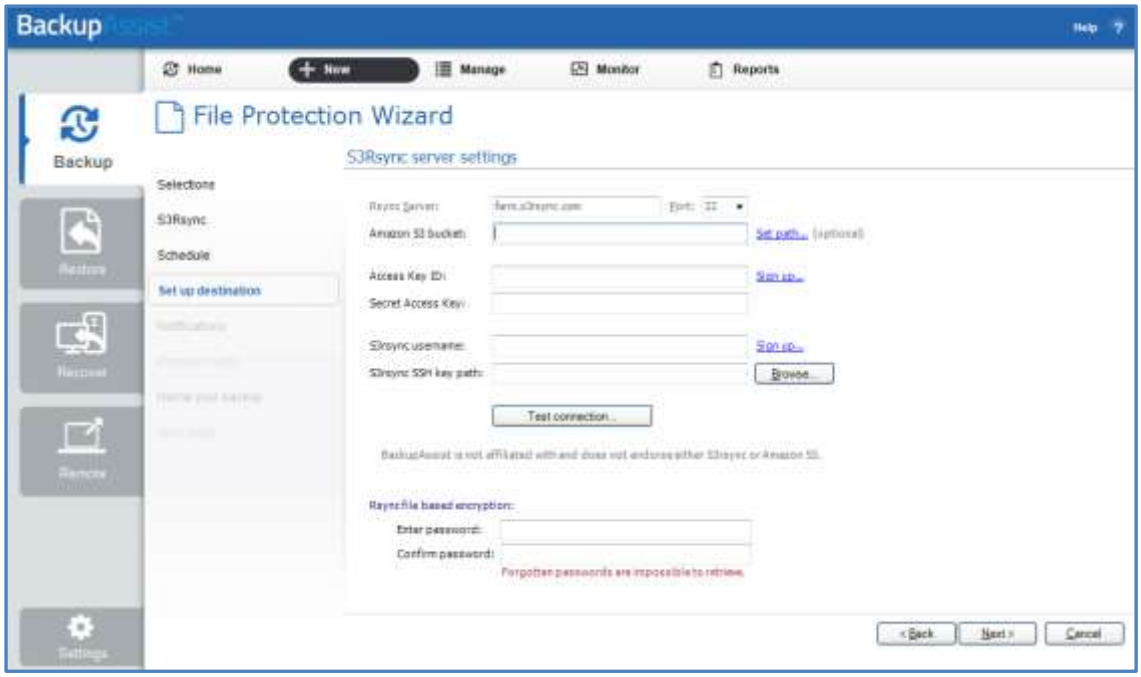

Figure 3: BackAssist File Protection – S3Rsync destination selection

The information below is provided when you create:

- An Amazon S3 account
- An S3rsync service account.

a. **Rsync Server**: This should be farm.s3rsync.com (the default setting) unless you have been advised otherwise by s3rsync (www.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**: Enter your S3 Access Key ID.

f. **Secret Access Key**: Enter your S3 Secret Access Key.

g. **S3rsync username**: Your username supplied by s3rsync (www.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 (www.s3rsync.com).

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

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

> **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.

For information on configuring S3Rsync, see the Rsync setup guide.

7. **Notifications**

   Once a backup job has completed, BackupAssist can send an email to inform the 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.

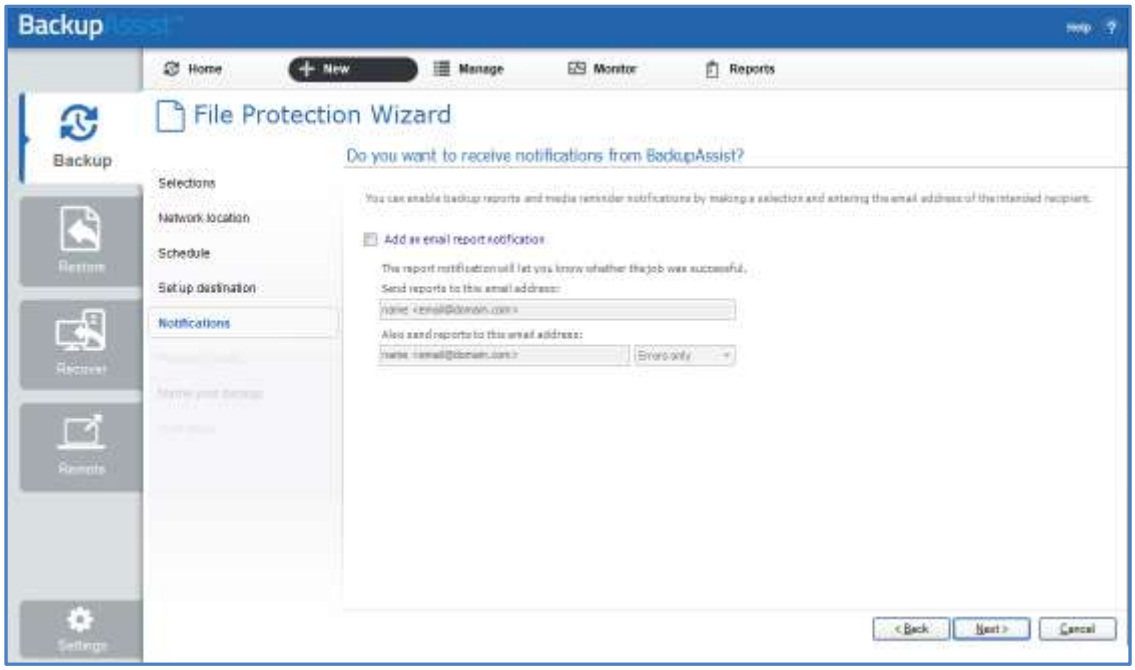

Figure 4: File Protection - Notifications setup

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.

To send email notifications, you will need to configure an SMTP mail server for BackupAssist. See the BackupAssist settings section to learn more or the Backup tab user guide for instructions.

8. **Prepare media** 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 the *backup* has been created, it should be checked. You can check the backup by performing a manual test restore, or using the Backup Verification feature. A manual restore is the only way to fully test a backup, and regular manual restores should be part of your backup solution.

4. Restoring from an rsync backup

The Restore tab displays the restore options available. This section provides instructions on how to use the *Local and Network Files* restore option, which is used to restore files and folders and VSS applications that do not have their own specific restore option.

The other restore options are documented in technology specific guides, as follows:

- For *Hyper-V Host File* and *Hyper-V Granular* restore, see the Hyper-V Protection guide.
- For *SQL Server* and *SQL Point-in-Time* restores, see the SQL Protection guide
- For *Exchange Server* and *Exchange Granular* restores, see the Exchange Protection guide

To restore data from a **File Protection rsync** backup, follow these steps:

1. **Select the Restore tab**

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

2. **Select Local and Network Files**

   This will display the volumes backed up by this installation of BackupAssist. It can also show backups from other machines added using the *Discover Backups* button, which is explained below.

   Expand a volume to display all of the backups available for that volume. There are tabs above each volume's backup list to help locate the required backup.

   - The *Last 7 days* and *Last 30 days* tabs can be used to display the backups within those ranges.
   - The *Custom* tab allows you to select a specific date range and display backups for that period.

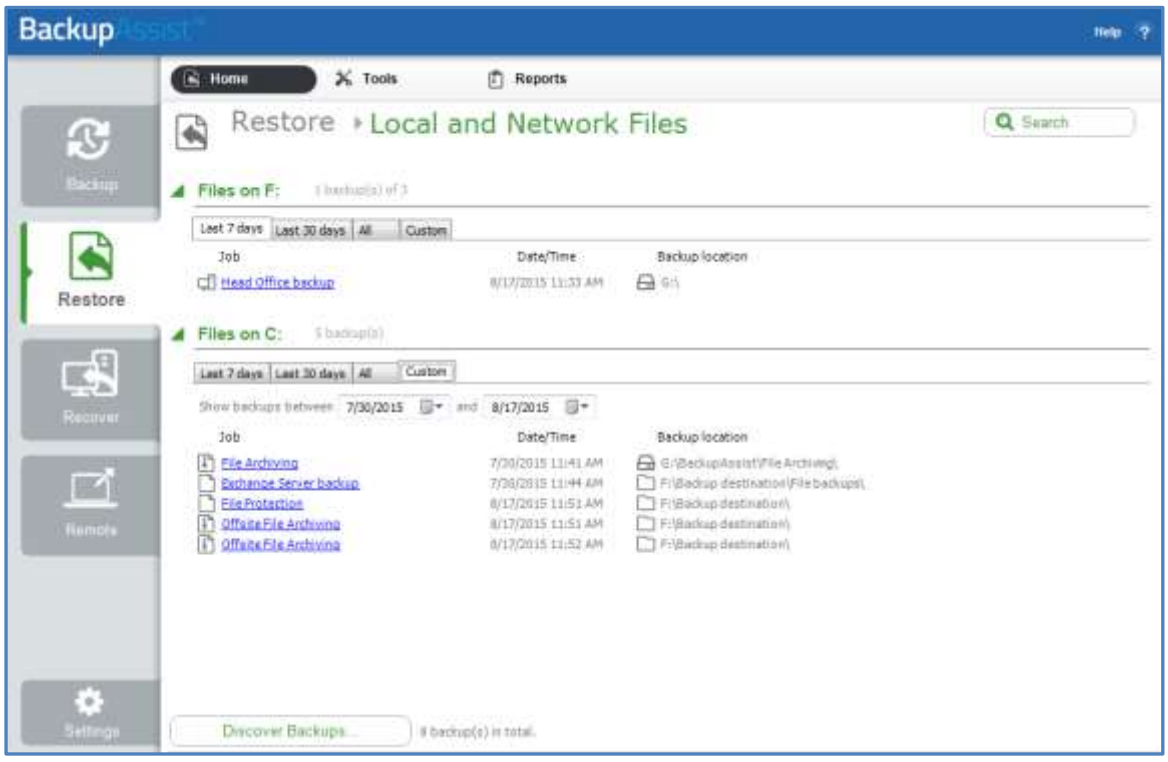

Figure 5: Restore tab – backup selection

The **Search** button allows you to locate files to restore across multiple backups. When you select Search, the Restore console will display the Search page.

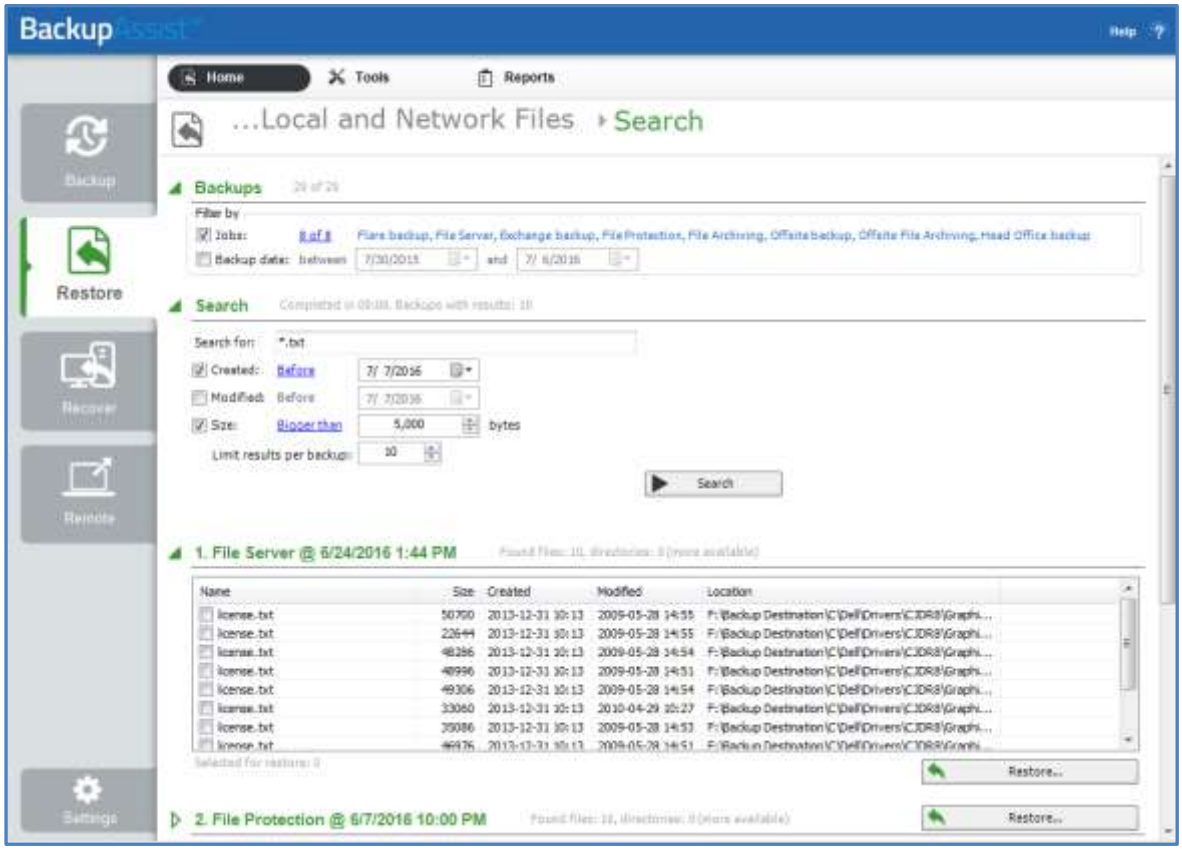

Figure 6: Restore Tab – Search page

a) The *Backups* section allows you to use the *Jobs* Filter, to limit the search to specific backup jobs. You can also use the *Backup Date* filter search within a specified date range.

b) The *Search* section is used to enter a search term associated with the name of the file you want to find. The *Search for* field will take the string provided and search for occurrences of that string within a file or directory name. The results of the search are displayed by backup.

   To refine the search, use the Created, Modified and Size options. Ticking any of these options will activate a drop down list of variables to select from. For Created and Modified, you can select a date using the Calendar selection fields. For Size, you can select the file size in bytes.

The **Discover Backups** button allows you to browse for backup catalogs created by deleted jobs and other servers. Selecting those backups will add them to the list of available backups.

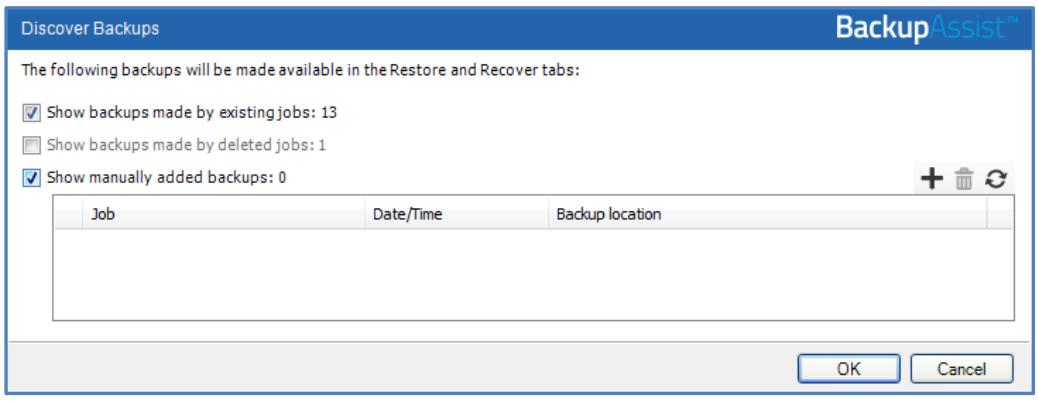

Figure 7: Discover Backups

3. **Select the backup that you want to restore from**

   Clicking on a backup's name will open the *Integrated Restore Console (IRC)*. The *Integrated Restore Console* is used to select the data to be restored, where to restore it to and the restore conditions.

4. **Select the files, folders or applications that you want to restore**

   - Use the left pane to locate and select the data that you want to restore.
   - The right pane will display the contents of the folder selected in the left pane.

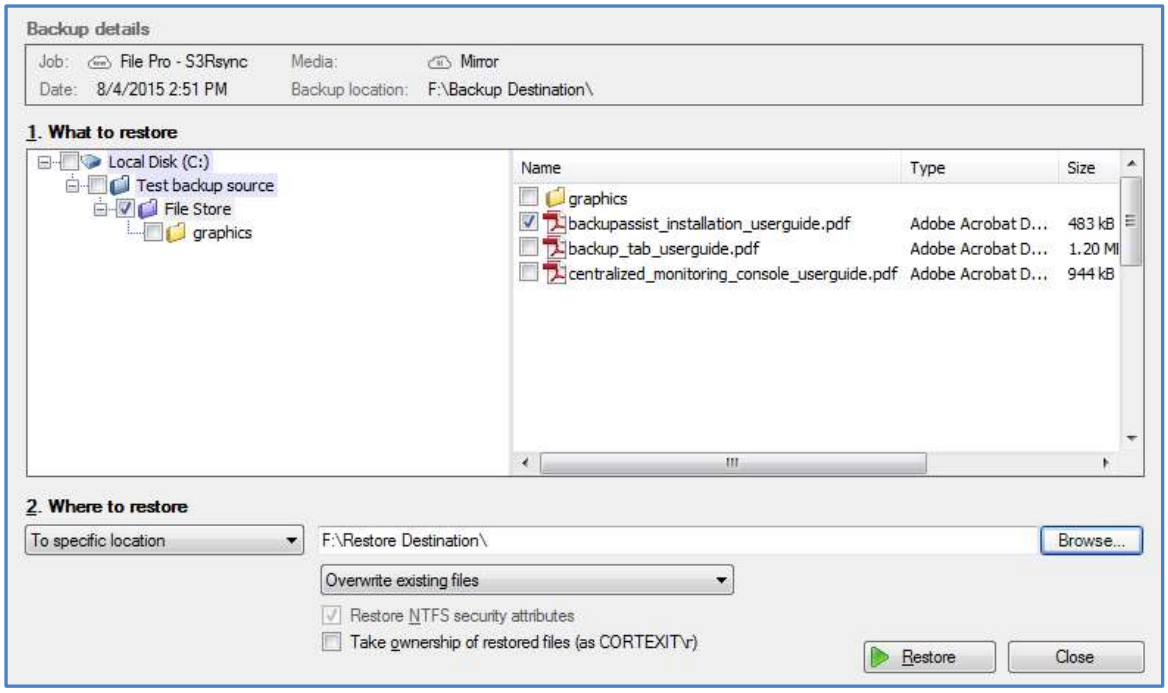

Figure 8: Integrated Restore Console

5. **Select Where to restore the data to**

   Follow these steps to select the restore destination and restore options:

   a) Under *Where to restore* select *To original location* or *To Specific location*.

   b) Use the *Browse* button to locate and select the restore destination.

   c) Use the drop down box to set the overwrite rules. The overwrite rules will apply if the files being restored encounter files with the same name in the restore destination.

      You can select:

      - *Overwrite existing files* - The restored files will overwrite files in the restore destination.
      - *Do not overwrite existing files* – The restored files will not overwrite files in the restore destination. This means the files will not be restored.
      - *Only overwrite older files* - If a source file has changed since the backup was made it will not be overwritten.

d) Review the *Restore NTFS security attributes* option

If you select this option, the NTFS security attributes the file had when it was backed up will be retained when the file is restored. The NTFS security attributes can be viewed in the Security tab on the file's Properties

e) Review the *Take ownership of restored files* option

Selecting the *Take ownership of restored files* tick box will give the current user ownership of the restored files. The user is shown to the right of the text box description.

6. **Select Restore**

When you select the *Restore* button, the restore process will begin. The *Integrated Restore Console* will display information about the restore job and provide status updates as the job runs.

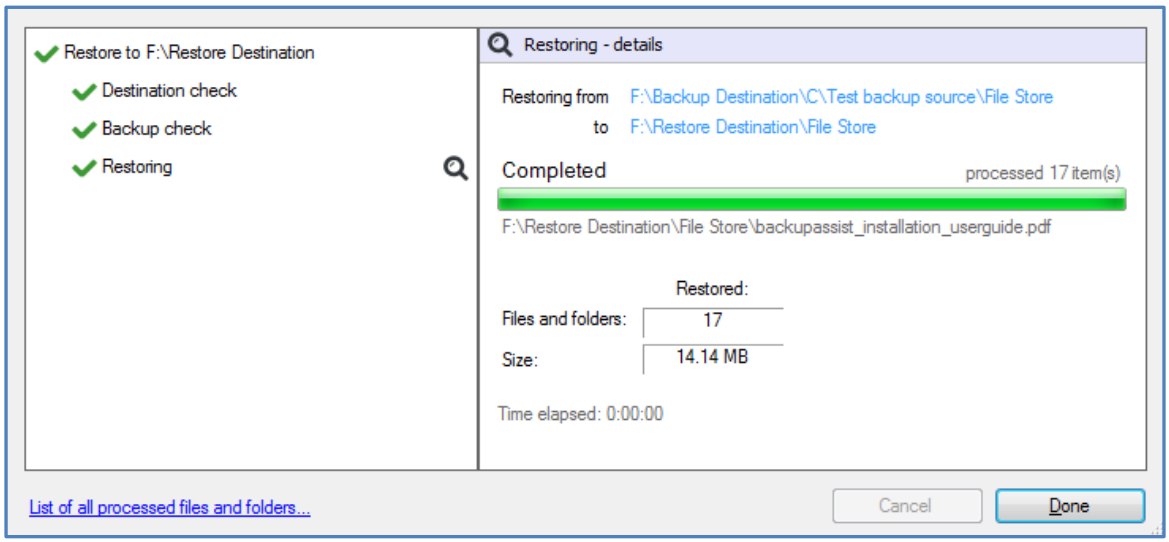

Figure 9: Integrated Restore console – restore monitor

Selecting *List all processed files and folders …* will open notepad and display a list of the files restored, including their full path.

7. **Select Done**

Once the restore has finished, selecting *Done* will return you to the main UI.

Your File Protection using rsync restore has now been completed.

Important: The *Integrated Restore Console* can restore encrypted files, but you will need to supply the password. It is important that you keep a copy of your password in a safe place, as we cannot assist you with opening password encrypted files if your password is lost or forgotten.

Helpful hint: These instructions explain how to restore data using the *Integrated Restore console*. If you do not have BackupAssist installed and need to restore a *File Protection* backup, you can manually browse the *rsync* destination and transfer data back using any method permissible by your host.

5. 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 guide.

Rsync options

Select the **Rsync options** item from the left hand menu. The *Rsync options* page contains 15 different configurations for backing up your data across the internet including:

- Rsync and SSH command line options.
- Data transfer limits.
- Backup permissions.
- Backup logging and a media usage report.
- Rsync encryption and encryption password.

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 an internet connection can take a long time. For this reason, two data seed options are available for rsync host servers that you have local access to. This may exclude third party vendors. 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's *Seed Backup* function, to automatically seed data offsite using a removable media, which can be physically transported to the data host so that the data can be uploaded locally.

To seed your data using this method, follow these steps:

1. Connect a removable media device to the machine running BackupAssist.
2. Select your backup job from the **Manage** Menu
3. Select **Edit** from the top menu.
4. Select the **Destination** left menu item.
5. Click the **Seed backup** button
6. Select the location of an empty folder on your portable media.

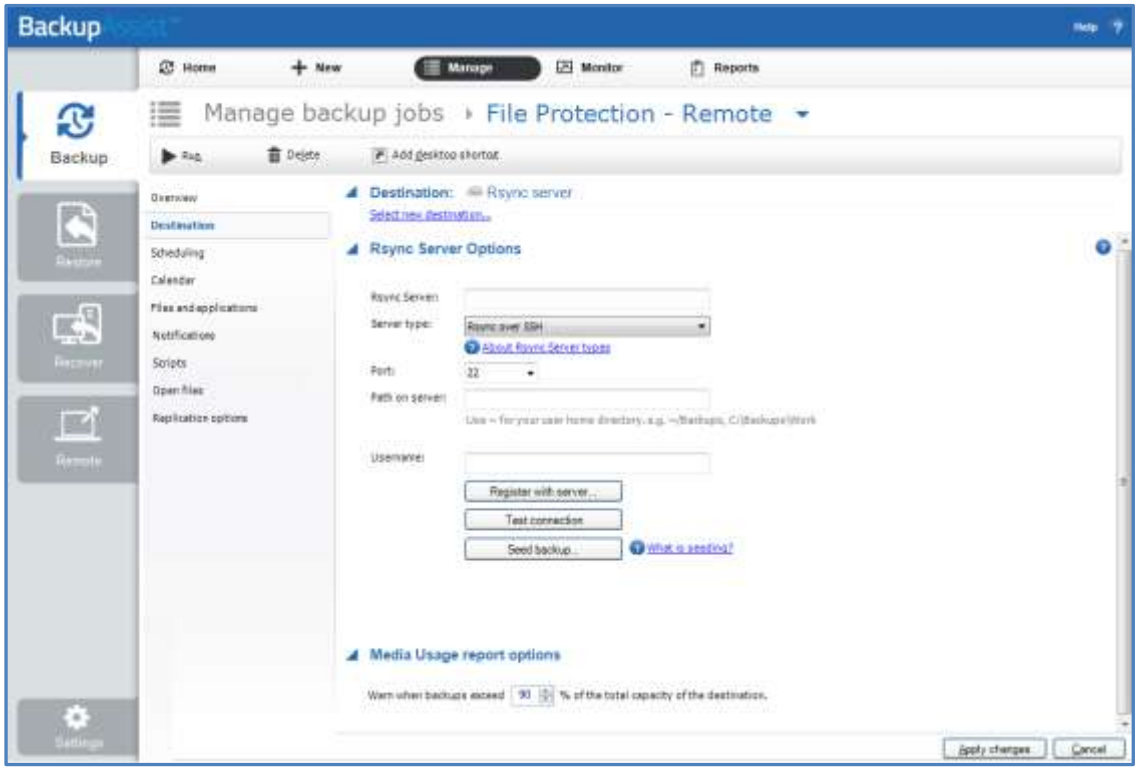

Figure 10: Manage backup jobs screen – Destination seeding

7. 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.
8. Transport the portable media containing the seed to the site where your rsync server is located.
9. Connect the device to the rsync host 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.

Scheduling

Selecting *Scheduling* will display the **Scheduling options.** You can use this screen to change the default time and days of your scheme's daily backups. If you selected a scheme with archive backups (e.g. weekly, monthly), you can specify when each archive backup will run. The current scheme is shown, along with two pop-up menus: *Select a new schedule* and *Customize schedule*.

Select a new Schedule: This will display the pre-configured backup schemes that you chose from during the creation of your backup job. The selections available will depend on the type of destination media you have selected. You can select a different scheme using this option.

Customize schedule: This selection can be used to modify each backup within your current schedule. The customizations available will depend on the type of backup media used. For File Protection backups, the *Method* field can only be set to *Automatic*. This is because single instance store provides the benefit of incremental backups in a full backup format. This technology is managed by BackupAssist and does not require further modification.

Files and applications: VSS

The Volume Shadow Copy Service (VSS) is a Microsoft Service that creates a copy of an application's data (e.g. Exchange and SQL) so the data can be backed up without interfering with the application. BackupAssist is a VSS-aware backup client, so it can backup application data using VSS. BackupAssist will automatically detect *locally* running VSS applications and list them for selection during the **Destination** step of the backup job creation.

VSS applications are displayed under the **Files and applications** menu item. You can modify your backup job by selecting entire VSS applications or drilling down to individual components. In some cases, only applications that are currently running will be detected. If an application is not listed, try re-starting it and then click the *Refresh* button in BackupAssist.

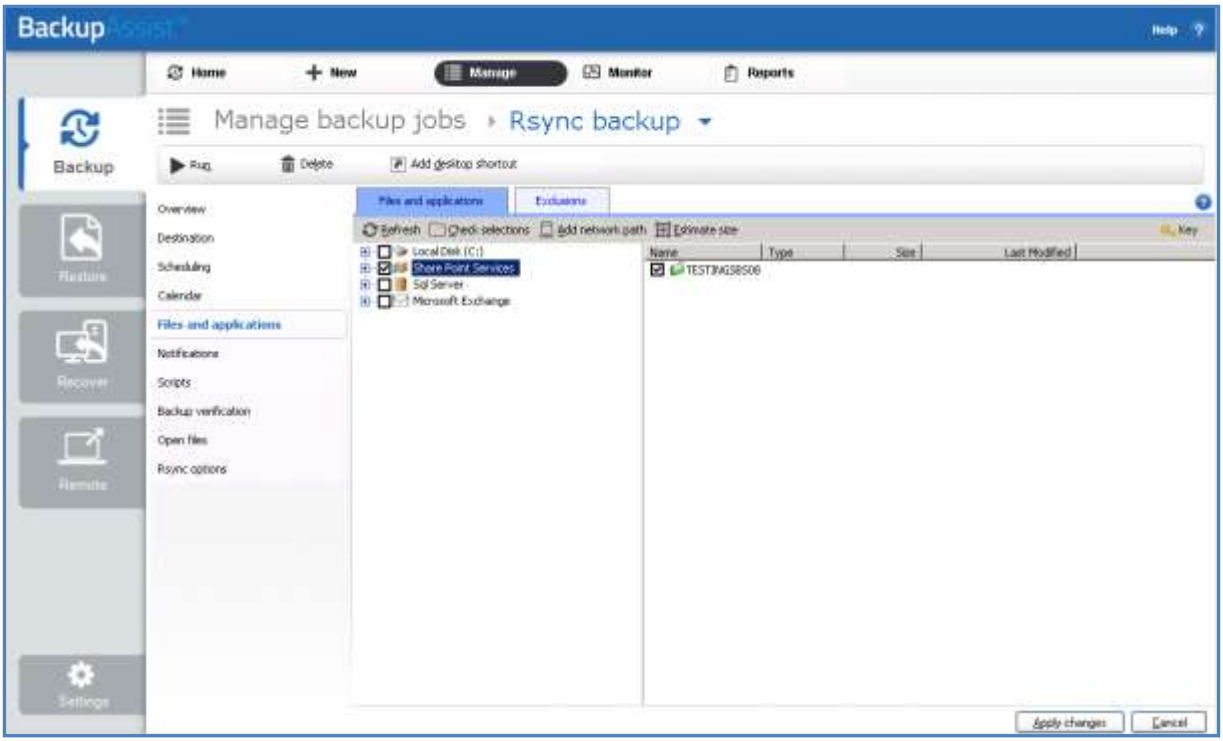

Figure 11: Manage backup jobs screen – File and applications option

6. Support and Resources

These resources can be used to help troubleshoot common rsync backup problems.

Troubleshooting FAQ

Test connection failed

Ensure that you are able to ping your rsync server from your BackupAssist server and that you have opened up the appropriate ports on your firewall. Make sure that the username can access the path you have specified.

SSH Connection Refused

Ensure that the services *Openssh SSHD* and *RsyncServer* are started on the data host machine (Administrative Tools > Services). Make sure your firewall is not blocking the attempt.

Register with server failed

Ensure that you have the correct username and password set up on your rsync server.

Appendix

Data host

The server that has been set up to host backup data.

Client

The machine that BackupAssist is installed on, that sends data to the data host.

SSH Authentication

For SSH communication, we use a public / private key method of authentication, meaning that you will only be asked for your password once (when registering with the server), and your public key will be uploaded to the server, enabling BackupAssist to log into the server in the future in a secure, password-less manner. For more information on public / private key authentication, visit the following Wikipedia article: Wikipedia Public Key Cryptography

Daemon Authentication

In Daemon mode, your password is stored in an encrypted format by BackupAssist and provided every time the backup runs. When running in Daemon mode, traffic will be unencrypted. For this reason, we recommend that you only use closed network environments, such as LANs or WANs connected by a secure VPN. VPNs inherently encrypt communication between nodes, so using rsync in Daemon mode over a VPN is still secure.