# System Protection for Hyper-V

# 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.

backupassist.com

# Contents

# 1. Overview

File Archiving, File Protection and System Protection can back up Hyper-V and Hyper-V CSV environments. System Protection backups can also be used to perform granular restores of Hyper-V guests using the Hyper-V Granular Restore Console. For this reason, System Protection is recommended for all Hyper-V backups. This whitepaper explains how to create, modify and restore Hyper-V backups using *System Protection*.

Hyper-V is Microsoft's implementation of virtual machine (VM) technology. Virtual machines run on a physical host machine and provide dedicated operating system environments (guests) to run applications (e.g. Exchange) or functions (e.g. a print server).

CSV is Microsoft's implementation of server clustering, which is designed for use with Hyper-V. CSV provides a shared disk that can be used by any guest (virtual machine) in the cluster. This means a guest can be moved within the cluster, and guests can share the same physical disk.

## Documentation

**More information on Hyper-V can be found in the** System Protection for Hyper-V whitepaper.

The whitepaper contains comprehensive information and should be referred to when planning a backup strategy using System Protection for Hyper-V environments.

For general advice on Hyper-V protection, see the Hyper-V backup implementation guide

Other BackupAssist documentation can be accessed through the Documentation webpage

## Licensing

To restore Hyper-V data on a guest at a granular level requires the *Hyper-V Granular Restore Add-on* license, once the initial trial period has expired. For instructions on how to activate license keys, visit out Licensing BackupAssist page.

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

## Hyper-V requirements

**Supported operating systems**

- Windows Server 2012/ R2, including Server Core and Hyper-V Server versions.
- Windows Server 2008R2, including Server Core and Hyper-V Server versions.
- Windows Server 2008R1

**Pre-requisites**

- Windows Backup - it is included with Windows Server but may not be installed.
- In a CSV environment, BackupAssist must be installed and licensed on each host.
- Partition sizes should be less than 2TB.
- Turn on the Backup Integration service for each Hyper-V guest as explained in this blog post.

> **Important**: Before implementing a Hyper-V protection solution, please refer to the Hyper-V implementation guide for advice and useful tips on backing up Hyper-V hosts and guests.

# 2. 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.

## Windows Server 2012

Windows Server 2012 introduces the CSVFS format, which allows a cluster to differentiate CSV storage from other NTFS storage. BackupAssist supports backing up Hyper-V guests located on CSV storage using the CSVFS file system.

- BackupAssist supports SMB 3.0 servers as a CSVFS backup destination.
- Bit Locker drive encryption is not supported.
- BackupAssist does not support backups of files from both CSVFS and NTFS locations in the same snapshot. For example, it is not possible to back up a Hyper-V guest with a source in a CSV along with another guest that is on an NTFS volume.

Hyper-V is Microsoft's implementation of virtual machine (VM) technology. Virtual machines run on a physical host machine and provide dedicated operating system environments.

- Hype-V Replica enables replication of a virtual machine across the internet. BackupAssist supports Hyper-V Replica and can back up and restore a primary Hyper-V guest using any backup type.

CSV Environments: CSV backups are supported by System Protection, File Protection and File Archiving backup types, but CSVFS and NTFS file systems cannot be mixed (as described above).

## 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.

# 3. Hyper-V in a Shared Clustered Volume

When stored in a Cluster Share Volume (CSV), a virtual machine's files (virtual disks and configuration files) present as a shared resource between the hosts. As the CSV's files are shared, access to the virtual machine's files must be coordinated so only the one host accesses a virtual machine's files at a time.
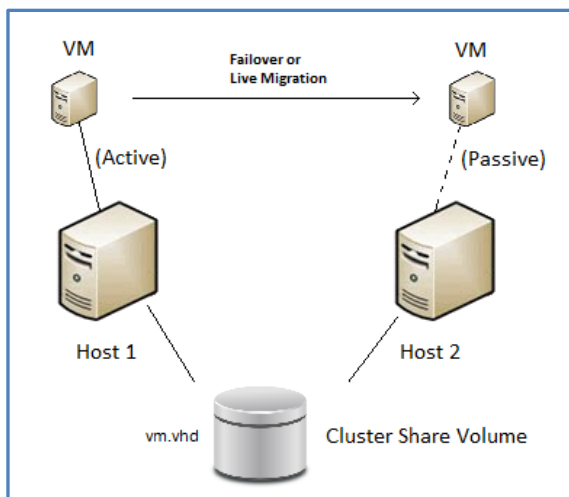


This diagram depicts a configuration with a small cluster of two hosts with a single virtual machine, which may fail-over at any time from Host 1 to 2.

This virtual machine will only be active on a single host at any time - referred to as the *active node*. The other node is called the *passive node*.

Other virtual machines added to the cluster may be active on either host at any point in time.

## Backing up a CSV environment:

Because a backup of a CSV must be coordinated, BackupAssist jobs on each host must have their start times staggered. The scheduled start time of the jobs on each host must be set at least 5 minutes apart. This will allow BackupAssist time to initialize and coordinate its access to the CSV. BackupAssist will automatically delay the second job so that it will not commence until the first job has completed. During a *System Protection* backup, the host and all guests are copied to an intermediate staging disk. The intermediate staging disk must be a local disk and it is overwritten each time a backup is run.

- BackupAssist v7 supports Hyper-V CSV on both Windows Server 2008R2 and 2012 machines.
- When a fail-over or migration occurs, one or more of the virtual machines will no longer be active on the original host. Therefore, BackupAssist must be installed and licensed on each host.
- A backup job must be created on each of the hosts to back up all virtual machines which are active on that host (at the time the backup runs). This can be done by selecting the entire Hyper-V VSS object for each job, which automatically finds active virtual machines at the time the backup runs.

## Restoring to a CSV environment

When a guest is successfully restored to a Hyper-V cluster (CSV), it will become available on the Hyper-V machine that it was restored to.

- If the restore replaces the existing guest, the CSV settings will be retained and the guest will be added back to the cluster.
- If the restore does not replace the existing guest, the CSV setting will need to be manually configured to add the guest to the cluster.
- This affects Hyper-V CSV implementations on Windows 2008/ 2008R2 / 2012.

**Note**: 1) The Hyper-V Granular Restore Console does <u>not</u> support mounting image backups (VHD files) that were backed up off a CSV. 2) Windows Server 2012R2 does not support CSV environments.

# 4. Creating a Hyper-V backup

The following instructions describe how to back up Hyper-V environments using System Protection.

File Archiving and File Protection backups can also detect and back up Hyper-V environments but they do not support granular restores of guests or full recoveries of Hyper-V servers.

Launch BackupAssist and follow the steps outlined below:

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

2.  Select **System 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 Hyper-V environments to be backed up.

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

    All Hyper-V environments detected will be displayed here as application containers. The example below is of a CSV environment, so there is also a CSV tab.

    **Files and Applications** tab.

    Select the Hyper-V environments to be backed up. You can also select the *Back up Entire System* option to create a backup that can be used for a system recovery.
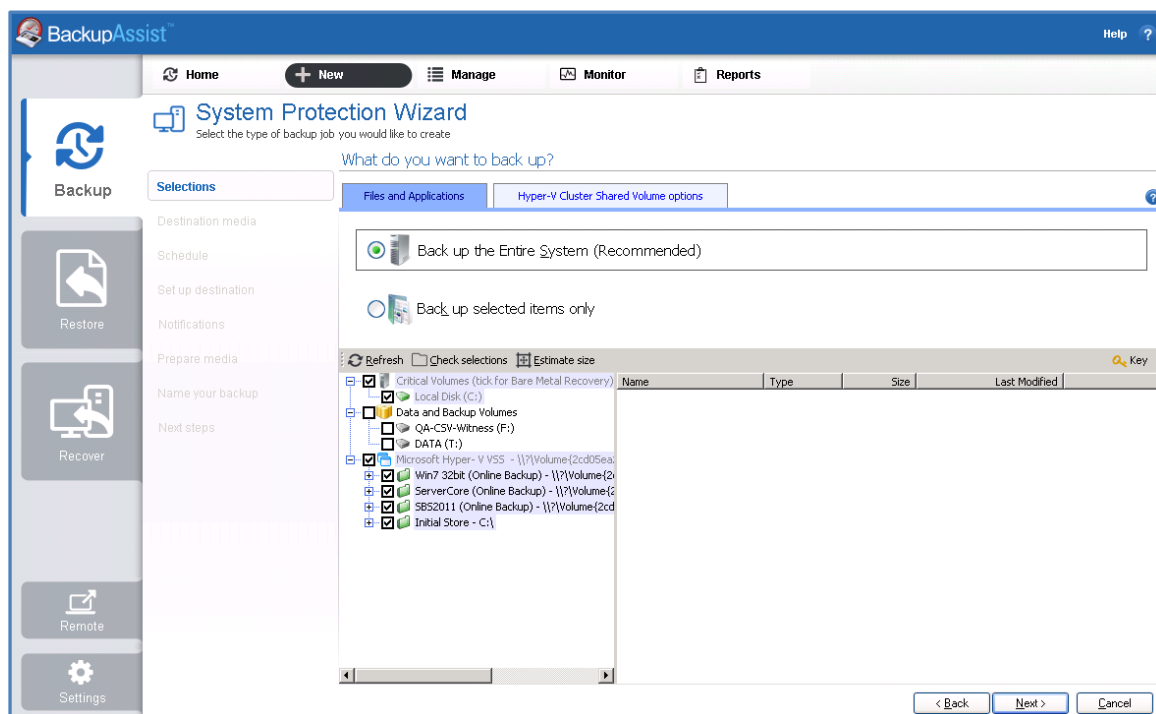


Figure 1: BackupAssist Hyper-V backup – Selections screen

- For a **non-clustered** host: Select **Next** and you will be taken to the *Destination media* step.

- For a **clustered** host: Select **Next** to proceed to the *Hyper-V Cluster Shared Volumes* tab.

**Hyper-V Clustered Shared Volume options** tab.

If you are creating a Hyper-V backup job for a shared clustered volume, then this tab will need to be completed. This step only applies to System Protection, Hyper-V backups using CSV. The tab will not appear in non-clustered environments.

The *Hyper-V Cluster Shared Volume* tab graphically shows the guests selected, and a field to select the intermediate backup location. The intermediate backup is used to copy the host and guest data to a single location. This is also known as the *staging disk*. An image backup will then be made. The image is of the host. Within that image, is an image of each guest residing on the host.

- The intermediate backup's location must be a local disk.
- The intermediate backup's disk is overwritten each time a backup is run.
- The image backup is taken from the whole disk.

Enter the drive to be used for the intermediate backup, and click **Next**.
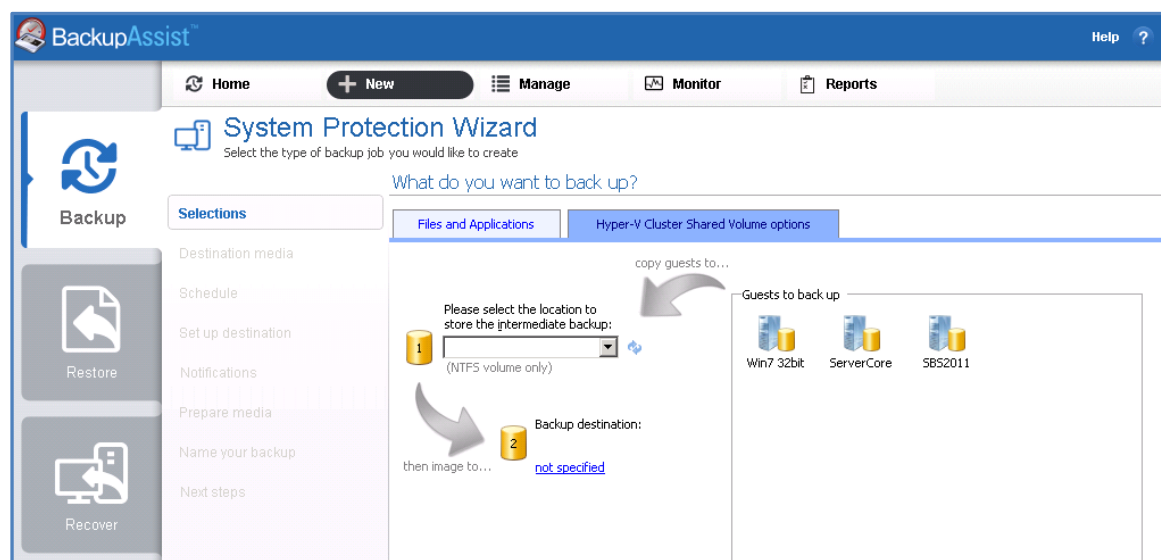


**Figure 2: Hyper-V intermediate (staging) disk selection**

**Note**: The entire contents of the staging disk will imaged to the backup destination each time the backup is run. Check that the disk is both appropriate and prepared for this step.

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 the media type selected, when you click next.

   The **Enable Data container**, option is available for the following destinations: *RDX drive*, *Local hard drive, Network location* and *External disk*. To learn more: [Data container resource page](Data container resource page).

   **Note**: *System Protection* backups on *RDX drives* cannot be used to restore individual files unless Data containers are used. This applies to Windows Server 2008R2 and later.

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. 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.
   - The schedule can be customized after the backup job has been created.

6.  **Set up destination:** This screen is used to configure the location of the media selected in step 4.

    -   The options presented will change with the type of media selected.
    -   If your media is removable, you can set the media to eject after the backup job has finished.
    -   Data containers are supported on Window Server 2008 R1/R2 and SBS 2008/2011 and later operating systems. If your destination is a Data container, the container size and location is set using this screen. It is important to review these settings to ensure they are correct, because the size of a Data container cannot be changed once the backup job has run. The *Use all available space* selection will use up to 2TB of space.

    Configure your backup destination, and click **Next**.

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:** If you selected a portable media device as your backup destination (such as an RDX drive or an external drive) you will be given the option to prepare the media for BackupAssist. BackupAssist will write a label onto the media so that it can recognise what media has been attached, and determine if it is the correct media for your backup schedule.

    To enable media detection:

    a.  Select, **Let BackupAssist keep track of your media.**
    b.  Select what you would like BackupAssist to do, *if the wrong media is inserted*.
    c.  Select what you would like BackupAssist to do, *if new or unrecognized media is inserted*.
    d.  Enter the name and drive designation to be used for each media device listed.
    e.  Select **Prepare** for each media device listed.

    BackupAssist will write the label to the media so that it is able to recognize the media and ensure that the correct media is being used on the correct day.

9.  **Name your backup:** Provide a name for your backup job, and click **Finish**.

10. **Next Steps:** If you are creating a backup of your entire system for use in a recovery, you can use this option to launch the RecoverAssist builder and create and bootable recovery media.

▶ **Your System Protection for Hyper-V backup job has now been created.**

**Important:** Once a backup job has been created, it should be reviewed and run using the *Manage* menu. See the section, Hyper-V 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. See the section Restoring a Hyper-V backup.

**Important:** Hyper-V Granular Restore does not support mounting image backups that were backed up off a CSV

# 5. Restoring a Hyper-V backup

This section provides instructions on how to restore files from within a guest or host (granular) or an entire guest or host environment using a System Protection backup image.

A System Protection backup image can be used by two different restore tools:

- The **Hyper-V Granular Restore** tool, which can restore files from within a guest (virtual machine).
- The **Restore Console**, which can restore a guest image and files from within a host.

When performing a restore, the steps required to locate the backup are the same for both restore tools. This section explains how to locate the backup and is followed by sections explaining how to use the *Restore console* and how to use the *Hyper-V Granular Restore* tool, to restore from the backup.

To restore Hyper-V data from a **System 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.

2. Select **Hyper-V** from the Home page.

3. The **Home page** will now display all backups that contain Hyper-V data.

   Hyper-V image backups will be displayed twice for each backup location. First in a group with the prefix *Entire* – associated to the *Restore Console* tool, and then in a group with the prefix, *File* – associated to the *Hyper-V Granular Restore* tool.

   ➢ Select **the Restore console** - to restore a guest (virtual machine) or files from within the host – and go to the section below, Hyper-V Guest & Host files restore,

   ➢ Select **the Hyper-V tool** – to restore individual files from a guest (virtual machine) – and go to the section below, Hyper-V granular restore,



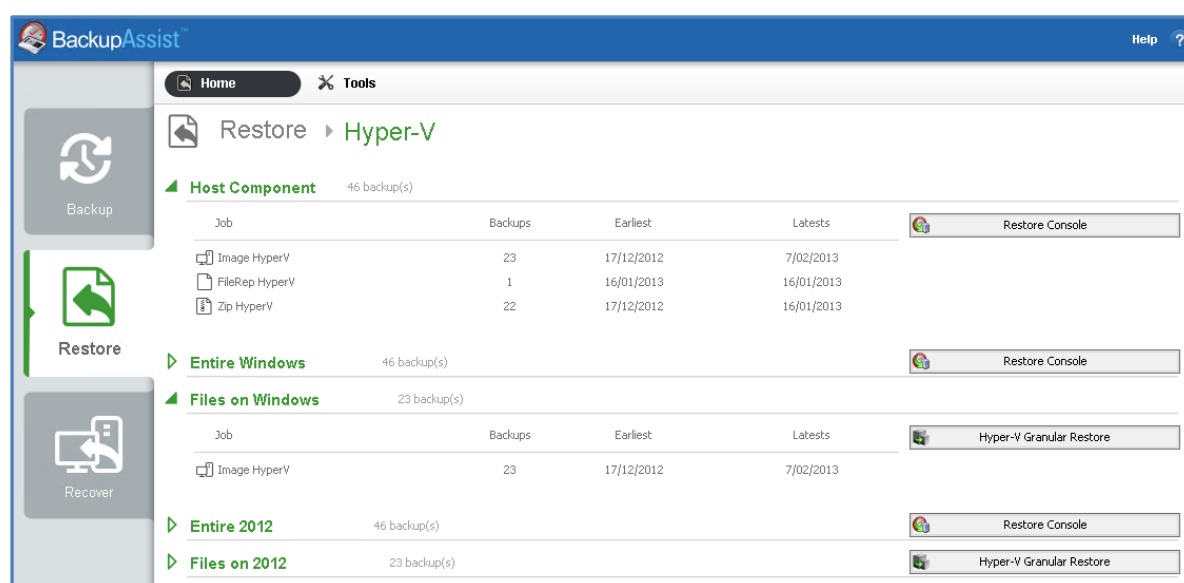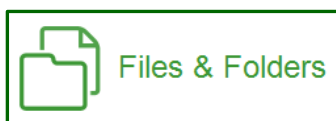**Figure 3: BackupAssist Restore screen - restore tool selection**

# Hyper-V guest & host files restore

Files & Folders

The following section explains how to use the BackupAssist *Restore Console*, to restore a *guest* (virtual machine) or files from within a Hyper-V *host* using a System Protection image backup.

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

The *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.
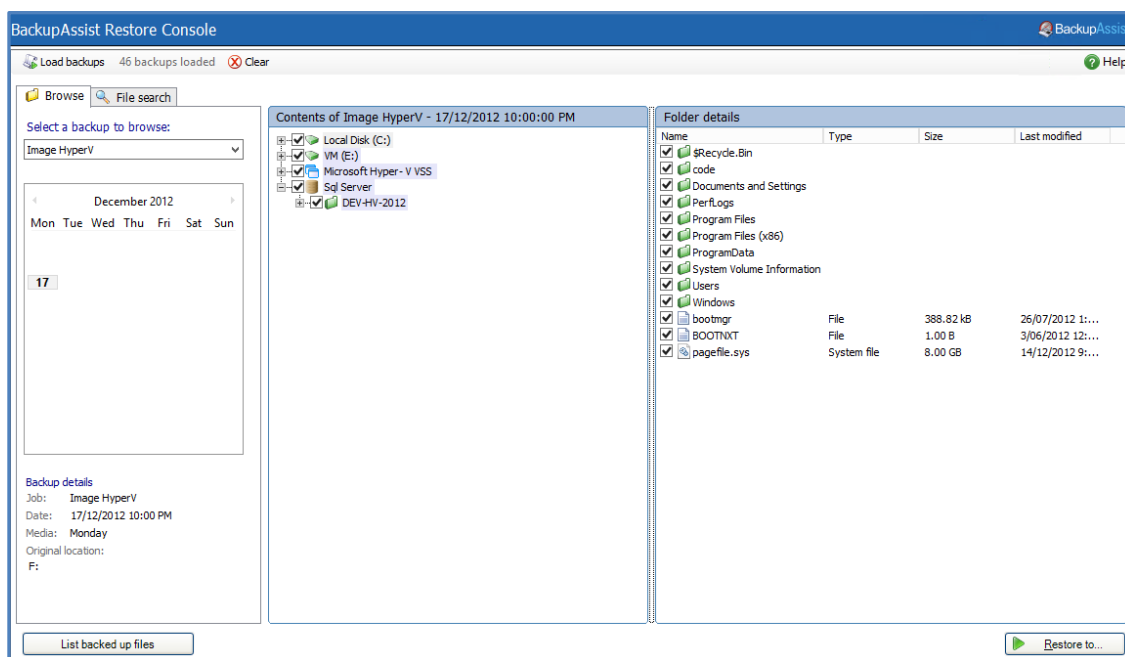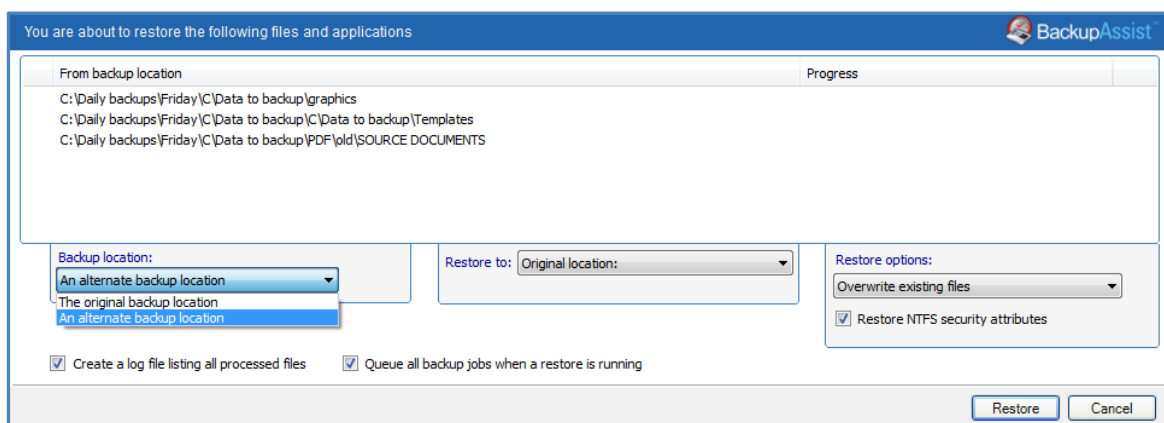


**Figure 4: BackupAssist Restore Console – backup and data selection**

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*.

**Figure 5: BackupAssist Restore Console – restore destination**

   a. Review **Backup location:** Change the selection if the backup was moved after it was created.

   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 to restore your data.

      If BackupAssist cannot access the backup location you will be prompted to either connect the appropriate media or specify an alternate location where the backup can be found.

      The restore will run from the destination window and a **Report** link will appear once the restore has finished.

   g. Select **Done**.

▶ **Your Hyper-V restore has now been completed.**

**Important:** Only backups made with BackupAssist v5.3 or later will show up in the Restore Console.

**Important**: Refer to the Restoring to a CSV environment section, if you are restoring a guest to a CSV.

# Hyper-V granular restore



The following section explains how to use the *Hyper-V Granular Restore* tool, to restore individual files from within a guest (virtual machine) using an image backup.

The BackupAssist *Hyper-V Granular Restore* tool will only run on 64-bit Window 2008 / 2012 Servers or a standalone Hyper-V Server. The machine must also have the Hyper-V role, which can be added to most Windows 2008 / 2012 Servers, and is the default for a Hyper-V Server.

ℹ️ Hyper-V Granular Restore does not support mounting image (VHD) backups that were backed up off a CSV.

1. **Select, Hyper-V Granular Restore**

   The tool will launch the *BackupAssist Image mounting tool*, which will open the *Hyper-V Granular Restore* console. The console is used to locate the image backup.

2. **Locate your backup**

   *Select a local drive* that contains your System Protection VHD backup (Press *Detect* if the image does not display), or select *Search a location* and use the *Browse* or *Search* functions to locate your VHD backups.

   Once the backups have been located, a list of VHD files will be displayed:

   - The *Time* column can be used to help determine what backup to select.
   - The *Available guests* column can be used to see what guests are inside the image backup.
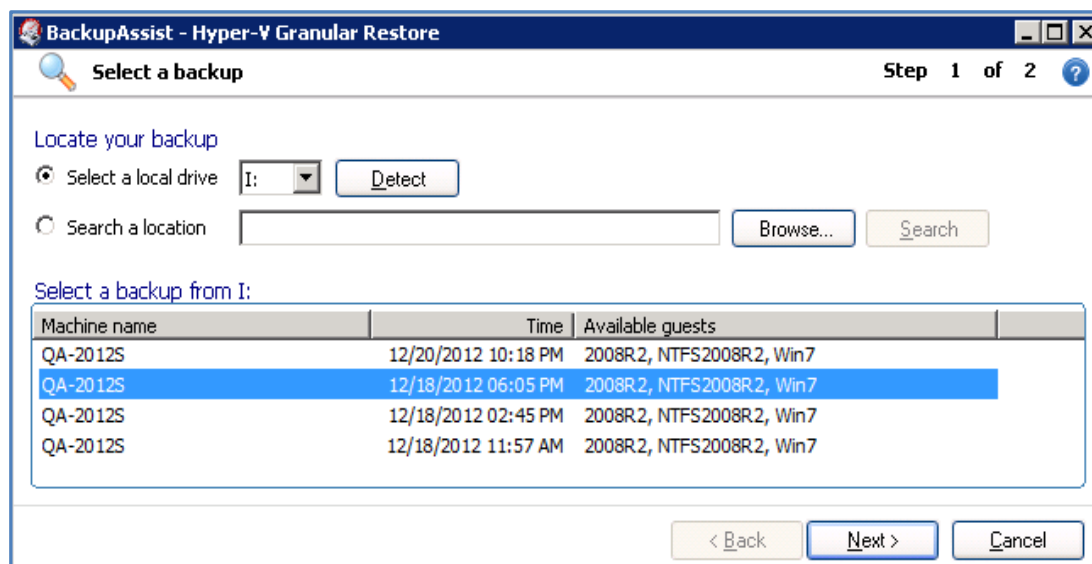


**Figure 6: BackupAssist disk mounting tool - detected Hyper-V backups**

3. **Select the host image to restore from**, and click **Next**.

   You will be presented with the **Mount volumes** screen, as shown in *Figure 11* below. The Mount Volumes screen is used to select the guest, within the VHD backup, that you want to restore data from.

4. **Select the guest image**, using the drop down list. The guest will appear as a selected item.

5. **Select the drive letter** to mount the guest image as.

6.  Select the appropriate action from the drop down list.

    - Option 1 –**Mount selected volumes to drive letter.**
    - Option 2 –**Export volumes as VHD files.**



**Figure 7: Selecting guest image to mount**

7.  **Option 1 –Mount selected volumes to drive letter.**

    - The Mounting Disk progress bar will be displayed until the image is mounted.
    - If there are two partitions, you will be asked to select which one.

    Once the *guest* VHD is mounted, a new drive will appear on your computer with the drive letter specified on the guest selection screen.

    a.  Open Windows Explorer to view the contents of the image.
    b.  Copy the files and folders you want to restore (from the mounted VHD file) to a network location that the guest machine can access.
    c.  Once you have finished copying data to a network location, you can then copy it back to the original guest machine.
    d.  Once your data restoration is complete, click **Unmount** in the Hyper-V Granular Restore Console, and click **Done**.

    **Option 2 –Export volumes as VHD files.**

    Use this option if you would like to mount the image on another machine.

    a.  Set the location that you would like to export the VHD files to in the field located at the bottom right corner.
    b.  Click **Finish** to export the selected guest image.

        A new drive will appear on the selected machine using the letter specified. Open Windows Explorer to view the contents of the image and follow the steps in option 1.

## Quick recovery from staging disk

When a backup of a Hyper-V in a CSV environment is created, an intermediate staging disc is used to create the image backup. If the guest VM needs to be recovered (or its image backup is rendered unusable) the intermediate staging disc and be used to recover the guest. Configure a new guest using the staging disk and the virtual machine will be operational again with minimal data loss. When you are ready to perform a migration back to the server, the differencing VHD stored on the staging disk along with the last back up from the staging disk can be merged together and installed on the server.

# 6. Hyper-V 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.

## 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. Some backup types will have additional configuration options. 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. 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 and the type of backup media used.

For additional information on *Scheduling*, please refer to the whitepaper for the type of backup you are using (e.g. System Protection), please refer to the Backup tab whitepaper.

## Files and applications

**File and applications tab**: A new *System Protection* backup job will back up an entire disk or application. However, under the *Files and applications* menu item, you can modify your backup job by selecting specific files and folders, or individual components within a VSS application.

**Hyper-V Shared cluster volume options tab:** Changes the intermediate (staging disk) backup location.

**Exclusions tab:** This is an additional tab where certain file types can be excluded from the backup. The exclusions can be applied to selected locations within the backup job. Use the *add* button to enter a *file type* and the *location* to apply the exclusion to. Pre-defined exclusions can also be added.