# Acronis Backup & Recovery 11

Update 0

# User Guide

Applies to the following editions:

☐ Advanced Server

☐ Virtual Edition

☐ Advanced Server SBS Edition

☐ Advanced Workstation

☐ Server for Linux

☑ Server for Windows

☑ Workstation

# Table of contents

# 1 Introducing Acronis Backup & Recovery 11

## 1.1 What's new in Acronis Backup & Recovery 11

Acronis Backup & Recovery 11 builds on the success that Acronis Backup & Recovery 10 has established by bringing enterprise-class capabilities to the small business market at an affordable price in an easy-to-use package.

Acronis Backup & Recovery 11 continues the trend of expanding the backup and recovery capabilities in physical, virtual and cloud environments. The following is a summary of the product's new features and enhancements.

- **Simplified installation**

  The new installer makes the installation procedure simple and clear.

- **Improved usability**

  The redesigned product's UI lets you perform any operation easier, faster and more intuitively.

- **Advanced replication and retention of backups** (p. 66)

  Store a backup in multiple locations (possibly off-site) for redundancy. Move or copy backups to a cheaper or off-site storage automatically. Set a replication time window if you do not want copying or moving to occur during business hours.

- **Data view for vaults** (p. 99)

  Select data from a vault by browsing either the archives and backups (in the **Archive view**) or the backed up data (in the **Data view**).

- **Alert notifications** (p. 200)

  A new alert system has been introduced for both local and centralized management. Select the alerts you want to observe. Set up e-mail notifications about various types of alerts.

- **UEFI/GPT support** (p. 114)

  Full support for UEFI-based systems and GPT disks. Recovery of BIOS-based systems to UEFI-based systems and vice versa.

- **4-KB drives support (p. 108)**

  When recovering disks or volumes, the software automatically eliminates volume misalignment – a situation that occurs when volume clusters are not aligned with disk sectors.

- **Partition (volume) alignment** (p. 108)

  Solid State Drives (SSD) require a specific partition alignment for optimal performance. The required alignment is set automatically during recovery, but you can change it manually if required.

- **Automatic disk/volume mapping** (p. 105)

  When recovering disks or volumes, the software automatically maps the selected disk/volumes to the target disks in the optimal manner.

- **Hardware snapshot provider support** (p. 95)

  The Volume Shadow Copy Service (VSS) can use hardware-based providers for taking snapshots.

- **Applying Acronis Universal Restore without recovery** (p. 112)

  Using bootable media, you can apply Acronis Universal Restore to an operating system without performing the recovery.

- **Exporting and importing backup plans** (p. 193)

  Export a backup plan to an .xml file and import it to a different machine.

- **Deploying backup plans as files** (p. 196)

  Export a backup plan from one machine and deploy it as an .xml file to multiple machines.

- **Disaster Recovery Plan** (p. 82)

  The software can generate a disaster recovery plan and send it via e-mail right after a backup creation. The plan contains step-by-step instructions on how to recover.

- **Converting a backup to full** (p. 157)

  Convert an incremental or differential backup to a full one.

- **New command line**

  Provides backup and recovery automation. Includes remote management.

- **Automatic check for updates**

  The management console automatically checks for updates upon each start and provides notification once the newer version is available.

# 1.2 Acronis Backup & Recovery 11 components

This section contains a list of Acronis Backup & Recovery 11 components with a brief description of their functionality.

## Components for a managed machine (agents)

These are applications that perform data backup, recovery and other operations on the machines managed with Acronis Backup & Recovery 11. Agents require a license to perform operations on each managed machine. Agents have multiple features, or add-ons, that enable additional functionality and so might require additional licenses.

## Console

The console provides Graphical User Interface to the agents. Usage of the console is not licensed. In stand-alone editions of Acronis Backup & Recovery 11, the console is installed together with the agent and cannot be disconnected from it.

## Bootable Media Builder

With Bootable Media Builder, you can create bootable media in order to use the agents and other rescue utilities in a rescue environment. In stand-alone editions of Acronis Backup & Recovery 11, Bootable Media Builder is installed together with the agent. All add-ons to the agent, if installed, will be available in a rescue environment.

# 1.2.1 Agent for Windows

This agent enables disk-level and file-level data protection under Windows.

## Disk backup

Disk-level data protection is based on backing up either a disk or a volume file system as a whole, along with all the information necessary for the operating system to boot; or all the disk sectors using the sector-by-sector approach (raw mode). A backup that contains a copy of a disk or a volume in a packaged form is called a disk (volume) backup or a disk (volume) image. It is possible to recover disks or volumes as a whole from such backup, as well as individual folders or files.

**File backup**

File-level data protection is based on backing up files and folders residing on the machine where the agent is installed or on a network share. Files can be recovered to their original location or to another place. It is possible to recover all files and folders that were backed up or select which of them to recover.

**Other operations**

**Conversion to a virtual machine**

Rather than converting a disk backup to a virtual disk file, which requires additional operations to bring the virtual disk into use, Agent for Windows performs the conversion by recovering a disk backup to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat Kernel-based Virtual Machine (KVM). Files of the fully configured and operational machine will be placed in the folder you select. You can start the machine using the respective virtualization software or prepare the machine files for further usage.

**Disk management**

Agent for Windows includes Acronis Disk Director Lite - a handy disk management utility. Disk management operations, such as cloning disks; converting disks; creating, formatting and deleting volumes; changing a disk partitioning style between MBR and GPT or changing a disk label, can be performed either in the operating system or using bootable media.

## 1.2.1.1    Universal Restore

The Universal Restore add-on enables you to use the restore to dissimilar hardware functionality on the machine where the agent is installed and create bootable media with this functionality. Universal Restore handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

## 1.2.2    Management Console

Acronis Backup & Recovery 11 Management Console is an administrative tool for local access to Acronis Backup & Recovery 11 agent. Remote connection to the agent is not possible.

## 1.2.3    Bootable Media Builder

Acronis Bootable Media Builder is a dedicated tool for creating bootable media (p. 229). The media builder that installs on Windows can create bootable media based on either Windows Preinstallation Environment, or Linux kernel.

The Universal Restore (p. 10) add-on enables you to create bootable media with the restore to dissimilar hardware functionality. Universal Restore handles differences in devices that are critical for Windows start-up, such as storage controllers, motherboard or chipset.

## 1.3    Supported file systems

Acronis Backup & Recovery 11 can back up and recover the following file systems with the following limitations:

- FAT16/32
- NTFS
- Ext2/Ext3/Ext4
- ReiserFS3 - particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 11 Storage Node
- ReiserFS4 - volume recovery without the volume resize capability; particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 11 Storage Node
- XFS - volume recovery without the volume resize capability; particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 11 Storage Node
- JFS - particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 11 Storage Node
- Linux SWAP

Acronis Backup & Recovery 11 can back up and recover corrupted or non-supported file systems using the sector-by-sector approach.

# 1.4   Technical Support

**Maintenance and Support Program**

If you need assistance with your Acronis product, please go to http://www.acronis.com/support/

**Product Updates**

You can download the latest updates for all your registered Acronis software products from our website at any time after logging into your **Account** (https://www.acronis.com/my) and registering the product. See **Registering Acronis Products at the Website** (http://kb.acronis.com/content/4834) and **Acronis Website User Guide** (http://kb.acronis.com/content/8128).

# 2 Getting started

## Step 1. Installation

These brief installation instructions enable you to start using the product quickly. For the complete description of installation methods and procedures, please refer to the Installation documentation.

Before installation, make sure that:

- Your hardware meets the system requirements.
- You have license keys for the edition of your choice.
- You have the setup program. You can download it from the Acronis Web site.

***To install Acronis Backup & Recovery 11***

Run the Acronis Backup & Recovery 11 setup program and follow the on-screen instructions.

## Step 2. Running

Run Acronis Backup & Recovery 11 by selecting **Acronis Backup & Recovery 11** from the **Start** menu.

For understanding of the GUI elements see "Using the management console" (p. 13).

## Step 3. Bootable media

To be able to recover an operating system that fails to start, or deploy it on bare metal, create bootable media.

1. Select **Tools >** **Create bootable media** in the menu.
2. Click **Next** in the welcome screen. Keep clicking **Next** until the list of components appears.
3. Proceed as described in "Linux-based bootable media" (p. 161).

## Step 4. Backup

**Back up now** (p. 29)

Click **Back up now** to do a one-time backup in a few simple steps. The backup process will start immediately after you perform the required steps.

*To save your machine to a file:*

Under **Where to back up**, click **Location**, and select the location where the backup will be saved. Click **OK** to confirm your selection. Click **OK** at the bottom of the window to start the backup.

**Tip.** Using the bootable media, you can do off-line ("cold") backups in the same way as in the operating system.

**Create backup plan** (p. 29)

Create a backup plan if you need a long-term backup strategy including backup schemes, schedules and conditions, timely deleting of backups, or moving them to different locations.

## Step 5. Recovery

**Recover** (p. 97)

To recover data, you need to select the backed up data and the destination the data will be recovered to. As a result, a recovery task will be created.

Recovery of a disk or volume over a volume locked by the operating system requires a reboot. After the recovery is completed, the recovered operating system goes online automatically.

If the machine fails to boot or if you need to recover a system to bare metal, boot the machine using the bootable media and configure the recovery operation in the same way as the recovery task.

## Step 6. Management

The **Navigation** pane (at the left part of the console) enables you to navigate across the product views that are used for different administering purposes.

- Use the **Backup plans and tasks** view to manage backup plans and tasks: run, edit, stop and delete plans and tasks, view their states and progress.

- Use the **Alerts** view to rapidly identify and solve the problems.

- Use the **Log** view to browse the operations log.

- The location where you store backup archives is called a vault (p. 240). Navigate to the **Vaults** (p. 135) view to obtain information about your vaults. Navigate further to the specific vault to view backups and their contents. You can also select the data to recover and perform manual operations with backups (mounting, validating, deleting).

# 2.1    Using the management console

As soon as the console connects to a managed machine (p. 237) or to a management server (p. 237), the respective items appear across the console's workspace (in the menu, in the main area with the **Welcome** screen, or in the **Navigation** pane) enabling you to perform agent-specific or server-specific operations.

**Acronis Backup & Recovery 11 Management Console - Welcome screen**

## Key elements of the console workspace

| | Name | Description |
|---|---|---|
| **1** | **Navigation** pane | Contains the **Navigation** tree and the **Shortcuts** bar. Lets you navigate to the different views. For details, see Navigation pane (p. 14). |
| **2** | Main area | Here you configure and monitor backup, recovery and other operations. The main area displays views and action pages (p. 15) depending on the items selected in the menu or **Navigation** tree. |
| **3** | **Menu** bar | Appears across the top of the program window. Lets you perform most of operations available in Acronis Backup & Recovery 11. The menu items change dynamically depending on the item selected in the **Navigation** tree and the main area. |

# 2.1.1   "Navigation" pane

The navigation pane includes the **Navigation** tree and the **Shortcuts** bar.

## Navigation tree

The **Navigation** tree enables you to navigate across the program views. You can choose between the **Full list** or the **Short list** of views. The **Short list** contains the most frequently used views from the **Full list**.

The **Short list** displays

- ▪ **[Machine name]**. This is the root of the tree also called a **Welcome** screen. It displays the name of the machine the console is currently connected to. Use this view for quick access to the main operations, available on the managed machine.

- ▪ 📋 **Backup plans and tasks**. Use this view to manage backup plans and tasks on the managed machine: run, edit, stop and delete plans and tasks, view their progress.

- ▪ 📁 **Vaults**. Use this view to manage personal vaults and archives stored in there, add new vaults, rename and delete the existing ones, validate vaults, explore backup content, perform operations on archives and backups, etc. If the machine is registered on the management server, you can browse the centralized vaults and perform operations on the archives for which you have the appropriate permissions.

- ▪ ⚡ **Alerts**. Use this view to examine warning messages for the managed machine.

The **Full list** additionally displays

- ▪ 🗑 **Disk management**. Use this view to perform operations on the machine's hard disk drives.

- ▪ 📄 **Log**. Use this view to examine information on operations performed by the program on the managed machine.

- ▪ 💧 **Mounted images**. This node is displayed if at least one volume is mounted. Use this view to manage mounted images.

## Shortcuts bar

The **Shortcuts** bar appears under the navigation tree. It offers you an easy and convenient way of connection to the machines in demand by adding them as shortcuts.

### To add a shortcut to a machine

1. Connect the console to a managed machine.
2. In the navigation tree, right-click the machine's name (a root element of the navigation tree), and then select **Create shortcut**.

   If the console and agent are installed on the same machine, the shortcut to this machine will be added to the shortcuts bar automatically as **Local machine [Machine name]**.

## Operations with pane

### How to expand/minimize panes

By default, the **Navigation** pane appears expanded. You might need to minimize the pane in order to free some additional workspace. To do this, click the chevron (⟨«⟩). The pane will be minimized and the chevron changes its direction (⟨»⟩). Click the chevron once again to expand the pane.

### How to change the panes' borders

1. Point to the pane's border.
2. When the pointer becomes a double-headed arrow, drag the pointer to move the border.

## 2.1.2 Main area, views and action pages

The main area is a basic place where you work with the console. Here you create, edit and manage backup plans, recovery tasks and perform other operations. The main area displays different views and action pages according the items you select in the menu, or **Navigation** tree.

## 2.1.2.1 Views

A view appears on the main area when clicking any item in the **Navigation** tree in the Navigation pane (p. 14).



"Log" view

**Common way of working with views**

Generally, every view contains a table of items, a table toolbar with buttons, and the **Information** panel.

- Use filtering and sorting (p. 16) capabilities to search the table for the item in question.
- In the table, select the desired item.
- In the information panel (collapsed by default), view the item's details. To expand the panel, click the arrow mark ( ▲ ).
- Perform actions on the selected item. There are several ways of performing the same action on selected items:
  - By clicking the buttons on the table toolbar.
  - By selecting the items in the **Actions** menu.
  - By right-clicking the item and selecting the operation in the context menu.

## Sorting, filtering and configuring table items

The following is a guideline to sort, filter and configure table items in any view.

| To | Do the following |
|----|------------------|
| Sort items by any column | Click a column's header to sort items in ascending order.<br>Click it once again to sort items in descending order. |

| Filter items by predefined column value | In a field below the corresponding column's header, select the required value from the drop-down list. |
|---|---|
| Filter items by entered value | In a field below the corresponding column's header, type a value. |
| | As a result you will see the list of values, fully or just partly coincide with the entered value. |
| Filter items by a predefined parameters | Depending on the view, you can filter a table items by some predefined parameters. To do this, click the respective buttons or links at the top of the table. |
| | For example: |
| | ■ In the **Log** view, you can filter the event entries by clicking buttons associated with the result: **Succeeded**, **Succeeded with warnings**, or **Failed**. |
| | ■ The **Log** view has the activity start time as the default parameter, and three predefined settings for filtering activities by this parameter (**All available**, **For last 3 months**, or **For custom period**) are placed at the top of the Log view. |
| Show or hide table columns | By default, any table has a fixed number of columns that are shown, others are hidden. If required, you can hide the shown columns and show the hidden ones. |
| | To show or hide columns |
| | 1. Right-click any column header to open the context menu. |
| | 2. Click the items you want to be displayed/hidden. |

## 2.1.2.2    Action pages

An action page appears in the main area when clicking any action item in the **Actions** menu. It contains steps you need to perform in order to create and launch any task or a backup plan.



**Action page - Create backup plan**

## Using controls and specifying settings

Use active controls to specify a backup plan or recovery task settings and parameters. By default, such fields as credentials, options, comments, and some others are hidden. Most settings are configured by clicking the respective **Show…** links. Others are selected from the drop-down list, or typed manually in the page's fields.



**Action page - Controls**

Acronis Backup & Recovery 11 remembers the changes you made on the action pages. For example, if you started to create a backup plan, and then for any reason switched to another view without accomplishing the plan creation, you can click the **Back** navigation button on the menu. Or, if you have passed several steps forward, click the **Down** arrow and select the page where you started the plan creation from the list. Thus, you can perform the remaining steps and accomplish the backup plan creation.



**Navigation buttons**

## 2.1.3 Console options

The console options define the way information is represented in the Graphical User Interface of Acronis Backup & Recovery 11.

To access the console options, select **Options > Console** options from the top menu.

### 2.1.3.1 Alert display options

The option specifies which alerts to show and which to hide in the **Alerts** view.

The preset is: **All alerts**.

To show (hide) alerts, select (clear) the check boxes next to the respective alert types.

## 2.1.3.2    Credentials cache

The option specifies whether to store the credentials entered while using the management console.

The preset is: **Disabled**.

If the option is disabled, access credentials for various locations that you enter during a console session are stored only until the console is closed.

If the option is enabled, the credentials are saved for use during later sessions. In Windows, the credentials are stored in the Windows Credential Manager. In Linux, the credentials are stored in a special encrypted file.

## 2.1.3.3    Fonts

The option defines the fonts to be used in the Graphical User Interface of Acronis Backup & Recovery 11. The **Menu font** setting affects the drop-down and context menus. The **Application font** setting affects all other GUI elements.

The preset is: **System Default** font for both the menus and the application interface items.

To make a selection, choose the font from the respective combo-box and set the font's properties. You can preview the font's appearance by clicking **Browse** to the right.

## 2.1.3.4    Pop-up messages

These options are effective when the console is connected to a managed machine or to the management server.

**The "Activities Need Interaction" dialog**

This option defines whether to display a pop-up window when one or more activities require user interaction. This window enables you to specify your decision, such as to confirm reboot or to retry after freeing-up the disk space, on all the activities in the same place. Until at least one activity requires interaction, you can open this window at any time from the managed machine's welcome screen. Alternatively, you can review the task execution states in the **Backup plans and tasks** view and specify your decision on each task in the information panel.

The preset is: **Enabled.**

To make a selection, select or clear the **The "Activities Need Interaction" dialog** check box.

**The "Feedback Confirmation" dialog**

This option defines whether to display a pop-up window with the information about your system after an error occurs. You can send this information to Acronis technical support.

The preset is: **Enabled.**

To make a selection, select or clear the **The "Feedback Confirmation" dialog** check box.

**Notify if bootable media is not created**

This option defines whether to display a pop-up window when the management console is launched on a machine and no bootable media has been created on that machine.

The preset is: **Enabled.**

To make a selection, select or clear the **Notify if bootable media is not created** check box.

**Notify when the management console is connected to a component of a different version**

This option defines whether to display a pop-up window when a console is connected to an agent/management server and their versions differ.

The preset is: **Enabled.**

To make a selection, select or clear the **Notify when the management console is connected to a component of a different version** check box.

**About the task execution results**

This option is effective only when the console is connected to a managed machine.

The option defines whether to display the pop-up messages about task run results: successful completion, failure or success with warnings. When the displaying of pop-up messages is disabled, you can review the task execution states and results in the **Backup plans and tasks** view.

The preset is: **Enabled** for all results.

To make a setting for each result (successful completion, failure or success with warnings) individually, select or clear the respective check box.

## 2.1.3.5    Startup page

This option defines whether to show the **Welcome** screen or the **Dashboard** view on the console connection to the management server.

The preset is: the **Welcome** screen.

To make a selection, select or clear the check box for **Show the "Dashboard" view**.

This option can also be set on the **Welcome** screen. If you select the check box for **At startup, show the Dashboard instead of the current view** on the **Welcome** screen, the setting mentioned above will be updated accordingly.

# 3   Understanding Acronis Backup & Recovery 11

This section attempts to give its readers a clear understanding of the product so that they can use the product in various circumstances without step-by-step instructions.

## 3.1   Owners and credentials

This section explains the concept of owner and the meaning of a backup plan's (or task's) credentials.

**Plan (task) owner**

A local backup plan owner is the user who created or last modified the plan.

A centralized backup plan owner is the management server administrator who created or last modified the centralized backup plan.

Tasks, belonging to a backup plan, either local or centralized, are owned by the backup plan owner.

Tasks that do not belong to a backup plan, such as the recovery task, are owned by the user who has created or last modified the task.

**Managing a plan (task) owned by another user**

Having Administrator privileges on the machine, a user can modify tasks and local backup plans owned by any user registered in the operating system.

When a user opens a plan or task for editing, which is owned by another user, all passwords set in the task are cleared. This prevents the "modify settings, leave passwords" trick. The program displays a warning each time you are trying to edit a plan (task) last modified by another user. On seeing the warning, you have two options:

- Click **Cancel** and create your own plan or task. The original task will remain intact.
- Continue editing. You will have to enter all credentials required for the plan or task execution.

**Archive owner**

An archive owner is the user who saved the archive to the destination. To be more precise, this is the user whose account was specified when creating the backup plan in the **Where to back up** step. By default, the plan's credentials are used.

**Plan's credentials and task credentials**

Any task running on a machine runs on behalf of a user. When creating a plan or a task, you have the option to explicitly specify an account under which the plan or the task will run. Your choice depends on whether the plan or task is intended for manual start or for executing on schedule.

**Manual start**

You can skip the **Plan's (Task) credentials** step. Every time you start the task, the task will run under the credentials with which you are currently logged on. Any person that has administrative privileges on the machine can also start the task. The task will run under this person's credentials.

The task will always run under the same credentials, regardless of the user who actually starts the task, if you specify the task credentials explicitly. To do so, on the plan (task) creation page:

1. In the **Plan parameters** (or **Task parameters**) section, click **Show plan's credentials, comments, label** (or **Show task credentials**).

2. Click **Plan's (Task) credentials**.

3. Enter the credentials under which the plan (task) will run.

**Scheduled or postponed start**

The plan (task) credentials are mandatory. If you skip the credentials step, you will be asked for credentials after finishing the plan (task) creation.

**Why does the program compel me to specify credentials?**

A scheduled or postponed task has to run anyway, regardless if any user is logged on or not (for example, the system is at the Windows "Welcome" screen) or a user other than the task owner is logged on. It is sufficient that the machine be on (that is, not in standby or hibernate) at the scheduled task start time. That's why the Acronis scheduler needs the explicitly specified credentials to be able to start the task.

# 3.2    User privileges on a managed machine

When managing a machine running Windows, the scope of a user's management rights depends on the user's privileges on the machine.

**Regular users**

A regular user, such as a member of the Users group, has the following management rights:

- Perform file-level backup and recovery of the files that the user has permissions to access—but without using a file-level backup snapshot.

- Create backup plans and tasks and manage them.

- View—but not manage—backup plans and tasks created by other users.

- View the local event log.

**Administrative users**

A user who has administrative privileges on the machine, such as a member of the Administrators or Backup Operators group, additionally has the following management rights:

- Back up and recover the entire machine or any data on the machine, with or without using a disk snapshot.

Members of the Administrators group also can:

- View and manage backup plans and tasks owned by any user on the machine.

# 3.3    Full, incremental and differential backups

Acronis Backup & Recovery 11 provides the capability to use popular backup schemes, such as Grandfather-Father-Son and Tower of Hanoi, as well as to create custom backup schemes. All backup schemes are based on full, incremental and differential backup methods. The term "scheme" in fact denotes the algorithm of applying these methods plus the algorithm of the archive cleanup.

Comparing backup methods with each other does not make much sense because the methods work as a team in a backup scheme. Each method should play its specific role according to its advantages. A competent backup scheme will benefit from the advantages of all backup methods and lessen the

influence of all the methods' shortcomings. For example, weekly differential backup facilitates archive cleanup because it can be easily deleted along with the weekly set of daily incremental backups depending on it.

Backing up with the full, incremental or differential backup method results in a backup (p. 228) of the corresponding type.

## Full backup

A full backup stores all data selected for backup. A full backup underlies any archive and forms the base for incremental and differential backups. An archive can contain multiple full backups or consist of only full backups. A full backup is self-sufficient - you do not need access to any other backup to recover data from a full backup.

It is widely accepted that a full backup is the slowest to do but the fastest to restore. With Acronis technologies, recovery from an incremental backup may be not slower than recovery from a full one.

A full backup is most useful when:

- you need to roll back the system to its initial state
- this initial state does not change often, so there is no need for regular backup.

Example: An Internet cafe, school or university lab where the administrator often undoes changes made by the students or guests but rarely updates the reference backup (in fact, after installing software updates only). The backup time is not crucial in this case and the recovery time will be minimal when recovering the systems from the full backup. The administrator can have several copies of the full backup for additional reliability.

## Incremental backup

An incremental backup stores changes to the data against the **latest backup**. You need access to other backups from the same archive to recover data from an incremental backup.

An incremental backup is most useful when:

- you need the possibility to roll back to any one of multiple saved states
- the data changes tend to be small as compared to the total data size.

It is widely accepted that incremental backups are less reliable than full ones because if one backup in the "chain" is corrupted, the next ones can no longer be used. However, storing multiple full backups is not an option when you need multiple prior versions of your data, because reliability of an oversized archive is even more questionable.

Example: Backing up a database transaction log.

## Differential backup

A differential backup stores changes to the data against the **latest full backup**. You need access to the corresponding full backup to recover the data from a differential backup. A differential backup is most useful when:

- you are interested in saving only the most recent data state
- the data changes tend to be small as compared to the total data size.

The typical conclusion is: "differential backups take longer to do and are faster to restore, while incremental ones are quicker to do and take longer to restore." In fact, there is no physical difference between an incremental backup appended to a full backup and a differential backup appended to

the same full backup at the same point of time. The above mentioned difference implies creating a differential backup after (or instead of) creating multiple incremental backups.

*An incremental or differential backup created after disk defragmentation might be considerably larger than usual because defragmentation changes file locations on the disk and the backup reflects these changes. It is recommended that you re-create a full backup after disk defragmentation.*

The following table summarizes the advantages and shortcomings of each backup type as they appear based on common knowledge. In real life, these parameters depend on numerous factors such as the amount, speed and pattern of data changes; the nature of the data, the physical specifications of the devices, the backup/recovery options you set, to name a few. Practice is the best guide to selecting the optimal backup scheme.

| Parameter | Full backup | Differential backup | Incremental backup |
|---|---|---|---|
| Storage space | Maximal | Medium | Minimal |
| Creation time | Maximal | Medium | Minimal |
| Recovery time | Minimal | Medium | Maximal |

# 3.4   What does a disk or volume backup store?

A disk or volume backup stores a disk or a volume file system as a whole, along with all the information necessary for the operating system to boot. It is possible to recover disks or volumes as a whole from such backup, as well as individual folders or files.

With the sector-by-sector (raw mode) option enabled, a disk backup stores all the disk sectors.

For supported file systems, with the sector-by-sector option turned off, a disk or volume backup stores only those sectors that contain data. This reduces the resulting backup size and speeds up the backup and recovery operations.

**Windows**

The swap file (pagefile.sys) and the file that keeps the RAM content when the machine goes into hibernation (hiberfil.sys) are not backed up. After recovery, the files will be re-created in the appropriate place with the zero size.

A volume backup stores all other files and folders of the selected volume independent of their attributes (including hidden and system files), the boot record, the file allocation table (FAT) if it exists, the root and the zero track of the hard disk with the master boot record (MBR). The boot code of GPT volumes is not backed up.

A disk backup stores all volumes of the selected disk (including hidden volumes such as the vendor's maintenance partitions) and the zero track with the master boot record.

**Linux**

A volume backup stores all files and folders of the selected volume independent of their attributes; a boot record and the file system super block.

A disk backup stores all disk volumes as well as the zero track with the master boot record.

# 3.5 Backup and recovery of dynamic volumes (Windows)

This section explains in brief how to back up and recover dynamic volumes (p. 234) using Acronis Backup & Recovery 11.

A dynamic volume is a volume located on dynamic disks (p. 233), or more exactly, on a disk group (p. 233). Acronis Backup & Recovery 11 supports the following dynamic volume types/RAID levels:

- simple/spanned
- striped (RAID 0)
- mirrored (RAID 1)
- a mirror of stripes (RAID 0+1)
- RAID-5.

## Backing up dynamic volumes

Dynamic volumes are backed up in the same way as basic volumes. When creating a backup plan through the GUI, all types of volumes are available for selection as **Items to back up**. When using the command line, specify the dynamic volumes with the DYN prefix.

**Command line examples**

```
acrocmd backup disk --volume=DYN1,DYN2 --loc=\\srv1\backups
--credentials=netuser1,pass1 --arc=dyn1_2_arc
```

This will back up volumes DYN1 and DYN2 to a network shared folder.

```
acrocmd backup disk --volume=DYN --loc=\\srv1\backups
--credentials=netuser1,pass1 --arc=alldyn_arc
```

This will back up all dynamic volumes of the local machine to a network shared folder.

## Recovering dynamic volumes

A dynamic volume can be recovered:

- Over any type of existing volume.
- To unallocated space of a disk group.
- To unallocated space of a basic disk.
- To a disk which has not been initialized.

**Recovery over an existing volume**

When a dynamic volume is recovered over an existing volume, either basic or dynamic, the target volume's data is overwritten with the backup content. The type of target volume (basic, simple/spanned, striped, mirrored, RAID 0+1, RAID-5) will not change. The target volume size has to be enough to accommodate the backup content.

**Recovery to disk group unallocated space**

When recovering a dynamic volume to disk group unallocated space, the software preserves the volume's original type and size. If the disk group configuration does not allow for the original volume type, the volume will be recovered as a simple or spanned volume. If this volume does not fit the unallocated space, the volume will be resized by decreasing its free space.

**Examples of when the disk group configuration does not allow the original type of the volume**

**Example 1**. The group contains fewer disks than is required for the dynamic volume. Assume you are going to recover an 80 GB RAID-5 volume that had resided on three disks, to a disk group consisting of two disks. The total size of unallocated space is 100 GB: 40 GB on the first disk and 60 GB on the second. The RAID-5 volume will be recovered as a spanned volume across two disks.

**Example 2**. Unallocated space distribution does not allow recovery of certain types of dynamic volumes. Assume you are going to recover a 30 GB striped volume to a disk group consisting of two disks. The total size of unallocated space is 50 GB: 10 GB on the first disk and 40 GB on the second. The striped volume will be recovered to the second disk as simple.

**Recovery to a disk that has not been initialized**

In this case, the target disk will be automatically initialized to the MBR partitioning style. The dynamic volumes will be recovered as basic ones. If the volumes cannot fit into unallocated space, they will be proportionally resized (by decreasing their free space).

The table below demonstrates the resulting volume types depending on the backed up source and the recovery target.

| | Backup (source): | |
|---|---|---|
| **Recovered to:** | **Dynamic volume** | **Basic volume** |
| **Dynamic volume** | Dynamic volume Type as of the target | Dynamic volume Type as of the target |
| **Unallocated space (disk group)** | Dynamic volume Type as of the source | Dynamic volume Simple |
| **Basic volume or unallocated space on a basic disk** | Basic volume | Basic volume |

**Moving and resizing volumes during recovery**

You can manually resize the resulting basic volume during recovery, or change the volume's location on the disk. A resulting dynamic volume cannot be moved or resized manually.

## Preparing disk groups and volumes

Before recovering dynamic volumes to bare metal you should create a disk group on the target hardware.

You also might need to create or increase unallocated space on an existing disk group. This can be done by deleting volumes or converting basic disks to dynamic.

You might want to change the target volume type (basic, simple/spanned, striped, mirrored, RAID 0+1, RAID 5). This can be done by deleting the target volume and creating a new volume on the resulting unallocated space.

Acronis Backup & Recovery 11 includes a handy disk management utility which enables you to perform the above operations both under the operating system and on bare metal. To find out more about Acronis Disk Director Lite, see the Disk management (p. 172) section.

# 3.6   Compatibility with encryption software

Acronis Backup & Recovery 11 fully retains its functionality when interacting with file-level encryption software.

Disk-level encryption software encrypts data on the fly. This is why data contained in the backup is not encrypted. Disk-level encryption software often modifies system areas: boot records, or partition tables, or file system tables. These factors affect disk-level backup and recovery, the ability of the recovered system to boot and access to Acronis Secure Zone.

Under some conditions, Acronis Backup & Recovery 11 is compatible with the following disk-level encryption software:

- Microsoft BitLocker Drive Encryption
- McAfee Endpoint Encryption
- PGP Whole Disk Encryption.

To ensure reliable disk-level recovery, follow the common rules and software-specific recommendations.

## Common installation rule

The strong recommendation is to install the encryption software before installing Acronis Backup & Recovery 11.

## The way of using Acronis Secure Zone

Acronis Secure Zone must not be encrypted with disk-level encryption. This is the only way to use Acronis Secure Zone:

1. Install encryption software; then, install Acronis Backup & Recovery 11.
2. Create Acronis Secure Zone.
3. Exclude Acronis Secure Zone when encrypting the disk or its volumes.

## Common backup rule

You can do a disk-level backup in the operating system. Do not try to back up using bootable media or Acronis Startup Recovery Manager.

## Software-specific recovery procedures

### Microsoft BitLocker Drive Encryption

To recover a system that was encrypted by BitLocker:

1. Boot from the bootable media.
2. Recover the system. The recovered data will be unencrypted.
3. Reboot the recovered system.
4. Turn on BitLocker.

If you only need to recover one partition of a multi-partitioned disk, do so under the operating system. Recovery under bootable media may make the recovered partition undetectable for Windows.

### McAfee Endpoint Encryption and PGP Whole Disk Encryption

You can recover an encrypted system partition by using bootable media only.

If the recovered system fails to boot, rebuild Master Boot Record as described in the following Acronis knowledge base article: "Restoring Windows Boot Loader Manually" and reboot.

# 3.7    Support for SNMP

**SNMP objects**

Acronis Backup & Recovery 11 provides the following Simple Network Management Protocol (SNMP) objects to SNMP management applications:

- Type of event

  Object identifier (OID): 1.3.6.1.4.1.24769.100.200.1.0

  Syntax: OctetString

  The value may be "Information", "Warning", 'Error" and "Unknown". "Unknown" is sent only in the test message.

- Text description of the event

  Object identifier (OID): 1.3.6.1.4.1.24769.100.200.2.0

  Syntax: OctetString

  The value contains the text description of the event (it looks identical to messages published by Acronis Backup & Recovery 11 in its log).

**Example of varbind values:**

1.3.6.1.4.1.24769.100.200.1.0:Information

1.3.6.1.4.1.24769.100.200.2.0:I0064000B

**Supported operations**

Acronis Backup & Recovery 11 **supports only TRAP operations**. It is not possible to manage Acronis Backup & Recovery 11 using GET- and SET- requests. This means that you need to use an SNMP Trap receiver to receive TRAP-messages.

**About the management information base (MIB)**

The MIB file **acronis-abr.mib** is located in the Acronis Backup & Recovery 11 installation directory. By default: %ProgramFiles%\Acronis\BackupAndRecovery in Windows and /usr/lib/Acronis/BackupAndRecovery in Linux.

This file can be read by a MIB browser or a simple text editor such as Notepad or vi.

**About the test message**

When configuring SNMP notifications, you can send a test message to check if your settings are correct.

The parameters of the test message are as follows:

- Type of event
  OID: 1.3.6.1.4.1.24769.100.200.1.0
  Value: "Unknown"
- Text description of the event
  OID: 1.3.6.1.4.1.24769.100.200.2.0
  Value: "?00000000"

# 4 Backup

## 4.1 Back up now

Use the **Back up now** feature to configure and run a one-time backup in a few simple steps. The backup process will start immediately after you perform the required steps and click **OK**.

For a long-time backup strategy that includes schedules and conditions, timely deleting of backups or moving them to different locations, consider creating a backup plan.

Configuring immediate backup is similar to creating a backup plan (p. 29) except for the following:

- There are no options to schedule backups and to set up retention rules.
- Simplified naming of backup files (p. 48) is used, if the backup destination supports it. Otherwise, the standard backup naming is used.

    The following locations do not support simplified file naming: managed vaults, tape, Acronis Secure Zone or Acronis Online Backup Storage.

- Conversion of a disk-level backup to a virtual machine is not available as a part of the backup operation. You can convert the resulting backup afterwards.

## 4.2 Creating a backup plan

Before creating your first backup plan (p. 229), please familiarize yourself with the basic concepts used in Acronis Backup & Recovery 11.

***To create a backup plan, perform the following steps.***

### What to back up

**Items to back up** (p. 31)

Select the type of data to back up and specify the data items. The type of data depends on the agents installed on the machine.

**Access credentials, exclusions**

To access these settings, click **Show access credentials, exclusions**.

**Access credentials** (p. 32)

Provide credentials for the source data if the plan's account does not have access permissions to the data.

**Exclusions** (p. 32)

Set up exclusions for the specific types of files you do not wish to back up.

### Where to back up

**Location** (p. 44)

Specify a path to the location where the backup archive will be stored and the archive name. The archive name has to be unique within the location. Otherwise, backups of the newly created backup plan will be placed to the existing archive that belongs to another backup plan. The default archive name is Archive(N) where N is the sequence number of the archive in the location you have selected.

**Backup file naming, access credentials, archive comments**

To access these settings, click **Show backup file naming, access credentials, archive comments**.

**File naming** (p. 48)

[Optional] Select the **Name backup files using the archive name, as in Acronis True Image Echo, rather than auto-generated names** check box if you want to use simplified file naming for the archive's backups.

*Not available when backing up to a managed vault, tape, Acronis Secure Zone or Acronis Online Backup Storage.*

**Access credentials** (p. 34)

[Optional] Provide credentials for the location if the plan account does not have access permissions to the location.

**Archive comments**

[Optional] Enter comments on the archive.

## How to back up

**Backup scheme** (p. 34)

Specify when and how often to back up your data; define for how long to keep the created backup archives in the selected location; set up schedule for the archive cleanup procedure (see "Replication and retention settings" below). Use well-known optimized backup schemes, such as Grandfather-Father-Son and Tower of Hanoi; create a custom backup scheme, or back up data once.

Replication and retention settings (p. 66)

*Not available when choosing simplified naming of backup files (p. 48).*

Define whether to copy (replicate) the backups to another location, and whether to move or delete them according to retention rules. The available settings depend on the backup scheme.

**2nd location, validation, convert to virtual machine**

To access these settings, click **Show 2nd location, validation, convert to virtual machine**.

**2nd location**

[Optional] To set up replication of backups, select the **Replicate just created backup to another location** check box. For more information about backup replication, see Setting up replication of backups (p. 68).

**When to validate** (p. 46)

[Optional] Depending on the selected backup scheme, define when and how often to perform validation and whether to validate the entire archive or the latest backup in the archive.

**Convert to virtual machine** (p. 72)

[Optional] Applies to: disk or volume backup, backup of entire virtual machines or volumes of a virtual machine.

Set up a regular conversion of a disk or volume backup to a virtual machine.

## Plan parameters

**Plan name**

[Optional] Enter a unique name for the backup plan. A conscious name lets you identify the plan among others.

**Backup options**

[Optional] Configure parameters of the backup operation, such as pre/post backup commands, maximum network bandwidth allocated for the backup stream or the backup archive compression level. If you do nothing in this section, the default values (p. 75) will be used.

After any of the settings is changed against the default value, a new line that displays the newly set value appears. The setting status changes from **Default** to **Reset to default**. Should you modify the setting again, the line will display the new value unless the new value is the default one. When the default value is set, the line disappears. Therefore, in this section you always see only the settings that differ from the default values.

To reset all the settings to the default values, click **Reset to default**.

**Plan's credentials, comments, label**

To access these settings, click **Show plan's credentials, comments, label**.

**Plan's credentials** (p. 46)

[Optional] The backup plan will run on behalf of the user who is creating the plan. You can change the plan's credentials if necessary.

**Comments**

[Optional] Type a description of the backup plan.

**Label** (p. 47)

[Optional] Type a text label for the machine you are going to back up. The label can be used to identify the machine in various scenarios.

After you have performed all the required steps, click **OK** to create the backup plan.

After that, you might be prompted for the password (p. 48).

The plan you have created will be accessible for examination and managing in the **Backup plans and tasks** (p. 188) view.

## 4.2.1 Selecting data to back up

***To select the data to back up***

1. In the **Data to back up** section, select the type of data you want to be backed up. The list of available data types depends on the agents running on the machine and the types of licenses:

   **Disks/volumes**

   Available if Acronis Backup & Recovery 11 Agent for Windows or Acronis Backup & Recovery 11 Agent for Linux is installed.

   Select this option to back up entire physical machine or its individual disks or volumes. To be able to back up this data, you must have Administrator or Backup operator privileges.

   A disk-level backup enables you to recover the entire system in case of severe data damage or hardware failure. The backup procedure is faster than copying files, and may significantly speed up the backup process when backing up large volumes of data.

   **Folders/files**

   Available if Acronis Backup & Recovery 11 Agent for Windows or Acronis Backup & Recovery 11 for Linux is installed.

   Select this option to back up specific files and folders.

   A file-level backup is not sufficient for recovery of the operating system. Choose file backup if you plan to keep safe only certain data (the current project, for example). This will reduce the archive size, thus saving storage space.

In order to recover your operating system along with all the settings and applications, you have to perform a disk backup.

2. In the tree below the **Data to back up** section, select the items to back up by selecting check boxes next to the items.

    Selecting a check box for a machine means backing up all the machine's disks. To select individual disks and/or volumes, expand the machine item and select check boxes next to the disks and/or volumes.

    **Note for Disks/volumes**

    ▪ If your operating system and its loader reside on different volumes, always include both volumes in the backup. The volumes must also be recovered together; otherwise there is a high risk that the operating system will not start.

3. Having specified the data to backup, click **OK**.

## 4.2.2    Access credentials for source

Specify the credentials required for access to the data you are going to backup.

***To specify credentials***

1. Select one of the following:

    ▪ **Use the plan's credentials**

       The program will access the source data using the credentials of the backup plan account specified in the **Plan parameters** section.

    ▪ **Use the following credentials**

       The program will access the source data using the credentials you specify.

       Use this option if the plan's account does not have access permissions to the data.

       Specify:

       ▪ **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).

       ▪ **Password**. The password for the account.

2. Click **OK**.

## 4.2.3    Source files exclusion

This option is effective for Windows and Linux operating systems and bootable media.

This option is effective for disk-level backup of NTFS and FAT file systems only. This option is effective for file-level backup of all supported file systems.

The option defines which files and folders to skip during the backup process and thus exclude from the list of backed-up items.

The preset is: **Exclude files matching the following criteria: *.tmp, *.~, *.bak.**

***To specify which files and folders to exclude:***

Set up any of the following parameters:

▪ **Exclude all hidden files and folders**

This option is effective only for file systems that are supported by Windows. Select this check box to skip files and folders with the **Hidden** attribute. If a folder is **Hidden**, all of its contents — including files that are not **Hidden** — will be excluded.

- **Exclude all system files and folders**

This option is effective only for file systems that are supported by Windows. Select this check box to skip files and folders with the **System** attribute. If a folder is **System**, all of its contents — including files that are not **System** — will be excluded.

*You can view file or folder attributes in the file/folder properties or by using the **attrib** command. For more information, refer to the Help and Support Center in Windows.*

- **Exclude files matching the following criteria**

Select this check box to skip files and folders whose names match any of the criteria — called file masks — in the list; use the **Add**, **Edit**, **Remove** and **Remove All** buttons to create the list of file masks.

You can use one or more wildcard characters * and ? in a file mask:

The asterisk (*) substitutes for zero or more characters in a file name; for example, the file mask Doc*.txt yields files such as Doc.txt and Document.txt

The question mark (?) substitutes for exactly one character in a file name; for example, the file mask Doc?.txt yields files such as Doc1.txt and Docs.txt — but not the files Doc.txt or Doc11.txt

To exclude a folder specified by a path containing the drive letter, add a backslash (\) to the folder name in the criterion; for example: C:\Finance\

**Exclusion examples**

| Criterion | Example | Description |
|---|---|---|
| **Windows and Linux** | | |
| By name | F.log | Excludes all files named "F.log" |
| | F | Excludes all folders named "F" |
| By mask (*) | *.log | Excludes all files with the .log extension |
| | F* | Excludes all files and folders with names starting with "F" (such as folders F, F1 and files F.log, F1.log) |
| By mask (?) | F???.log | Excludes all .log files with names consisting of four symbols and starting with "F" |
| **Windows** | | |
| By file path | C:\Finance\F.log | Excludes the file named "F.log" located in the folder C:\Finance |
| By folder path | C:\Finance\F\ | Excludes the folder C:\Finance\F (be sure to specify the full path starting from the disk letter) |
| **Linux** | | |
| By file path | /home/user/Finance/F.log | Excludes the file named "F.log" located in the folder /home/user/Finance |
| By folder path | /home/user/Finance/ | Excludes the folder /home/user/Finance |

The above settings are not effective for the files or folders that were explicitly selected for backup. For example, assume that you selected the folder MyFolder and the file MyFile.tmp outside that folder, and selected to skip all .tmp files. In this case, all .tmp files in the folder MyFolder will be skipped during the backup process, but the file MyFile.tmp will not be skipped.

## 4.2.4 Access credentials for archive location

Specify credentials required for access to the location where the backup archive will be stored. The user whose name is specified will be considered as the archive owner.

**To specify credentials**

1. Select one of the following:

   ▪ **Use the plan's credentials**

   The program will access the source data using the credentials of the backup plan account specified in the **Plan parameters** section.

   ▪ **Use the following credentials**

   The program will access the source data using the credentials you specify.

   Use this option if the plan account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

   Specify:

   ▪ **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).

   ▪ **Password**. The password for the account.

2. Click **OK**.

*Warning: According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

## 4.2.5 Backup schemes

Choose one of the available backup schemes:

▪ **Simple** – to schedule when and how often to backup data and specify retention rules.

▪ **Run now** - to perform the backup immediately right after you click the **OK** button.

▪ **Grandfather-Father-Son** – to use the Grandfather-Father-Son backup scheme. The scheme does not allow data to be backed up more than once a day. You set the days of week when the daily backup will be performed and select from these days the day of weekly/monthly backup. Then you set the retention periods for the daily (referred to as "sons"), weekly (referred to as "fathers") and monthly (referred to as "grandfathers") backups. The expired backups will be deleted automatically.

▪ **Tower of Hanoi** – to use the Tower of Hanoi backup scheme. This scheme allows you to schedule when and how often to back up (sessions) and select the number of backup levels (up to 16). The data can be backed up more than once a day. By setting up the backup schedule and selecting backup levels, you automatically obtain the rollback period – the guaranteed number of sessions that you can go back at any time. The automatic cleanup mechanism maintains the required rollback period by deleting the expired backups and keeping the most recent backups of each level.

- **Custom** – to create a custom scheme, where you are free to set up a backup strategy in the way your enterprise needs it most: specify multiple schedules for different backup types, add conditions and specify the retention rules.
- **Manual start** – to create a backup task for manual start OR schedule one-time task execution in the future.
- **Initial seeding** - to save locally a full backup whose final destination is Acronis Online Backup Storage.

## 4.2.5.1    Simple scheme

With the simple backup scheme, you just schedule when and how often to back up data. Other steps are optional.

To set up the simple backup scheme, specify the appropriate settings as follows.

**Schedule**

Set up when and how often to back up the data. To learn more about setting up the schedule, see the Scheduling (p. 53) section.

**Retention rules**

Specify how long to store backups in the location and whether to move or delete them afterward. The retention rules are applied after creating a backup. The **Keep backups indefinitely** is set by default, which means that no backups will be deleted automatically. For more information about retention rules, see Setting up retention of backups (p. 68).

**Backup type**

To access this setting, click **Show backup type, 2nd location, validation, convert to virtual machine**.
Select the backup type.

- **Full** - selected by default for all backup locations (except for Acronis Online Backup Storage).
- **Incremental**. At the first time a full backup will be created. The next backups will be incremental. Selected as the one and only backup type for Acronis Online Backup Storage.

*Note: When the **Incremental** backup type is selected along with retention rules, the archive will be cleaned up using consolidation (p. 231), which is a more time-consuming and resource-intensive operation.*

## 4.2.5.2    Run now scheme

With the **Run now** scheme, the backup will be performed immediately, right after you click the **OK** button at the bottom of the **Create Backup Plan** page.

In the **Backup type** field, select whether you want to create a full, incremental or differential backup (p. 22).

## 4.2.5.3    Grandfather-Father-Son scheme

**At a glance**

- Daily ("Son") incremental, weekly ("Father") differential, and monthly ("Grandfather") backups
- Custom day for weekly and monthly backups

- Custom retention periods for backups of each type

## Description

Let us suppose that we want to set up a backup plan that will regularly produce a series of daily (D), weekly (W), and monthly (M) backups. Here is a natural way to do this: the following table shows a sample two-month period for such a plan.

| | Mo | Tu | We | Th | Fr | Sa | Su |
|---|---|---|---|---|---|---|---|
| Jan 1—Jan 7 | D | D | D | D | W | - | - |
| Jan 8—Jan 14 | D | D | D | D | W | - | - |
| Jan 15—Jan 21 | D | D | D | D | W | - | - |
| Jan 22—Jan 28 | D | D | D | D | M | - | - |
| Jan 29—Feb 4 | D | D | D | D | W | - | - |
| Feb 5—Feb 11 | D | D | D | D | W | - | - |
| Feb 12—Feb 18 | D | D | D | D | W | - | - |
| Feb 19—Feb 25 | D | D | D | D | M | - | - |
| Feb 26—Mar 4 | D | D | D | D | W | - | - |

Daily backups run every workday except Friday, which is left for weekly and monthly backups. Monthly backups run every fourth Friday, and weekly backups run on all other Fridays.

## Parameters

You can set up the following parameters of a Grandfather-Father-Son (GFS) scheme.

| | |
|---|---|
| **Start backup at** | Specifies when to start a backup. The default value is 12:00 PM. |
| **Back up on** | Specifies the days on which to perform a backup. The default value is Workdays. |
| **Weekly/Monthly** | Specifies which of the days selected in the **Back up on** field you want to reserve for weekly and monthly backups. A monthly backup will be performed every fourth such day. The default value is Friday. |
| **Keep backups** | Specifies how long you want the backups to be stored in the archive. A term can be set in hours, days, weeks, months, or years. For monthly backups, you can also select **Keep indefinitely** if you want them to be saved forever.<br><br>The default values for each backup type are as follows.<br><br>Daily: 5 days (recommended minimum)<br><br>Weekly: 7 weeks<br><br>Monthly: indefinitely<br><br>The retention period for weekly backups must exceed that for daily backups; the monthly backups' retention period must be greater than the weekly backups' retention period.<br><br>We recommend setting a retention period of at least one week for daily backups. |

| Backup type | Specifies the types of daily, weekly and monthly backups |
|---|---|
| | ■ **Always full** - all the daily, weekly and monthly backups will be always full. This is the default selection for cases when a tape drive is selected as a backup location. |
| | ■ **Full/Differential/Incremental** - daily backups are incremental, weekly backups are differential, and monthly backups are full. |
| **Advanced settings** | Available only for advanced editions of Acronis Backup & Recovery 11 when creating a centralized backup plan. See the "Advanced scheduling settings" section for details. |

*A backup is not deleted until all backups that directly depend on it become subject to deletion as well. This is why you might see a weekly or a monthly backup remain in the archive for a few days past its expected expiration date.*

*If the schedule starts with a daily or a weekly backup, a full backup is created instead.*

## Examples

### Each day of the past week, each week of the past month

Let us consider a GFS backup scheme that many may find useful.

■ Back up files every day, including weekends

■ Be able to recover files as of any date over the past seven days

■ Have access to weekly backups of the past month

■ Keep monthly backups indefinitely.

Backup scheme parameters can then be set up as follows.

■ Start backup at: **11:00 PM**

■ Back up on: **All days**

■ Weekly/monthly: **Saturday** (for example)

■ Keep backups:

   ■ Daily: **1 week**

   ■ Weekly: **1 month**

   ■ Monthly: **indefinitely**

As a result, an archive of daily, weekly, and monthly backups will be created. Daily backups will be available for seven days since creation. For instance, a daily backup of Sunday, January 1, will be available through next Sunday, January 8; the first weekly backup, the one of Saturday, January 7, will be stored on the system until February 7. Monthly backups will never be deleted.

### Limited storage

If you do not want to arrange a vast amount of space to store a huge archive, you may set up a GFS scheme so as to make your backups more short-lived, at the same time ensuring that your information can be recovered in case of an accidental data loss.

Suppose that you need to:

■ Perform backups at the end of each working day

■ Be able to recover an accidentally deleted or inadvertently modified file if this has been discovered relatively quickly

■ Have access to a weekly backup for 10 days after it was created

- Keep monthly backups for half a year.

Backup scheme parameters can then be set up as follows.

- Start backup at: **6:00 PM**
- Back up on: **Workdays**
- Weekly/monthly: **Friday**
- Keep backups:
    - Daily: **1 week**
    - Weekly: **10 days**
    - Monthly: **6 months**

With this scheme, you will have a week to recover a previous version of a damaged file from a daily backup; as well as 10-day access to weekly backups. Each monthly full backup will be available for six months since the creation date.

**Work schedule**

Suppose you are a part-time financial consultant and work in a company on Tuesdays and Thursdays. On these days, you often make changes to your financial documents, statements, and update the spreadsheets etc. on your laptop. To back up this data, you may want to:

- Track changes to the financial statements, spreadsheets, etc. performed on Tuesdays and Thursdays (daily incremental backup).
- Have a weekly summary of file changes since last month (Friday weekly differential backup).
- Have a monthly full backup of your files.

Moreover, assume that you want to retain access to all backups, including the daily ones, for at least six months.

The following GFS scheme suits such purposes:

- Start backup at: **11:30 PM**
- Back up on: **Tuesday, Thursday, Friday**
- Weekly/monthly: **Friday**
- Keep backups:
    - Daily: **6 months**
    - Weekly: **6 months**
    - Monthly: **5 years**

Here, daily incremental backups will be created on Tuesdays and Thursdays, with weekly and monthly backups performed on Fridays. Note that, in order to choose **Friday** in the **Weekly/monthly** field, you need to first select it in the **Back up on** field.

Such an archive would allow you to compare your financial documents as of the first and the last day of work, and have a five-year history of all documents, etc.

**No daily backups**

Consider a more exotic GFS scheme:

- Start backup at: **12:00 PM**
- Back up on: **Friday**

- Weekly/monthly: **Friday**
- Keep backups:
    - Daily: **1 week**
    - Weekly: **1 month**
    - Monthly: **indefinitely**

Backup is thus performed only on Fridays. This makes Friday the only choice for weekly and monthly backups, leaving no other date for daily backups. The resulting "Grandfather-Father" archive will hence consist only of weekly differential and monthly full backups.

Even though it is possible to use GFS to create such an archive, the Custom scheme is more flexible in this situation.

## 4.2.5.4  Custom backup scheme

**At a glance**

- Custom schedule and conditions for backups of each type
- Custom schedule and retention rules

**Parameters**

| Parameter | Meaning |
|---|---|
| **Full backup schedule** | Specifies on what schedule and under which conditions to perform a full backup. |
| | For example, the full backup can be set up to run every Sunday at 1:00 AM as soon as all users are logged off. |
| **Incremental backup schedule** | Specifies on what schedule and under which conditions to perform an incremental backup. |
| | If the archive contains no backups at the time of the task run, a full backup is created instead of the incremental backup. |
| **Differential backup schedule** | Specifies on what schedule and under which conditions to perform a differential backup. |
| | If the archive contains no full backups at the time of the task run, a full backup is created instead of the differential backup. |
| **Clean up archive** | Specifies how to get rid of old backups: either to apply retention rules (p. 69) regularly or clean up the archive during a backup when the destination location runs out of space. |
| | By default, the retention rules are not specified, which means older backups will not be deleted automatically. |
| | **Using retention rules** |
| | Specify the retention rules and when to apply them. |
| | This setting is recommended for backup destinations such as shared folders or centralized vaults. |
| | **When there is insufficient space while backing up** |
| | The archive will be cleaned up only during backup and only if there is not enough space to create a new backup. In this case, the software will act as follows: |
| | - Delete the oldest full backup with all dependent incremental/differential |

| | backups |
|---|---|
| | ▪ If there is only one full backup left and a full backup is in progress, then delete the last full backup with all dependent incremental/differential backups |
| | ▪ If there is only one full backup left, and an incremental or differential backup is in progress, an error occurs saying there is a lack of available space |
| | This setting is recommended when backing up to a USB drive or Acronis Secure Zone. This setting is not applicable to managed vaults, FTP and SFTP servers. |
| | This setting enables deletion of the last backup in the archive, in case your storage device cannot accommodate more than one backup. However, you might end up with no backups if the program is not able to create the new backup for some reason. |
| **Apply retention rules**<br><br>(only if the retention rules are set) | Specifies when to apply the retention rules (p. 69).<br><br>For example, the cleanup procedure can be set up to run after each backup, and also on schedule.<br><br>This option is available only if you have set at least one retention rule in **Retention rules**. |
| Cleanup schedule<br><br>(only if **On schedule** is selected) | Specifies a schedule for archive cleanup.<br><br>For example, the cleanup can be scheduled to start on the last day of each month.<br><br>This option is available only if you selected **On schedule** in **Apply retention rules**. |
| **2nd location**, **3rd location**, and so on | Specifies where to copy or move (p. 66) the backups from the current location.<br><br>This option is available only if you selected either the **Replicate just created backup to another location** check box under **How to back up**, or **Move the oldest backups to another location** in the **Retention rules** window. |

## Examples

### Weekly full backup

The following scheme yields a full backup performed every Friday night.

**Full backup**: Schedule: **Weekly**, every **Friday**, at **10:00 PM**

Here, all parameters except **Schedule** in **Full backup** are left empty. All backups in the archive are kept indefinitely (no archive cleanup is performed).

### Full and incremental backup plus cleanup

With the following scheme, the archive will consist of weekly full backups and daily incremental backups. We further require that a full backup begin only after all users have logged off.

**Full backup: Schedule: Weekly**, every **Friday**, at **10:00 PM**

**Full backup: Conditions: User is logged off**

**Incremental: Schedule: Weekly, every workday**, at **9:00 PM**

Also, let all backups older than one year be deleted from the archive, and let the cleanup be performed upon creating a new backup.

**Retention rules:** Delete backups older than **12 months**

**Apply the rules: After backing up**

By default, a one-year-old full backup will not be deleted until all incremental backups that depend on it become subject to deletion too. For more information, see Retention rules (p. 69).

**Monthly full, weekly differential, and daily incremental backups plus cleanup**

This example demonstrates the use of all options available in the Custom scheme.

Suppose that we need a scheme that will produce monthly full backups, weekly differential backups, and daily incremental backups. Then the backup schedule can look as follows.

**Full backup: Schedule: Monthly**, every **Last Sunday** of the month, at **9:00 PM**

**Incremental: Schedule: Weekly**, every **workday**, at **7:00 PM**

**Differential: Schedule: Weekly**, every **Saturday**, at **8:00 PM**

Further, we want to add conditions that have to be satisfied for a backup task to start. This is set up in the **Conditions** fields for each backup type.

**Full backup**: **Conditions**: **Location available**

**Incremental**: **Conditions**: **User is logged off**

**Differential**: **Conditions**: **User is idle**

As a result, a full backup—originally scheduled at 9:00 PM—may actually start later: as soon as the backup location becomes available. Likewise, backup tasks for incremental and differential backups will wait until all users are logged off and users are idle, respectively.

Finally, we create retention rules for the archive: let us retain only backups that are no older than six months, and let the cleanup be performed after each backup task and also on the last day of every month.

**Retention rules**: Delete backups older than **6 months**

**Apply the rules**: **After backing up**, **On schedule**

**Cleanup schedule**: **Monthly**, on the **Last day** of **All months**, at **10:00 PM**

By default, a backup is not deleted as long as it has dependent backups that must be kept. For example, if a full backup has become subject to deletion, but there are incremental or differential backups that depend on it, the deletion is postponed until all the dependent backups can be deleted as well.

For more information, see Retention rules (p. 69).

## 4.2.5.5    Tower of Hanoi scheme

**At a glance**

- Up to 16 levels of full, differential, and incremental backups
- Next-level backups are twice as rare as previous-level backups
- One backup of each level is stored at a time

- Higher density of more recent backups

## Parameters

You can set up the following parameters of a Tower of Hanoi scheme.

| Schedule | Set up a daily (p. 54), weekly (p. 56), or monthly (p. 58) schedule. Setting up schedule parameters allows for the creation of simple schedules (example of a simple daily schedule: a backup task will be run every 1 day at 10 AM) as well as more complex schedules (example of a complex daily schedule: a task will be run every 3 days, starting from January 15. During the specified days the task will be repeated every 2 hours from 10 AM to 10 PM). Thus, complex schedules specify the sessions on which the scheme should run. In the discussion below, "days" can be replaced with "scheduled sessions". |
|---|---|
| Number of levels | Select from 2 to 16 backup levels. See the example stated below for details. |
| Roll-back period | The guaranteed number of sessions that one can go back in the archive at any time. Calculated automatically, depending on the schedule parameters and the numbers of levels you select. See the example below for details. |
| Backup type | Specifies what backup types the backup levels will have<br><br>■ **Always full** - all levels of backups will be full. This is the default selection for cases when a tape drive is selected as a backup location.<br><br>■ **Full/Differential/Incremental** - backups of different levels will have different types:<br><br>- Last-level backups are full<br><br>- Backups of intermediate levels are differential<br><br>- First-level backups are incremental |

## Example

**Schedule** parameters are set as follows

- Recur: Every 1 day
- Frequency: Once at 6 PM

**Number of levels**: 4

**Backup type**: Full/Differential/Incremental

This is how the first 14 days (or 14 sessions) of this scheme's schedule look. Shaded numbers denote backup levels.

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 4 | 1 | 2 | 1 | 3 | 1 | 2 | 1 | 4 | 1 | 2 | 1 | 3 | 1 |

Backups of different levels have different types:

- *Last-level* (in this case, level 4) backups are full;
- Backups of *intermediate levels* (2, 3) are differential;
- *First-level* (1) backups are incremental.

A cleanup mechanism ensures that only the most recent backups of each level are kept. Here is how the archive looks on day 8, a day before creating a new full backup.

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|---|---|---|---|---|---|---|---|
| 4 | 1 | 2 | 1 | 3 | 1 | 2 | 1 |

The scheme allows for efficient data storage: more backups accumulate toward the current time. Having four backups, we could recover data as of today, yesterday, half a week, or a week ago.

**Roll-back period**

The number of days we can go back in the archive is different on different days. The minimum number of days we are guaranteed to have is called the roll-back period.

The following table shows full backup and roll-back periods for schemes of various levels.

| Number of levels | Full backup every | On different days, can go back | Roll-back period |
|---|---|---|---|
| 2 | 2 days | 1 to 2 days | 1 day |
| 3 | 4 days | 2 to 5 days | 2 days |
| 4 | 8 days | 4 to 11 days | 4 days |
| 5 | 16 days | 8 to 23 days | 8 days |
| 6 | 32 days | 16 to 47 days | 16 days |

Adding a level doubles the full backup and roll-back periods.

To see why the number of recovery days varies, let us return to the previous example.

Here are the backups we have on day 12 (numbers in gray denote deleted backups).

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 4 | 1 | 2 | 1 | 3 | 1 | 2 | 1 | 4 | 1 | 2 | 1 |

A new level 3 differential backup has not yet been created, so the backup of day five is still stored. Since it depends on the full backup of day one, that backup is available as well. This enables us to go as far back as 11 days, which is the best-case scenario.

The following day, however, a new third-level differential backup is created, and the old full backup is deleted.

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 4 | 1 | 2 | 1 | 3 | 1 | 2 | 1 | 4 | 1 | 2 | 1 | 3 |

This gives us only a four day recovery interval, which turns out to be the worst-case scenario.

On day 14, the interval is five days. It increases on subsequent days before decreasing again, and so on.

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 4 | 1 | 2 | 1 | 3 | 1 | 2 | 1 | 4 | 1 | 2 | 1 | 3 | 1 |

The roll-back period shows how many days we are guaranteed to have even in the worst case. For a four-level scheme, it is four days.

## 4.2.5.6    Manual start

With the **Manual start** scheme, you do not have to specify the backup schedule. You can run the backup plan from the **Plans and Tasks** view manually at any time afterwards.

Specify the appropriate settings as follows.

**Backup type**

Select the type of backup

- **Full -** selected by default for all backup locations (except for Acronis Online Backup Storage).
- **Incremental.** At the first time a full backup will be created. The next backups will be incremental. Selected as the one and only backup type for Acronis Online Backup Storage.
- **Differential.** At the first time a full backup will be created. The next backups will be differential.

## 4.2.5.7   Initial seeding

This backup scheme is only available when you have an Initial Seeding license and selected the Online Backup Storage as the backup destination.

Initial seeding enables you to transfer the first backup, which is full and usually the largest, to the online storage on a hard drive instead of over the Internet. Subsequent backups, which are all incremental and thus usually much smaller, can be transferred over the Internet after the full backup has arrived in the online storage.

If you back up a large amount of data, initial seeding ensures faster delivery of the backed-up data and lower traffic costs.

Please refer to the "Initial Seeding FAQ (p. 210)" section for more details.

## 4.2.6   Backup location selection

Specify where the archive will be stored.

### 1. Selecting the destination

In the **Path** field, enter the full path to the destination, or select the desired destination in the location tree as described in "Selecting backup destinations" (p. 45).

### 2. Using the archives table

To assist you with choosing the right destination, the table displays the names of the archives contained in each location you select. While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

### 3. Naming the new archive

Once you select the archive destination, the program generates a name for the new archive and displays it in the **Name** field. The name commonly looks like *Archive(N)*, where *N* is a sequence number. The generated name is unique within the selected location. If you are satisfied with the automatically generated name, click **OK**. Otherwise enter another unique name.

### Backing up to an existing archive

You can configure the backup plan to back up to an existing archive. To do so, select the archive in the archives table or type the archive name in the **Name** field. If the archive is protected with a password, the program will ask for it in the pop-up window.

By selecting the existing archive, you are meddling in the area of another backup plan that uses the archive. This is not an issue if the other plan is discontinued. However, you should generally follow

the rule: "one backup plan - one archive". Doing the opposite will not prevent the program from functioning but is not practical or efficient, except for some specific cases.

**Why two or more plans should not back up to the same archive**

1. Backing up different sources to the same archive makes it difficult to use archive. When it comes to recovery, every second counts, and you might be "lost" in the archive content.

   Backup plans that operate with the same archive should back up the same data items (say, both plans back up volume C.)

2. Applying multiple retention rules to an archive makes the archive content unpredictable. Since each of the rules will be applied to the entire archive, the backups belonging to one backup plan can be easily deleted along with the backups belonging to the other. You should not expect the classic behavior of the GFS and Tower of Hanoi backup schemes.

   Normally, each complex backup plan should back up to its own archive.

## 4.2.6.1 Selecting backup destinations

Acronis Backup & Recovery 11 lets you back up data to various physical storages.

| Destination | Details |
|---|---|
| Online backup storage | To back up data to Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then, expand the **Online backup storage** group and select the account. |
| | Prior to backing up to the online storage, you need to buy a subscription (p. 220) to the online backup service and activate (p. 221) the subscription on the machine(s) you want to back up. |
| | Online backup is not available under bootable media. |
| | *Note: Acronis Backup & Recovery Online might be unavailable in your region. To find more information, click here: http://www.acronis.com/my/backup-recovery-online/* |
| Personal | To back up data to a personal vault, expand the **Vaults** group and click the vault. Acronis Secure Zone is considered as a personal vault available to all users that can log on the system. |
| Machine | Local machine |
| Local folders | To back up data to the local folders of the machine, expand the **<Machine name>** group and select the required folder. |
| CD, DVD, etc. | To back up data to optical media such as CD or DVD, expand the **<Machine name>** group, then select the required drive. |
| Tape device | To back up data to a locally attached tape device, expand the **<Machine name>** group, then click the required device. |
| | Tape devices are available only if you have upgraded from Acronis Backup & Recovery 10. For information about using tapes, see the "Tape devices" section of the product Help. |
| Network folders | To back up data to the network folder, expand the **Network folders** group, select the required networked machine and, then click the shared folder. |
| | If the network share requires access credentials, the program will ask for them. |

| Destination | Details |
| --- | --- |
| FTP, SFTP | To back up data to FTP or SFTP, type the server name or address in the **Path** field as follows:<br><br>**ftp://ftp_server:port _number** or **sftp://sftp_server:port number**<br><br>If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.<br><br>After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.<br><br>You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials. |
|  | *Note: According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.* |

## 4.2.7    Archive validation

Set up the validation task to check if the backed up data is recoverable. If the backup could not pass the validation successfully, the validation task fails and the backup plan gets the Error status.

To set up validation, specify the following parameters

1. **When to validate** – select when to perform the validation. As the validation is a resource-intensive operation, it makes sense to **schedule** the validation to the managed machine's off-peak period. On the other hand, if the validation is a major part of your data protection strategy and you prefer to be immediately informed whether the backed up data is not corrupted and can be successfully recovered, think of starting the validation right after backup creation.
2. **What to validate** – select either to validate the entire archive or the latest backup in the archive. Validation of a file backup imitates recovery of all files from the backup to a dummy destination. Validation of a volume backup calculates a checksum for every data block saved in the backup. Validation of the archive will validate all the archive's backups and may take a long time and a lot of system resources.
3. **Validation schedule** (appears only if you have selected the on schedule in step 1) - set the schedule of validation. For more information see the Scheduling (p. 53) section.

## 4.2.8    Backup plan's credentials

Provide the credentials for the account under which the plan's tasks will run.

***To specify credentials***

1. Select one of the following:
   - **Run under the current user**

     The tasks will run under the credentials with which the user who starts the tasks is logged on. If any of the tasks has to run on schedule, you will be asked for the current user's password on completing the plan creation.
   - **Use the following credentials**

     The tasks will always run under the credentials you specify, whether started manually or executed on schedule.

     Specify:

- **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain)
- **Password**. The password for the account.

2.  Click **OK**.

To learn more about operations available depending on the user privileges, see the Users' privileges on a managed machine (p. 22) section.

# 4.2.9    Label (Preserving machine properties in a backup)

Any time data on a machine is backed up, information about the machine name, operating system, Windows service pack and security identifier (SID) is added to the backup, along with the user-defined text label. The label may include the department or machine owner's name or similar information that can be used as a tag or a key.

If you recover (p. 97) the machine to a VMware ESX(i) using Agent for ESX(i), or convert (p. 72) the backup to a ESX(i) virtual machine, these properties will be transferred to the virtual machine's configuration. You can view them in the virtual machine settings: **Edit settings** > **Options** > **Advanced** > **General** > **Configuration parameters**. You can select, sort and group the virtual machines with the help of these custom parameters. This can be useful in various scenarios.

**Example:**

Let's assume you migrate your office or datacenter to a virtual environment. By using third-party software that can access configuration parameters through VMware API, you can automatically apply security policies to each machine even before powering it on.

***To add a text label to a backup:***

1.  On the **Create backup plan** (p. 29) page, click **Show plan's credentials, comments, label**.
2.  In **Label**, enter the text label or select it from the drop-down menu.

## Parameters specification

| Parameter | Value | Description |
|---|---|---|
| **acronisTag.label** | <string> | A user-defined label. The label can be set by a user when creating a backup plan. |
| **acronisTag.hostname** | <string> | Host name (FQDN) |
| **acronisTag.os.type** | <string> | Operating system |
| **acronisTag.os.servicepack** | 0, 1, 2... | The version of the Service Pack installed in the system. For Windows OS only. |
| **acronisTag.os.sid** | <string> | Machine's SID. For example: S-1-5-21-874133492-782267321-3928949834. For Windows OS only. |

## Values of the "acronisTag.os.type" **parameter**

| | |
|---|---|
| Windows NT 4 | winNTGuest |
| Windows 2000 Professional | win2000ProGuest |

| | |
|---|---|
| Windows 2000 Server | win2000ServGuest |
| Windows 2000 Advanced Server | win2000ServGuest |
| Windows XP All Editions | winXPProGuest |
| Windows XP All Editions (64 bit) | winXPPro64Guest |
| Windows Server 2003, All Editions | winNetStandardGuest |
| Windows Server 2003, All Editions (64 bit) | winNetStandard64Guest |
| Windows 2008 | winLonghornGuest |
| Windows 2008 (64 bit) | winLonghorn64Guest |
| Windows Vista | winVistaGuest |
| Windows Vista (64 bit) | winVista64Guest |
| Windows 7 | windows7Guest |
| Windows 7 (64 bit) | windows7_64Guest |
| Windows Server 2008 R2 (64 bit) | windows7Server64Guest |
| Linux | otherLinuxGuest |
| Linux (64 bit) | otherLinux64Guest |
| Other Operating System | otherGuest |
| Other Operating System (64 bit) | otherGuest64 |

**Example**

acronisTag.label = "DEPT:BUCH; COMP:SUPERSERVER; OWNER:EJONSON"

acronisTag.hostname = "superserver.corp.local"

acronisTag.os.type = "windows7Server64Guest"

acronisTag.os.servicepack = "1"

acronisTag.os.sid = "S-1-5-21-874133492-782267321-3928949834"

## 4.2.10   Why is the program asking for the password?

A scheduled or postponed task has to run regardless of users being logged on. In case you have not explicitly specified the credentials, under which the task(s) will run, the program proposes using your account. Enter your password, specify another account or change the scheduled start to manual.

## 4.3   Simplified naming of backup files

When creating a backup plan (p. 29), you can choose between standard and simplified naming of backup files.

If you select the **Name backup files using the archive name...** check box:

- The file name of the first (full) backup in the archive will consist of the archive name; for example: **MyData.tib.** The file names of subsequent (incremental or differential) backups will have an index. For example: **MyData2.tib**, **MyData3.tib**, and so on.

  This simple naming scheme enables you to create a portable image of a machine on a detachable media or move the backups to a different location by using a script.

- Before creating a new full backup, the software will delete the entire archive and start a new one.

  This behavior is useful when you rotate USB hard drives and want each drive to keep a single full backup (p. 50) or all backups created during a week (p. 51). But you might end up with no backups if a full backup to your only drive fails.

  This behavior can be suppressed by adding the [Date] variable (p. 52) to the archive name.

If you *do not* select the **Name backup files using the archive name...** check box:

- Each backup will have a unique file name with the exact time stamp and the backup type. For example: **MyData_2010_03_26_17_01_38_960D.tib**. This standard file naming allows for a wider range of backup destinations and backup schemes.

**Restrictions**

When using simplified file naming, the following functionality is not available:

- Setting up full, incremental and differential backups within a single backup plan. You need to create separate backup plans for each type of backup
- Backup to a managed vault, tape, Acronis Secure Zone or Acronis Online Backup Storage
- Setting up retention rules
- Setting up regular conversion of backups to a virtual machine

*Tip. The FAT16, FAT32, and NTFS file systems do not allow the following characters in the file name: backslash (\), slash (/), colon (:), asterisk (\*), question mark (?), quotation mark ("), less than sign (<), greater than sign (>), and pipe (|).*

## 4.3.1   Usage examples

This section provides examples of how you can use simplified file naming.

### 4.3.1.1   Example 1. Daily backup replacing the old one

Consider the following scenario:

- You want to perform a daily full backup of your machine.
- You want to store the backup locally in the file **MyMachine.tib**.
- You want each new backup to replace the old one.

In this scenario, create a backup plan with a daily schedule. When creating the backup plan, specify **MyMachine** as the archive name, select the **Name backup files using the archive name...** check box, and select **Full** as the backup type.

**Result.** The archive consists of a single file: MyMachine.tib. This file is deleted before creating a new backup.

### 4.3.1.2   Example 2. Daily full backups with a date stamp

Consider the following scenario:

- You want to perform a daily full backup of your machine.
- You want to move older backups to a remote location by using a script.

In this scenario, create a backup plan with a daily schedule. When creating the backup plan, specify **MyMachine-[DATE]** as the archive name, select the **Name backup files using the archive name...** check box, and select **Full** as the backup type.

**Result:**

- The backups of January 1, 2011, January 2, 2011, and so on, are stored respectively as MyMachine-1.1.2011.tib, MyMachine-1.2.2011.tib, and so on.
- Your script can move older backups based on the date stamp.

See also "The [Date] variable" (p. 52).

## 4.3.1.3  Example 3. Hourly backups within a day

Consider the following scenario:

- You want to perform hourly backups of your server's critical files every day.
- You want the first backup of each day to be full and to run at midnight; and the subsequent backups of the day to be differential and to run at 01:00, 02:00, and so on.
- You want to keep older backups in the archive.

In this scenario, create a backup plan with a daily schedule. When creating the backup plan, specify **ServerFiles([Date])** as the archive name, select the **Name backup files using the archive name...** check box, specify **Differential** as the backup type, and schedule the backups to run every hour from midnight.

**Result:**

- The 24 backups of January 1, 2011, will be stored as ServerFiles(1.1.2011).tib, ServerFiles(1.1.2011)2.tib, and so on up to ServerFiles(1.1.2011)24.tib.
- The following day, the backups will start with the full backup ServerFiles(1.2.2011).tib.

See also "The [Date] variable" (p. 52).

## 4.3.1.4  Example 4. Daily full backups with daily drive swaps

Consider the following scenario:

- You want to perform daily full backups of your machine to the file **MyMachine.tib** on an external hard disk drive.
- You have two such drives. Either of them has the drive letter **D** when attached to the machine.
- You want to swap the drives before each backup, so that one drive contains today's backup and the other drive yesterday's backup.
- You want each new backup to replace the backup on the currently attached drive.

In this scenario, create a backup plan with a daily schedule. When creating the backup plan, specify **MyMachine** as the archive name and **D:\** as the archive location, select the **Name backup files using the archive name...** check box, and select **Full** as the backup type.

**Result.** Each hard disk drive will contain one full backup. While one drive is attached to the machine, you can keep the other drive off-site for extra data protection.

## 4.3.1.5    Example 5. Daily backups with weekly drive swaps

Consider the following scenario:

- You want to perform daily backups of your machine: a full backup each Monday and incremental backups on Tuesday through Sunday.
- You want to back up to the archive **MyMachine** on an external hard disk drive.
- You have two such drives. Either of them has drive letter **D** in the operating system when attached to the machine.
- You want to swap the drives each Monday, so that one drive contains backups of the current week (Monday through Sunday), and the other drive those of the previous week.

In this scenario, you need to create two backup plans as follows:

a) When creating the first backup plan, specify **MyMachine** as the archive name and **D:\** as the archive location, select the **Name backup files using the archive name...** check box, select **Full** as the backup type, and schedule the backups to run every week on Monday.

b) When creating the second backup plan, specify the same settings as in the first backup plan, but select **Incremental** as the backup type and schedule the backups to run every week on Tuesday through Sunday.

**Result:**

- Before creating a Monday backup (by the first backup plan), all backups will be deleted from the currently attached drive.
- While one drive is attached to the machine, you can keep the other drive off-site for extra data protection.

## 4.3.1.6    Example 6. Backups within working hours

Consider the following scenario:

- You want to back up your server's critical files every day.
- You want the first backup of each day to be full and to run at 01:00 AM.
- You want the backups during working hours to be differential and to run every hour from 8:00 AM through 5:00 PM.
- You want to include a creation date in the name of each backup file.

In this scenario, you need to create two backup plans as follows:

a) When creating the first backup plan, specify **ServerFiles([DATE])** as the archive name, select the **Name backup files using the archive name...** check box, select **Full** as the backup type, and schedule the backups to run every day at 01:00:00 AM.

b) When creating the second backup plan, specify the same settings as in the first backup plan, but select **Differential** as the backup type and schedule the backups as follows:

- **Run the task**: **Daily**
- **Every**: **1 Hour(s)**
- **From**: **08:00:00 AM**
- **Until**: **05:01:00 PM**

**Result:**

- The full backup of January 31, 2011, will be stored as ServerFiles(1.31.2011).tib.

- The 10 differential backups of January 31, 2011, will be stored as ServerFiles(1.31.2011)2.tib, ServerFiles(1.31.2011)3.tib, and so on up to ServerFiles(1.31.2011)11.tib.
- The following day, February 1, the backups will start with the full backup ServerFiles(2.1.2011).tib. The differential backups will start with ServerFiles(2.1.2011)2.tib.

See also "The [Date] variable" (p. 52).

## 4.3.2    The [DATE] variable

If you specify the **[DATE]** variable in the archive name, the file name of each backup will include that backup's creation date.

When using this variable, the first backup of a new day will be a full backup. Before creating the next full backup, the software deletes all backups taken earlier that day. Backups taken before that day are kept. This means you can store multiple full backups with or without incremental ones, but no more than one full backup per day. You can sort the backups by date, copy, move, delete the backups manually or by using a script.

The date format is *m*.*d*.*yyyy*. For example, it is 1.31.2011 for January 31, 2011. (Note absence of leading zeros.)

You can place this variable anywhere in the archive name. You can use both lowercase and uppercase letters in this variable.

**Examples**

**Example 1.** Suppose that you perform incremental backups twice a day (at midnight and noon) for two days, starting on January 31, 2011. If the archive name is **MyArchive-[DATE]-**, here is the list of backup files after day two:

**MyArchive-1.31.2011-.tib** (full, created on January 31 at midnight)
**MyArchive-1.31.2011-2.tib** (incremental, created on January 31 at noon)
**MyArchive-2.1.2011-.tib** (full, created on February 1 at midnight)
**MyArchive-2.1.2011-2.tib** (incremental, created on February 1 at noon)

**Example 2.** Suppose that you perform full backups, with the same schedule and archive name as in the previous example. Then, the list of backup files after day two is the following:

**MyArchive-1.31.2011-.tib** (full, created on January 31 at noon)
**MyArchive-2.1.2011-.tib** (full, created on February 1 at noon)

This is because the full backups created at midnight were replaced by new full backups of the same day.

## 4.3.3    Backup splitting and simplified file naming

When a backup is split according to backup splitting (p. 81) settings, the same indexing is used to also name parts of the backup. The file name for the next backup will have the next available index.

For example, suppose that the first backup of the archive **MyData** has been split in two parts. Then, the file names for this backup are **MyData1.tib** and **MyData2.tib**. The second backup (supposing that it is not split) will be named **MyData3.tib**.

# 4.4 Scheduling

Acronis scheduler helps the administrator adapt backup plans to the company's daily routine and each employee's work style. The plans' tasks will be launched systematically keeping the critical data safely protected.

The scheduling is available when creating a backup plan (p. 29) with any of the following backup schemes: Simple, Custom or Tower of Hanoi. The schedule also can be set for validation tasks (p. 143).

The scheduler uses local time of the machine the backup plan exists on. Before creating a schedule, be sure the machine's date and time settings are correct.

## Schedule

To define when a task has to be executed, you need to specify an event or multiple events. The task will be launched as soon as any of the events occurs. The table below lists the events available under Windows operating systems.

| Event |
|---|
| Time: Daily, Weekly, Monthly |
| Time passed since the last successful backup has completed<br><br>(specify the length of time) |
| User logon<br><br>(any user, current user, specify the user's account) |
| User logoff*<br><br>(any user, current user, specify the user's account)<br><br>*Shutting down is not the same as logging off. The task will not run at a system shutdown. |
| System startup |
| Free space change<br><br>(specify the amount of free space change on any volume selected for backup or containing data selected for backup) |
| An event in Windows event log<br><br>(specify the parameters of the event) |
| Upon an Acronis Drive Monitor alert |

## Condition

For backup operations only, you can specify a condition or multiple conditions in addition to the events. Once any of the events occurs, the scheduler checks the condition and runs the task if the condition is met. With multiple conditions, all of them must be met simultaneously to enable task execution. The table below lists the conditions available under Windows operating systems.

| Condition: run the task only if |
|---|
| User is idle (a screen saver is running or the machine is locked) |
| Location's host is available |

| The task run time is within the specified time interval |
| The specified period of time has passed since the last successful backup completed |

| The task run time is within the specified time interval |
| --- |
| All users are logged off |
| The specified period of time has passed since the last successful backup completed |

The scheduler behavior, in case the event occurs but the condition (or any of multiple conditions) is not met is defined by the Task start conditions (p. 94) backup option.

**What-ifs**

- **What if an event occurs (and a condition, if any, is met) while the previous task run has not completed?**

  The event will be ignored.

- **What if an event occurs while the scheduler is waiting for the condition required by the previous event?**

  The event will be ignored.

- **What if the condition is not met for a very long time?**

  If delaying a backup is getting risky, you can force the condition (tell the users to log off) or run the task manually. To automatically handle this situation, you can set the time interval after which the task will run regardless of the condition.

## 4.4.1 Daily schedule

Daily schedule is effective in Windows and Linux operating systems.

***To specify a daily schedule***

In the **Schedule** area, select the appropriate parameter as follows:

| Every: <...> day(s) | Set up the certain number of days you want the task to be run. For example, if you set Every 2 day(s), the task will be started on every other day. |
| --- | --- |

In the **During the day execute the task...** area, select one of the following:

| Once at: <...> | Set up the time at which the task will be run once. |
| --- | --- |
| Every: <...> <br> From: <...> Until: <...> | Set up how many times the task will be restarted during the specified time interval. For example, setting the task frequency to Every 1 hour From 10:00:00 AM until 10:00:00 PM allows the task to run 12 times: from 10 AM to 10 PM during one day. |

In the **Effective...** area, set the following settings:

| From: <...> | Set up a date when this schedule will be enabled (an effective date). If this check box is cleared, the task will be started on the nearest day and time you have specified above. |
| --- | --- |
| To: <...> | Set up a date when this schedule will be disabled. If this check box is cleared, the task will be run for an indefinite number of days. |

Advanced scheduling settings are available only for machines registered on Acronis Backup & Recovery 11 Management Server. To specify these settings, click **Change** in the **Advanced settings** area.

All the settings you made are displayed in the **Result** field at the bottom of the window.

## Examples

**"Simple" daily schedule**

Run the task every day at 6PM.

The schedule's parameters are thus set up as follows.

1. Every: **1** day(s).
2. Once at: **06:00:00 PM**.
3. Effective:

   From: **not set**. The task will be started on the current day, if it has been created before 6PM. If you have created the task after 6 PM, the task will be started for the first time on the next day at 6 PM.

   To: **not set**. The task will be performed for an indefinite number of days.

**"Three-hour time interval lasting for three months" schedule**

Run the task every three hours. The task starts on a certain date (say, September 15, 2009), and ends after three months.

The schedule's parameters are thus set up as follows.

1. Every: **1** day(s).
2. Every: **3** hours

   From: **12:00:00 AM** (midnight) Until: **09:00:00 PM** - thus, the task will be performed 8 times a day with a 3 hour time interval. After the last daily recurrence at 9 PM, the next day comes and the task starts over again from midnight.
3. Effective:

   From: **09/15/2009**. If September 15, 2009 is the current date of the task's creation and, say, 01:15 PM is the task's creation time, the task will be started when the nearest time interval comes: at 03:00 PM in our example.

   To: **12/15/2009**. On this date the task will be performed for the last time, but the task itself is still available in the **Tasks** view.

**Several daily schedules for one task**

There are some cases when you might need the task to be run several times a day, or even several times a day with different time intervals. For such cases, consider adding several schedules to a single task.

For example, suppose that the task has to be run every 3rd day, starting from 09/20/2009, five times a day:

- first at 8 AM
- second at 12 PM (noon)
- third at 3 PM
- fourth at 5 PM
- fifth at 7 PM

The obvious way is to add five simple schedules. If you spend one minute for examination, you can think out a more optimal way. As you can see, the time interval between the first and the second

task's recurrences is 4 hours, and between the third, fourth and fifth is 2 hours. In this case, the optimal way is to add two schedules to the task.

**First daily schedule**

1. Every: **3** day(s).
2. Every: **4** hours.

   From: **08:00:00 AM** Until: **12:00:00 PM**.
3. Effective:

   From: **09/20/2009**.

   To: **not set**.

**Second daily schedule**

1. Every: **3** day(s).
2. Every: **2** hour(s).

   From: **03:00:00 PM** Until: **07:00:00 PM**.
3. Effective:

   From: **09/20/2009**.

   To: **not set**.

## 4.4.2    Weekly schedule

Weekly schedule is effective in Windows and Linux operating systems.

### To specify a weekly schedule

In the **Schedule** area, select the appropriate parameter as follows:

| Every: <...> week(s) on: <...> | Specify a certain number of weeks and the days of the week you want the task to be run. For example, with the Every **2** week(s) on **Mon** setting, the task will be performed on Monday of every other week. |
|---|---|

In the **During the day execute the task...** area, select one of the following:

| Once at: <...> | Set up the time at which the task will be run once. |
|---|---|
| Every: <...><br><br>From: <...> Until: <...> | Set up how many times the task will be run during the specified time interval. For example, setting the task frequency to Every **1** hour From **10:00:00 AM** until **10:00:00 PM** allows the task to be run 12 times from 10 AM to 10 PM during one day. |

In the **Effective...** area, set the following settings:

| From: <...> | Set up a date when this schedule will be enabled (an effective date). If this check box is cleared, the task will be started on the nearest day and time you have specified above. |
|---|---|
| To: <...> | Set up a date when this schedule will be disabled. If this check box is cleared, the task will be run for an indefinite number of weeks. |

Advanced scheduling settings are available only for machines registered on Acronis Backup & Recovery 11 Management Server. To specify these settings, click **Change** in the **Advanced settings** area.

All the settings you made are displayed in the **Result** field at the bottom of the window.

## Examples

**"One day in the week" schedule**

Run the task every Friday at 10PM, starting from a certain date (say 05/14/2009) and ending after six months.

The schedule's parameters are thus set up as follows.

1. Every: **1** week(s) on: **Fri**.
2. Once at: **10:00:00 PM**.
3. Effective:

   From: **05/13/2009**. The task will be started on the nearest Friday at 10 PM.

   To: **11/13/2009**. The task will be performed for the last time on this date, but the task itself will still be available in the Tasks view after this date. (If this date were not a Friday, the task would be last performed on the last Friday preceding this date.)

This schedule is widely used when creating a custom backup scheme. The "One day in the week"-like schedule is added to the full backups, while the incremental backups are scheduled to be performed on workdays. For more details, see the Full and incremental backups plus cleanup example in the Custom backup scheme (p. 39) section.

**"Workdays" schedule**

Run the task every week on workdays: from Monday through Friday. During a workday, the task starts only once at 9 PM.

The schedule's parameters are thus set up as follows.

1. Every: **1** week(s) on: **<Workdays>** - selecting the <Workdays> check box automatically selects the corresponding check boxes (**Mon**, **Tue**, **Wed**, **Thu**, and **Fri**), and leaves the remaining ones unchanged.
2. Once at: **09:00:00 PM**.
3. Effective:

   From: **empty**. If you have created the task, say on Monday at 11:30 AM, the task will be started on the same day at 9 PM. If the task was created, say on Friday after 9 PM, then it will be started for the first time on the nearest workday (Monday in our example) at 9 PM.

   End date: **empty**. The task will be restarted for an indefinite number of weeks.

This schedule is widely used when creating a custom backup scheme. The "Workdays"-like schedule is added to the incremental backups, while the full backup is scheduled to be performed one day in the week. For more details, see the Full and incremental backups plus cleanup example in the Custom backup scheme (p. 39) section.

**Several weekly schedules for one task**

In the case when the task needs to be run on different days of the weeks with different time intervals, consider adding a dedicated schedule to every desired day of the week, or to several days.

For example, you need the task to be run with the following schedule:

▪ Monday: twice at 12 PM (noon) and 9 PM

▪ Tuesday: every 3 hours from 9 AM till 9 PM

▪ Wednesday: every 3 hours from 9 AM till 9 PM

- Thursday: every 3 hours from 9 AM till 9 PM
- Friday: twice at 12 PM and 9 PM (i.e. same as on Monday)
- Saturday: once at 9 PM
- Sunday: once at 9 PM

Combining the identical times, the following three schedules can be added to the task:

**First schedule**

1. Every: **1** week(s) on: **Mon**, **Fri**.
2. Every: **9** hours

   From: **12:00:00 PM** Until: **09:00:00 PM**.
3. Effective:

   From: **not set**.

   To: **not set**.

**Second schedule**

1. Every **1** week(s) on: **Tue**, **Wed**, **Thu**.
2. Every **3** hours

   From **09:00:00 AM** until **09:00:00 PM**.
3. Effective:

   From: **not set**.

   To: **not set**.

**Third schedule**

1. Every: **1** week(s) on: **Sat, Sun**.
2. Once at: **09:00:00 PM**.
3. Effective:

   From: **not set**.

   To: **not set**.

## 4.4.3    Monthly schedule

Monthly schedule is effective in Windows and Linux operating systems.

### To specify a monthly schedule

In the **Schedule** area, select the appropriate parameter as follows:

| | |
|---|---|
| **Months: <...>** | Select a certain month(s) you want to run the task in. |
| **Days: <...>** | Select specific days of the month to run the task on. You can also select the last day of the month, irrespective of its actual date. |
| **On: <...> <...>** | Select specific days of the weeks to run the task on. |

In the **During the day execute the task...** area, select one of the following:

| | |
|---|---|
| **Once at: <...>** | Set up the time at which the task will be run once. |

| Every: <...><br><br>From: <...> Until: <...> | Set up how many times the task will be run during the specified time interval. For example, setting the task frequency to Every **1** hour From **10:00:00 AM** until **10:00:00 PM** allows the task to be run 12 times from 10 AM to 10 PM during one day. |
|---|---|

In the **Effective...** area, set the following settings:

| From: <...> | Set up a date when this schedule will be enabled (an effective date). If this check box is cleared, the task will be started on the nearest day and time you have specified above. |
|---|---|
| To: <...> | Set up a date when this schedule will be disabled. If this check box is cleared, the task will be run for an indefinite number of months. |

Advanced scheduling settings are available only for machines registered on Acronis Backup & Recovery 11 Management Server. To specify these settings, click **Change** in the **Advanced settings** area.

All the settings you made are displayed in the **Result** field at the bottom of the window.

## Examples

**"Last day of every month" schedule**

Run the task once at 10 PM on the last day of every month.

The schedule's parameters are set up as follows.

1. Months: **<All months>**.
2. Days: **Last**. The task will run on the last day of every month despite its actual date.
3. Once at: **10:00:00 PM**.
4. Effective:

   From: **empty**.

   To: **empty**.

This schedule is widely used when creating a custom backup scheme. The "Last day of every month" schedule is added to the full backups, while the differential backups are scheduled to be performed once a week and incremental on workdays. For more details, see the Monthly full, weekly differential, and daily incremental backups plus cleanup example in the Custom backup scheme (p. 39) section.

**"Season" schedule**

Run the task on all workdays during the northern autumn seasons of 2009 and 2010. During a workday, the task is performed every 6 hours from 12 AM (midnight) till 6 PM.

The schedule's parameters are set up as follows.

1. Months: **September, October, November**.
2. On: **<all> <workdays>**.
3. Every: **6** hours.

   From: **12:00:00 AM** Until: **06:00:00 PM**.
4. Effective:

   From: **08/30/2009**. Actually the task will be started on the first workday of September. By setting up this date we just define that the task must be started in 2009.

To: **12/01/2010**. Actually the task will end on the last workday of November. By setting up this date we just define that the task must be discontinued in 2010, after autumn ends in the northern hemisphere.

**Several monthly schedules for one task**

In the case when the task needs to be run on different days or weeks with different time intervals depending on the month, consider adding a dedicated schedule to every desired month or several months.

Suppose that the task goes into effect on 11/01/2009.

- During northern winter, the task runs once at 10PM on every workday.
- During northern spring and autumn, the task runs every 12 hours on all workdays.
- During northern summer, the task runs every first and fifteenth of every month at 10 PM.

Thus, the following three schedules are added to the task.

**First schedule**

1. Months: **December**, **January**, **February**.
2. On: **<All> <All workdays>**
3. Once at: **10:00:00 PM**.
4. Effective:

   From: **11/01/2009**.

   To: **not set**.

**Second schedule**

1. Months: **March**, **April**, **May**, **September**, **October**, **November**.
2. On: **<All> <All workdays>**.
3. Every: **12 hours**

   From: **12:00:00 AM** Until: **12:00:00 PM**.
4. Effective:

   From: **11/01/2009**.

   To: **not set**.

**Third schedule**

1. Months: **June**, **July**, **August**.
2. Days: **1**, **15**.
3. Once at: **10:00:00 PM**.
4. Effective:

   From: **11/01/2009**.

   To: **not set**.

## 4.4.4    At Windows Event Log event

This type of schedule is effective only in Windows operating systems.

You can schedule a backup task to start when a certain Windows event has been recorded in one of the event logs such as the Application, Security, or System log.

For example, you may want to set up a backup plan that will automatically perform an emergency full backup of your data as soon as Windows discovers that your hard disk drive is about to fail.

### Parameters

**Log name**

Specifies the name of the log. Select the name of a standard log (**Application**, **Security**, or **System**) from the list, or type a log name—for example: **Microsoft Office Sessions**

**Event source**

Specifies the event source, which typically indicates the program or the system component that caused the event—for example: **disk**

**Event type**

Specifies the event type: **Error**, **Warning**, **Information**, **Audit success**, or **Audit failure**.

**Event ID**

Specifies the event number, which typically identifies the particular kind of events among events from the same source.

For example, an **Error** event with Event source **disk** and Event ID **7** occurs when Windows discovers a bad block on a disk, whereas an **Error** event with Event source **disk** and Event ID **15** occurs when a disk is not ready for access yet.

### Examples

**"Bad block" emergency backup**

One or more bad blocks that have suddenly appeared on a hard disk usually indicate that the hard disk drive will soon fail. Suppose that you want to create a backup plan that will back up hard disk data as soon as such a situation occurs.

When Windows detects a bad block on a hard disk, it records an event with the event source **disk** and the event number **7** into the **System** log; the type of this event is **Error**.

When creating the plan, type or select the following in the **Schedule** area:

- **Log name**: **System**
- **Event source**: **disk**
- **Event type**: Error
- **Event ID**: **7**

*Important: To ensure that such a task will complete despite the presence of bad blocks, you must make the task ignore bad blocks. To do this, in **Backup options**, go to **Error handling**, and then select the **Ignore bad sectors** check box.*

**Pre-update backup in Vista**

Suppose that you want to create a backup plan that will automatically perform a backup of the system—for example, by backing up the volume where Windows is installed—every time that Windows is about to install updates.

Having downloaded one or more updates and scheduled their installation, the Microsoft Windows Vista operating system records an event with the event source

**Microsoft-Windows-WindowsUpdateClient** and event number **18** into the **System** log; the type of this event is **Information**.

When creating the plan, type or select the following in the **Schedule** area:

- **Log name**: **System**
- **Event source**: **Microsoft-Windows-WindowsUpdateClient**
- **Event type**: **Information**
- **Event ID**: **18**

*Tip: To set up a similar backup plan for machines running Microsoft Windows XP, replace the text in **Event source** with **Windows Update Agent** and leave the remaining fields the same.*

### How to view events in Event viewer

*To open a log in Event Viewer*

1. On the Desktop or in the **Start** menu, right-click **My Computer**, and then click **Manage**.
2. In the **Computer Management** console, expand **System Tools**, and then expand **Event Viewer**.
3. In **Event Viewer**, click the name of a log that you want to view—for example, **Application**.

   *Note: To be able to open the security log (**Security**), you must be a member of the Administrators group.*

*To view properties of an event, including the event source and event number*

1. In **Event Viewer**, click the name of a log that you want to view—for example, **Application**.

   *Note: To be able to open the security log (**Security**), you must be a member of the Administrators group.*

2. In the list of events in the right pane, double-click the name of an event whose properties you want to view.
3. In the **Event Properties** dialog box, view the event's properties such as the event source, shown in the **Source** field; and the event number, shown in the **Event ID** field.

When you are finished, click **OK** to close the **Event Properties** dialog box.

## 4.4.5    Upon an Acronis Drive Monitor alert

This schedule is effective in Windows operating systems when Acronis® Drive Monitor™ is installed.

Acronis Drive Monitor reports on hard disk health by using the hard disk's internal monitoring system (S.M.A.R.T.). Based on alerts from Acronis Drive Monitor, you can set up emergency backups of your data in addition to regular backups. The emergency backup will start when a hard disk with your data is about to fail.

The backup starts as soon as disk health reaches a warning level or a critical level. You can see the disk health indicator (as a percentage) for each disk by opening Acronis Drive Monitor.

Alerts about disk temperature do not start the backup.

*Tip: If your backup plan uses the custom backup scheme (p. 39), you can set up this emergency backup simply by adding an extra schedule to the same backup plan. When using a different backup scheme, you will need to create a separate backup plan.*

## 4.4.6  Conditions

Conditions add more flexibility to the scheduler, enabling to execute backup tasks with respect to certain conditions. Once a specified event occurs (see the "Scheduling (p. 53)" section for the list of available events), the scheduler checks the specified condition and executes the task if the condition is met.

The scheduler behavior in case the event occurs but the condition (or any of multiple conditions) is not met, is defined by the **Task start conditions** (p. 94) backup option. There, you can specify how important the conditions are for the backup strategy:

- conditions are obligatory - put the backup task run on hold until all the conditions are met.
- conditions are preferable, but a backup task run has higher priority - put the task on hold for the specified time interval. If the time interval lapses and the conditions are still not met, run the task anyway. With this setting, the program will automatically handle the situation when the conditions are not met for too long and further delaying the backup is undesirable.
- backup task start time matters - skip the backup task if the conditions are not met at the time when the task should be started. Skipping the task run makes sense when you need to back up data strictly at the specified time, especially if the events are relatively often.

*Conditions are available only when the custom backup scheme (p. 39) is used. You can set conditions for full, incremental and differential backup separately.*

### Adding multiple conditions

Multiple conditions must be met simultaneously to enable task execution.

**Example:**

It is required to run the backup task after free space on the managed machine is changed by at least 1 GB, but only if all users are logged off and more than 12 hours have passed since the last backup.

Set the schedule, conditions and the **Task start conditions** backup option as follows:

- Schedule: **When free space changed**; Value: Run task if free space has changed by at least: **1 GB**.
- Condition: **User logged off**; Value: Run the task on schedule only if all users are logged off.
- Condition: **Time since last backup**; Value: Time since the last backup: **12 hour(s)**.
- Task start conditions: **Wait until the conditions are met**.

If the free space changes by more than 1 GB, the scheduler will wait until both conditions are met at the same time and then run the backup task.

## 4.4.6.1  User is idle

*Applies to: Windows*

"User is idle" means that a screen saver is running on the managed machine or the machine is locked.

### Example:

Run the backup task on the managed machine every day at 9PM, preferably when the user is idle. If the user is still active by 11PM, run the task anyway.

- Event: **Daily**, every **1** day(s); Once at: **09:00:00 PM**.
- Condition: **User is idle**.

- Task start conditions: **Wait until the conditions are met**, Run the task anyway after **2** hour(s).

As a result,

(1) If the user becomes idle before 9PM, the backup task will start at 9PM.

(2) If the user becomes idle between 9PM and 11PM, the backup task will start immediately after the user becomes idle.

(3) If the user is still active at 11PM, the backup task starts anyway.

## 4.4.6.2    Location's host is available

*Applies to: Windows, Linux*

"Location's host is available" means that the machine hosting the destination for storing archives on a networked drive is available.

### Example:

Backing up data to the networked location is performed on workdays at 9:00 PM. If the location's host is not available at that moment (for instance, due to maintenance work), skip the backup and wait for the next workday to start the task. It is assumed that the backup task should not be started at all rather than failed.

- Event: **Weekly**, Every **1** week(s) on **<workdays>**; Once at **09:00:00 PM**.
- Condition: **Location's host is available**
- Task start conditions: **Skip the task execution**.

As a result,

(1) If 9:00 PM comes and the location's host is available, the backup task starts right on time.

(2) If 9:00 PM comes but the host is unavailable at the moment, the backup task will start on the next workday if the location's host is available.

(3) If the location's host will never be available on workdays at 9:00 PM, the task never starts.

## 4.4.6.3    Fits time interval

*Applies to: Windows, Linux*

Restricts a backup task's start time to a specified interval.

### Example

A company uses different locations on the same network-attached storage for backing up users data and servers. The workday starts at 8AM and ends at 5 PM. Users' data should be backed up as soon as the users log off, but not earlier than 4:30 PM and not later than 10 PM. Every day at 11 PM the company's servers are backed up. So, all the users' data should be preferably backed up before this time, in order to free network bandwidth. By specifying the upper limit as 10 PM, it is supposed that the backing up of users' data does not take more than one hour. If a user is still logged on within the specified time interval, or logs off at any other time – do not back up the users' data, i.e. skip task execution.

- Event: **When logging off**, The following user: **Any user**.

- Condition: **Fits the time interval**, from **04:30:00 PM** until **10:00:00 PM**.
- Task start conditions: **Skip the task execution**.

As a result,

(1) if the user logs off between 04:30:00 PM and 10:00:00 PM, the backup task will start immediately following the logging off.

(2) if the user logs off at any other time, the task will be skipped.

**What if...**

*What if a task is scheduled to be executed at a certain time and this time is outside the specified time interval?*

For example:

- Event: **Daily**, Every **1** day(s); Once at **03:00:00 PM**.
- Condition: **Fits time interval**, from **06:00:00 PM** until **11:59:59 PM**.

In this case, whether and when the task will run depends on the task start conditions:

- If the task start conditions are **Skip the task execution**, the task will never run.
- If the task start conditions are **Wait until the conditions are met** and the **Run the task anyway after** check box is *cleared*, the task (scheduled to run at 3:00 PM) will start at 6:00 PM—the time when the condition is met.
- If the task start conditions are **Wait until the conditions are met** and the **Run the task anyway after** check box is *selected* with, say, the **1 Hour** waiting time, the task (scheduled to run at 3:00 PM) will start at 4:00 PM—the time when the waiting period ends.

## 4.4.6.4    User logged off

*Applies to: Windows*

Enables to put a backup task run on hold until all users log off from Windows on the managed machine.

**Example**

Run the backup task at 8 PM on the first and third Friday of every month, preferably when all users are logged off. If one of the users is still logged on at 11 PM, run the task anyway.

- Event: **Monthly**, Months: **<All>**; On: **<First>**, **<Third> <Friday>**; Once at **08:00:00 PM**.
- Condition: **User logged off**.
- Task start conditions: **Wait until the conditions are met**, Run the task anyway after **3** hour(s).

As a result,

(1) If all users are logged off at 8PM, the backup task will start at 8PM.

(2) If the last user logs off between 8PM and 11PM, the backup task will start immediately after the user has logged off.

(3) If any of the users is still logged on at 11PM, the backup task starts anyway.

## 4.4.6.5 Time since last backup

*Applies to: Windows, Linux*

Enables to put a backup task run on hold until the specified time interval since the last successful backup completion passes.

**Example:**

Run the backup task at system startup, but only if more than 12 hours have passed since the last successful backup.

- Event: **At startup**, Start the task on machine startup.
- Condition: **Time since last backup**, Time since the last backup: **12** hour(s).
- Task start conditions: **Wait until the conditions are met**.

As a result,

(1) if the machine is restarted before 12 hours pass since the completion of the latest successful backup, the scheduler will wait until 12 hours pass, and then will start the task.

(2) if the machine is restarted after 12 hours have passed since the completion of the latest successful backup, the backup task will start immediately.

(3) if the machine is never restarted, the task will never start. You can start the backup manually, if need be, in the **Backup plans and tasks** view.

# 4.5 Replication and retention of backups

When creating a backup plan (p. 29), you specify the primary location for the backups. In addition, you can do the following:

- Replicate (copy) each backup to a second location immediately after creation.
- Retain the backups according to the retention rules you specify, and then either move them to a second location or delete them.

Similarly, you can copy or move backups from a second location to a third location and so on. Up to five consecutive locations are supported (including the primary one).

*Note: The replication feature replaces and enhances the **Dual destination** option, which was available in Acronis Backup & Recovery 10.*

**Example.** You back up your machine to a local folder. The backup is immediately copied to a network folder. In the original local folder, the backup is stored for just one month.

The following picture illustrates this example.

**Usage scenarios**

- **Reliable disaster recovery** (p. 71)

  Store your backups both on-site (for immediate recovery) and off-site (to secure the backups from local storage failure or a natural disaster).

- **Keeping only the latest recovery points** (p. 71)

  Delete older backups from a fast storage according to retention rules, in order to not overuse expensive storage space.

- **Reduced costs of storing the backed up data**

  Store your backups on a fast storage for as long as a need to access them is likely. Then, move them to a lower-cost storage to keep them there for a longer term. This enables you to meet legal requirements on data retention.

**Replication and retention in backup schemes**

The following table shows availability of replication and retention rules in various backup schemes.

| Backup scheme | Can copy backups | Can move backups | Can delete backups |
|---|---|---|---|
| **Run now** (p. 35) | Yes | No | No |
| **Manual start** (p. 43) | Yes | No | No |
| **Simple** (p. 35) | Yes | Yes | Yes |
| **Grandfather-Father-Son (GFS)** (p. 35) | Yes | No | Yes |
| **Tower of Hanoi** (p. 41) | Yes | No | Yes |
| **Custom** (p. 39) | Yes | Yes | Yes |
| **Initial seeding** (p. 44) | No | No | No |

**Notes:**

- Setting up both copying and moving backups from the same location is not possible.
- With simplified naming of backup files (p. 48), neither replication nor use of retention rules is available.

## 4.5.1 Supported locations

You can copy or move a backup *from* any of these locations:

- A local folder on a fixed or removable drive
- A network folder
- An FTP or SFTP server
- Acronis Secure Zone

You can copy or move a backup *to* any of these locations:

- A local folder on a fixed or removable drive
- A network folder
- An FTP or SFTP server

Backups that were copied or moved to the next location do not depend on the backups remaining in the original location and vice versa. You can recover data from any backup without access to other locations.

**Restrictions**

- Copying or moving backups *to and from* optical discs (CD, DVD, Blu-ray discs) is not supported.
- You cannot specify the same location more than once. For example, you cannot move a backup from one folder to another and then back to the original folder.

## 4.5.2    Setting up replication of backups

Setting up replication of backups is available when creating a backup plan (p. 29).

- To set up replication from the primary location, select the **Replicate just created backup to another location** check box.
- To set up replication from the second or a further location, select the **Replicate backups to another location as soon as they appear in this location** check box.

Next, select the location where to replicate the backups. A backup is replicated to the next location as soon as it appears in the previous location.

If allowed by the backup scheme, you can also specify when to automatically delete the backups from each of the locations.

## 4.5.3    Setting up retention of backups

You can set retention rules for backups when creating a backup plan (p. 29). The available retention rules depend on the chosen backup scheme.

Applying retention rules can be restricted by the **Replication/cleanup inactivity time** (p. 70) option.

**Simple scheme**

Each backup is retained until its age exceeds a limit you specify. Then, it is either deleted or moved.

To set up deleting the backups:

- In **Retention rules**, select **Delete backups older than…**, and then specify the retention period.

To set up moving the backups:

- In **Retention rules**, select **Move backups older than…**, specify the retention period. Under **Where to replicate/move backups**, specify the location.

The retention rules are applied after creating a backup. For the second and next locations, creating a backup means copying or moving a backup there from the previous location.

**Grandfather-Father-Son (GFS) scheme**

Backups of each type (daily, weekly, and monthly) are retained for the periods you specify in **Keep backups**, and then deleted.

The retention rules are applied after creating a backup. They are applied sequentially in the primary, the second and all next locations.

**Tower of Hanoi scheme**

Each backup is retained based on its level (p. 41), and then deleted. You specify the number of levels in **Number of levels**.

The retention rules are applied after creating a backup. They are applied sequentially in the primary, the second and all next locations.

**Custom scheme**

Each backup is retained until the rules you specify are met. Then, it is either deleted or moved.

To set up deleting the backups:

- In **Clean up archive**, select **Using retention rules**. In the **Retention Rules** window (p. 69), specify the rules and select **If the specified conditions are met: Delete the oldest backups**.
- In **Apply retention rules**, specify when to apply the rules.

To set up moving the backups:

- In **Clean up archive**, select **Using retention rules**. In the **Retention Rules** window (p. 69), specify the rules and select **If the specified conditions are met: Move the oldest backups to another location**. Click **OK** and then specify the location under **Where to replicate/move backups**.
- In **Apply retention rules**, specify when to apply the rules.

You can choose to apply the retention rules before creating a backup, after creating a backup, on a schedule, or combine these options. For the second and next locations, creating a backup means copying or moving a backup there from the previous location.

## 4.5.4    Retention rules for the Custom scheme

In the **Retention Rules** window, you can select how long to store backups in the location and whether to move or delete them afterward.

The rules will be applied to all the backups taken on the *specific machine* and put in this *specific location* by this *specific backup plan*. In Acronis Backup & Recovery 11, such set of backups is called *an archive*.

***To set up retention rules for backups:***

1. Specify one of the following (options (a) and (b) are mutually exclusive):
   a. **Backups older than...** and/or **Archive size greater than...**.
      A backup will be stored until the specified condition (or both of the conditions) are met.
      **Example:**
      **Backups older than 5 days**
      **Archive size greater than 100 GB**
      With these settings, a backup will be stored until it is older than five days *and* the size of the archive containing it exceeds 100 GB.
   b. **Number of backups in the archive exceeds...**
      If the number of backups exceeds the specified value, one or more of the oldest backups will be moved or deleted. The minimal setting is 1.
2. Select whether to delete the backups or to move them to another location if the specified conditions are met.

You will be able to specify the location where to move the backups and set up retention rules for that location after you click **OK**.

### Deleting the last backup in the archive

The retention rules are effective if the archive contains more than one backup. This means that the last backup in the archive will be kept, even if a retention rule violation is detected. Please do not try to delete the only backup you have by applying the retention rules *before* backup. This will not work. Use the alternative setting **Clean up archive** > **When there is insufficient space while backing up** (p. 39) if you accept the risk of losing the last backup.

### Deleting or moving backups with dependencies

To access this setting, click **Show advanced settings** in the **Retention Rules** window.

Retention rules presume deleting or moving some backups while retaining the others. What if the archive contains incremental and differential backups that depend on each other and on the full backups they are based on? You cannot, say, delete an outdated full backup and keep its incremental "children".

When deletion or movement of a backup affects other backups, one of the following rules is applied:

- **Retain the backup until all dependent backups become subject to deletion (movement)**

  The outdated backup will be kept until all backups that depend on it also become outdated. Then, all the chain will be deleted at once during the regular cleanup. If you chose moving outdated backups to the next location, the backup will be copied there without delay. Only its deletion from the current location is postponed.

  This mode helps to avoid the potentially time-consuming consolidation but requires extra space for storing backups whose deletion is postponed. The archive size and/or the backup age or number can exceed the values you specify.

- **Consolidate these backups**

  The software will consolidate the backup that is subject to deletion or movement, with the next dependent backup. For example, the retention rules require to delete a full backup but to retain the next incremental one. The backups will be combined into a single full backup which will be dated with the incremental backup date. When an incremental or differential backup from the middle of the chain is deleted, the resulting backup type will be incremental.

  This mode ensures that after each cleanup the archive size and the age or number of backups are within the bounds you specify. The consolidation, however, may take a lot of time and system resources. You still need some extra space in the vault for temporary files created during consolidation.

  **What you need to know about consolidation**

  Please be aware that consolidation is just a method of deletion but not an alternative to deletion. The resulting backup will not contain data that was present in the deleted backup and was absent from the retained incremental or differential backup.

## 4.5.5 Replication/cleanup inactivity time

This option is effective only if you set up replication or retention rules (p. 66) for the backups.

This option defines a time period when starting replication or applying retention rules is not allowed. The operations will be performed when the inactivity time ends, if the machine is powered on at that

moment. The operations that had started before the inactivity time began continue without interruption.

The inactivity time affects all locations, including the primary one.

The preset is: **Disabled.**

To specify the inactivity time, select the **Do not start replication/cleanup within the following time** check box, and then select the days and the time period during the day.

**Usage example**

You may want to use this option to separate the backup process from replication or cleanup. For example, suppose that you back up machines locally during the day and replicate the backups to a network folder. Make the inactivity time contain the working hours. Replication will be performed after the working hours, when network load is lower.

## 4.5.6    Usage examples

This section provides examples of how you can replicate backups and set up retention rules for them.

### 4.5.6.1    Example 1. Replicating backups to a network folder

Consider the following scenario:

- You want to perform a full backup of your machine manually.
- You want to store the backups in Acronis Secure Zone (p. 139) on the machine.
- You want to store a copy of the backups in a network folder.

In this scenario, create a backup plan with the **Manual start** scheme. When creating the backup plan, specify Acronis Secure Zone in the **Path** field, select **Full** in the **Backup type** field, select the **Replicate just created backup to another location** check box, and then specify the network folder in the **2nd location** field.

**Result:**

- You can recover the machine's volumes or files from a readily available local backup, which is stored in a dedicated area of the hard disk.
- You can recover the machine from the network folder if the machine's hard disk drive fails.

### 4.5.6.2    Example 2. Limiting the age and total size of stored backups

Consider the following scenario:

- You want to perform a weekly full backup of your machine.
- You want to keep all backups that are younger than a month.
- You want to keep even older backups, as long as the total size of all backups stays below 200 GB.

In this scenario, create a backup plan with the **Custom** scheme. When creating the backup plan, specify a weekly schedule for the full backup. In **Clean up archive**, select **Using retention rules**.

Click **Retention rules**, select the **Backups older than** and the **Archive size greater than** check boxes, and specify respectively **1 month** and **200 GB**. In **If the specified conditions are met**, select **Delete the oldest backups**.

Click **OK**. In **Apply retention rules**, select the **After backup** check box.

**Result:**

- Backups that are younger than one month are kept, regardless of their total size.
- Backups that are older than one month are kept only if the total size of all backups (older plus younger) does not exceed 200 GB. Otherwise, the software deletes some or all of the older backups, starting from the oldest one.

# 4.6 Setting up regular conversion to a virtual machine

When creating a backup plan (p. 29), you can set up regular conversion of a disk or volume backup to a virtual machine. By setting up regular conversion, you obtain a copy of your server or workstation on a virtual machine which can be readily powered on in case the original machine fails.

**Restrictions**

- Conversion is not available on machines running Linux.
- Conversion of a backup from the following locations is not available: CD, DVD, tape devices and Acronis Online Backup Storage.
- If you have set up copying or moving backups to other locations (p. 66), you cannot set up conversion from any of these locations. Conversion is always performed from the first location of a backup.

The following section provides information that helps you make the appropriate settings.

**Convert to virtual machine**

### When to convert

Depending on the selected backup scheme, specify whether to convert every full, every incremental or every differential backup or convert the last created backup on schedule. Specify the **conversion schedule (p. 73)** if required.

### Browse (p. 73)

Specify the machine that will perform the conversion. The machine has to have Acronis Backup & Recovery 11 Agent for ESX(i), Agent for Hyper-V or Agent for Windows installed.

The conversion can be performed by the agent that performs the backup or by an agent installed on another machine. If the latter is the case, the archive must be stored in a shared location such as a network folder or a managed vault, so that the other machine can access the archive.

### VM type (p. 122)

Here you select the resulting virtual machine type and location. Available options depend on the host you selected in the previous step.

### Storage

Choose the storage on the virtualization server or the folder to place the virtual machine files in.

### Resultant VMs

Specify the name of the virtual machine. The default name is **Backup_of_[Machine Name]**. You can add more variables to the name. The following templates are supported:

[Plan Name]
[Machine Name]
[Virtual Host Name]

[Virtual Machine Name]
[Virtualization Server Type]

**Folder on VMware vCenter**

If the management server is integrated with vCenter Server, the resultant virtual machines will appear in the **Acronis Backups** folder on the vCenter. You can specify a subfolder for the machines resulting from execution of the plan.

## 4.6.1   Setting up a conversion schedule

A disk backup (p. 233) created while executing a backup plan can be converted to a virtual machine immediately, on schedule, or combining both methods.

The conversion task will be created on the machine being backed up, and will use this machine's date and time. If the agent that backs up the machine is installed outside it (such is the case when a ESX(i) or Hyper-V virtual machine is backed up at a hypervisor level), the task will be created on the machine where the agent is.

The target virtual machine must be powered off by the time of conversion, otherwise the conversion task will fail. If this happens, you can restart the conversion task manually after powering off the machine. Any changes made to the machine while it was powered on, will be overwritten.

## 4.6.2   Selecting a machine that will perform conversion

Take into account the following considerations.

**Which agent is installed on the machine?**

The resulting virtual machine type and location depend on the agent that resides on the selected machine.

- **Agent for ESX(i)** is installed on the machine

  If the agent manages more than one ESX(i) host, you can choose the host where the virtual machine will be created. This selection is done in the **VM type** (p. 122) step.

  In the **Storage** step, you can select the storage where the virtual machine will be created.

  Virtual machines created as a result of backup cannot be added to a backup plan. On the management server they appear as unmanageable or do not appear at all (if integration with vCenter Server is not enabled).

- **Agent for Hyper-V** is installed on the machine

  You can choose between creating a virtual machine on the Hyper-V server and creating a VMware Workstation, Microsoft Virtual PC or Parallels Workstation machine. This selection is done in the **VM type** (p. 122) step.

  In the **Storage** step, you can select the virtual machine path.

  Virtual machines created on the server as a result of backup do not appear on the management server, because such machines are not supposed to be backed up.

- **Agent for Windows** is installed on the machine

  You can choose the virtual machine type: VMware Workstation, Microsoft Virtual PC, or Parallels Workstation. This selection is done in the **VM type** (p. 122) step.

  In the **Storage** step, you can select the virtual machine path.

**What is the machine's processing power?**

Conversion will take the selected machine's CPU resource. Multiple conversion tasks will be queued on that machine and it may take considerable time to complete them all. Consider this when creating a centralized backup plan with conversion for multiple machines or multiple local backup plans using the same machine for conversion.

**What storage will be used for the virtual machines?**

**Network usage**

As opposed to ordinary backups (TIB files), virtual machine files are transferred uncompressed through the network. Therefore, using a SAN or a storage local to the machine that performs conversion is the best choice from the network usage standpoint. A local disk is not an option though, if the conversion is performed by the same machine that is backed up. Using a NAS also makes good sense.

**Storage space**

Disks of the resulting virtual machine will use as much storage space as the original data occupies. Assuming that the original disk size is 100 GB and the disk stores 10 GB of data, the corresponding virtual disk will occupy about 10 GB. VMware calls this format "thin provisioning", Microsoft uses the "dynamically expanding disk" term.

Since the space is not pre-allocated, the physical storage is expected to have sufficient free space for the virtual disks to increase in size.

## 4.6.3    How regular conversion to VM works

The way the repeated conversions work depends on where you choose to create the virtual machine.

- **If you choose to save the virtual machine as a set of files:** each conversion re-creates the virtual machine from scratch.
- **If you choose to create the virtual machine on a virtualization server:** when converting an incremental or differential backup, the software updates the existing virtual machine instead of re-creating it. Such conversion is normally faster. It saves network traffic and CPU resource of the host that performs the conversion. If updating the virtual machine is not possible, the software re-creates it from scratch.

The following is a detailed description of both cases.

**If you choose to save the virtual machine as a set of files**

As a result of the first conversion, a new virtual machine will be created. Every subsequent conversion will re-create this machine from scratch. First, a new (temporary) virtual machine is created. If this operation succeeds, the old machine is replaced. If an error occurs during creation of the temporary machine, the temporary machine is deleted. This way, the conversion always ends up with a single machine. However, extra storage space is required during conversion to store the temporary machine.

**If you choose to create the virtual machine on a virtualization server**

The first conversion creates a new virtual machine. Any subsequent conversion works as follows:

- If there has been a *full backup* since the last conversion, the virtual machine is re-created from scratch. This involves creating a temporary virtual machine, as described earlier in this section.

- Otherwise, the existing virtual machine is updated to reflect changes since the last conversion. If updating is not possible (for example, if you deleted the intermediate snapshots, see below), the virtual machine is re-created from scratch.

**Intermediate snapshots**

To be able to update the virtual machine, the software stores a few intermediate snapshots of it. They are named **Backup…** and **Replica…** and should be kept. Unneeded snapshots are deleted automatically.

The latest **Replica…** snapshot corresponds to the result of the latest conversion. You can go to this snapshot if you want to return the machine to that state; for example, if you worked with the machine and now want to discard the changes made to it.

Other snapshots are for internal use by the software.

# 4.7 Default backup options

Each Acronis agent has its own default backup options. Once an agent is installed, the default options have pre-defined values, which are referred to as **presets** in the documentation. When creating a backup plan, you can either use a default option, or override the default option with the custom value that will be specific for this plan only.

You can also customize a default option itself by changing its value against the pre-defined one. The new value will be used by default in all backup plans you will create later on this machine.

To view and change the default backup options, connect the console to the managed machine and then select **Options** > **Default backup and recovery options** > **Default backup options** from the top menu.

## Availability of the backup options

The set of available backup options depends on:

- The environment the agent operates in (Windows, bootable media)
- The type of the data being backed up (disk, file)
- The backup destination (networked location or local disk)
- The backup scheme (manual start or using the scheduler)

The following table summarizes the availability of the backup options.

| | Agent for Windows | | Bootable media (Linux-based or PE-based) | |
|---|---|---|---|---|
| | Disk backup | File backup | Disk backup | File backup |
| Additional settings (p. 77): | | | | |
| Ask for the first media while backing up to removable media | Dest: removable media | Dest: removable media | Dest: removable media | Dest: removable media |
| Use FTP in Active mode | Dest: FTP server | Dest: FTP server | Dest: FTP server | Dest: FTP server |

| | Agent for Windows | | Bootable media (Linux-based or PE-based) | |
|---|---|---|---|---|
| | Disk backup | File backup | Disk backup | File backup |
| Reset archive bit | - | + | - | + |
| Restart the machine automatically after backup is finished | - | - | + | + |
| Archive protection (p. 78) (password + encryption) | + | + | + | + |
| Backup cataloging (p. 79) | + | + | - | - |
| Backup performance: | | | | |
| Backup priority (p. 80) | + | + | - | - |
| HDD writing speed (p. 80) | Dest: HDD | Dest: HDD | Dest: HDD | Dest: HDD |
| Network connection speed (p. 80) | Dest: network share | Dest: network share | Dest: network share | Dest: network share |
| Backup splitting (p. 81) | + | + | + | + |
| Compression level (p. 82) | + | + | + | + |
| Disaster recovery plan (p. 82) | + | + | - | - |
| Error handling (p. 83): | | | | |
| Do not show messages and dialogs while processing (silent mode) | + | + | + | + |
| Re-attempt if an error occurs | + | + | + | + |
| Ignore bad sectors | + | + | + | + |
| Event tracing: | | | | |
| Windows events log (p. 84) | + | + | - | - |
| SNMP (p. 84) | + | + | - | - |
| Fast incremental/differential backup (p. 85) | + | - | + | - |
| File-level backup snapshot (p. 85) | - | + | - | - |
| File-level security (p. 86): | | | | |
| Preserve files' security settings in archives | - | + | - | - |
| In archives, store encrypted files in decrypted state | - | + | - | - |

| | Agent for Windows | | Bootable media (Linux-based or PE-based) | |
|---|---|---|---|---|
| | Disk backup | File backup | Disk backup | File backup |
| Media components (p. 86) | Dest: removable media | Dest: removable media | - | - |
| Mount points (p. 87) | - | + | - | - |
| Multi-volume snapshot (p. 88) | + | + | - | - |
| Notifications: | | | | |
| E-mail (p. 88) | + | + | - | - |
| Win Pop-up (p. 89) | + | + | - | - |
| Pre/Post backup commands (p. 90) | + | + | PE only | PE only |
| Pre/Post data capture commands (p. 91) | + | + | - | - |
| Replication/cleanup inactivity time (p. 70) | + | + | - | - |
| Sector-by-sector backup (p. 93) | + | - | + | - |
| Task failure handling (p. 94) | + | + | - | - |
| Task start conditions (p. 94) | + | + | - | - |
| Volume Shadow Copy Service (p. 95) | + | + | - | - |

## 4.7.1    Additional settings

Specify the additional settings for the backup operation by selecting or clearing the following check boxes.

**Ask for the first media while backing up to removable media**

This option is effective only when backing up to removable media.

The option defines whether to display the **Insert First Media** prompt when backing up to removable media.

The preset is: **Enabled**.

When the option is enabled, backing up to removable media may be not possible if the user is away, because the program will wait for someone to press OK in the prompt box. Hence, you should disable the prompt when scheduling a backup to removable media. Then, if the removable media is available (for example, a DVD is inserted), the task can run unattended.

## Reset archive bit

The option is effective only for file-level backup in Windows operating systems and in bootable media.

The preset is: **Disabled**.

In Windows operating systems, each file has the **File is ready for archiving** attribute, available by selecting **File** -> **Properties** -> **General** -> **Advanced** -> **Archive and Index attributes**. This attribute, also known as the archive bit, is set by the operating system each time the file is changed and can be reset by backup applications each time they include the file in a backup. The archive bit value is used by various applications such as databases.

When the **Reset archive bit** check box is selected, Acronis Backup & Recovery 11 will reset the archive bits of all files being backed up. Acronis Backup & Recovery 11 itself does not use the archive bit value. When performing incremental or differential backup, it determines whether a file has changed by the file size and the date/time when the file was last saved.

## Restart the machine automatically after backup is finished

This option is available only when operating under bootable media.

The preset is: **Disabled**.

When the option is enabled, Acronis Backup & Recovery 11 will restart the machine after the backup process is completed.

For example, if the machine boots from a hard disk drive by default and you select this check box, the machine will be restarted and the operating system will start as soon as the bootable agent has finished creating the backup.

## Use FTP in Active mode

The preset is: **Disabled**.

Enable this option if the FTP server supports active mode and you want this mode to be used for file transfers.

## 4.7.2    Archive protection

This option is effective for Windows and Linux operating systems and bootable media.

This option is effective for both disk-level and file-level backup.

This option defines whether the archive will be protected with a password and whether the archive's content will be encrypted.

This option is not available when the archive already contains backups. For example, this option may not be available:

- When you specify an already existing archive as the destination of the backup plan.
- When you edit a backup plan that has already resulted in a backup.

The preset is: **Disabled.**

***To protect the archive from unauthorized access***

1.   Select the **Set password for the archive** check box.

2. In the **Enter the password** field, type a password.

3. In the **Confirm the password** field, re-type the password.

4. Select one of the following:

   ▪ **Do not encrypt** – the archive will be protected with the password only

   ▪ **AES 128** – the archive will be encrypted using the Advanced Encryption Standard (AES) algorithm with a 128-bit key

   ▪ **AES 192** – the archive will be encrypted using the AES algorithm with a 192-bit key

   ▪ **AES 256** – the archive will be encrypted using the AES algorithm with a 256-bit key.

5. Click **OK**.

The AES cryptographic algorithm operates in the Cipher-block chaining (CBC) mode and uses a randomly generated key with a user-defined size of 128, 192 or 256 bits. The larger the key size, the longer it will take for the program to encrypt the archive and the more secure your data will be.

The encryption key is then encrypted with AES-256 using a SHA-256 hash of the password as a key. The password itself is not stored anywhere on the disk or in the backup file; the password hash is used for verification purposes. With this two-level security, the backup data is protected from any unauthorized access, but recovering a lost password is not possible.

## 4.7.3 Backup cataloging

Cataloging a backup adds the contents of the backup to the data catalog. Using the data catalog, you can easily find the required version of data and select it for recovery.

The **Backup cataloging** option defines whether the backups will be cataloged automatically as soon as they are created.

The preset is: **Enabled.**

After the cataloging has been completed, the catalog will show all data contained in the just created backup, namely:

▪ For a disk-level backup - disks, volumes, files and folders.

▪ For a file-level backup - files and folders.

You may want to disable the automatic cataloging if it tends to affect the performance of the managed machine, or your backup window is too narrow. If the **Backup cataloging** option is disabled, the following data will be displayed in the catalog:

▪ For a disk-level backup - only disks and volumes.

▪ For a file-level backup - nothing.

To add the full content of already existing backups to the catalog, you can start the cataloging manually when appropriate.

For more information about using data catalog, see the Data catalog (p. 101) section.

## 4.7.4 Backup performance

Use this group of options to specify the amount of network and system resources to allocate to the backup process.

Backup performance options might have a more or less noticeable effect on the speed of the backup process. This depends on the overall system configuration and the physical characteristics of devices the backup is being performed from or to.

## 4.7.4.1    Backup priority

This option is effective for both Windows and Linux operating systems.

The priority of a process running in a system determines the amount of CPU and system resources allocated to that process. Decreasing the backup priority will free more resources for other applications. Increasing the backup priority might speed up the backup process by requesting the operating system to allocate more resources like the CPU to the backup application. However, the resulting effect will depend on the overall CPU usage and other factors like disk in/out speed or network traffic.

The preset is: **Low.**

***To specify the backup process priority***

Select one of the following:

- **Low** – to minimize resources taken by the backup process, leaving more resources to other processes running on the machine
- **Normal** – to run the backup process with normal speed, allocating resources on a par with other processes
- **High** – to maximize the backup process speed by taking resources from other processes.

## 4.7.4.2    HDD writing speed

This option is effective for Windows and Linux operating systems and bootable media.

This option is available when an internal (fixed) hard disk of the machine being backed up is selected as the backup destination

Backing up to a fixed hard disk (for example, to Acronis Secure Zone) may slow performance of the operating system and applications because of the large amounts of data that needs to be written to the disk. You can limit the hard disk usage by the backup process to the desired level.

The preset is: **Maximum.**

***To set the desired HDD writing speed for backup***

Do any of the following:

- Click **Writing speed stated as a percentage of the maximum speed of the destination hard disk**, and then drag the slider or select a percentage in the box
- Click **Writing speed stated in kilobytes per second**, and then enter the writing speed in kilobytes per second.

## 4.7.4.3    Network connection speed

This option is effective for Windows and Linux operating systems and bootable media.

This option is available when a location on the network (network share, managed vault or an FTP/SFTP server) is selected as the backup destination.

The option defines the amount of network connection bandwidth allocated for transferring the backup data.

By default the speed is set to maximum, i.e. the software uses all the network bandwidth it can get when transferring the backup data. Use this option to reserve a part of the network bandwidth to other network activities.

The preset is: **Maximum.**

***To set the network connection speed for backup***

Do any of the following:

- Click **Transferring speed stated as a percentage of the estimated maximum speed of the network connection**, and then drag the slider or type a percentage in the box
- Click **Transferring speed stated in kilobytes per second**, and then enter the bandwidth limit for transferring backup data in kilobytes per second.

## 4.7.5    Backup splitting

This option is effective for Windows and Linux operating systems and bootable media.

The option defines how a backup can be split.

The preset is: **Automatic.**

The following settings are available.

### Automatic

With this setting, Acronis Backup & Recovery 11 will act as follows.

- **When backing up to a hard disk:**

  A single backup file will be created if the destination disk's file system allows the estimated file size.

  The backup will automatically be split into several files if the destination disk's file system does not allow the estimated file size. Such might be the case when the backup is placed on FAT16 and FAT32 file systems that have a 4GB file size limit.

  If the destination disk runs out of free space while creating the backup, the task enters the **Need interaction** state. You have the ability to free additional space and retry the operation. If you do so, the resulting backup will be split into the parts created before and after the retry.

- **When backing up to removable media** (CD, DVD or a tape device locally attached to the managed machine):

  The task will enter the **Need interaction** state and ask for a new media when the previous one is full.

### Fixed size

Enter the desired file size or select it from the drop-down list. The backup will then be split into multiple files of the specified size. This comes in handy when creating a backup that you plan to burn to multiple CDs or DVDs later on. You might also want to split the backup destined to an FTP server, since data recovery directly from an FTP server requires the backup to be split into files no more than 2GB in size.

## 4.7.6    Compression level

This option is effective for Windows and Linux operating systems and bootable media.

The option defines the level of compression applied to the data being backed up.

The preset is: **Normal.**

The optimal data compression level depends on the type of data being backed up. For example, even maximum compression will not significantly reduce the archive size if the archive contains essentially compressed files, such as .jpg, .pdf or .mp3. However, formats such as .doc or .xls will be compressed well.

### *To specify the compression level*

Select one of the following:

- **None** – the data will be copied as is, without any compression. The resulting backup size will be maximal.

- **Normal** – recommended in most cases.

- **High** – the resulting backup size will typically be less than for the **Normal** level.

- **Maximum** – the data will be compressed as much as possible. The backup duration will be maximal. You may want to select maximum compression when backing up to removable media to reduce the number of blank disks required.

## 4.7.7    Disaster recovery plan (DRP)

This option is effective for Windows and Linux but is not applicable to bootable media.

Disaster recovery plan (DRP) contains a list of backed up data items and detailed instructions that guide a user through a process of recovering these items from a backup.

If the **Disaster recovery plan (DRP)** option is enabled, a DRP is created and sent by e-mail to the specified list of users after the first successful backup performed by the backup plan. The DRP will be created and sent again after the first successful backup in the following cases:

- The backup plan has been edited so that the DRP parameters changed.
- The backup contains new data items or does not contain items previously backed up. (This does not apply to such data items as files or folders.)

If multiple machines are protected by a backup plan, then a separate DRP is sent for each machine.

**DRP and post-backup commands**

Note that the DRP will not automatically change if post-backup commands in your backup plan copy or move the backups from the original location. The DRP points only to the locations specified in the backup plan.

**Adding information to a DRP template**

You can append additional information to a DRP template if you are well familiar with XML and HTML. The default paths to the DRP template are:

- **%ProgramFiles%\Acronis\BackupAndRecovery\drp.xsl** - in 32-bit Windows
- **%ProgramFiles(x86)%\Acronis\BackupAndRecovery\drp.xsl -** in 64-bit Windows
- **/usr/lib/Acronis/BackupAndRecovery/drp.xsl -** in Linux

***To set up sending DRPs:***

1. Select the **Send disaster recovery plan** check box.

2. Enter the e-mail address in the **E-mail Address** field. You can enter several e-mail addresses in a semicolon-delimited format.

3. [Optional] Change the default value of the **Subject** field, if necessary.

   If you back up multiple machines with one centralized backup plan and want each machine user to receive a separate DRP e-mail about his/her machine only:

   a. Use the *%MachineName%* variable to show the name of the certain machine in the e-mail subject.

   b. Set up your mail server or client to filter or forward e-mails using the **Subject** field.

4. Enter the parameters of access to the SMTP server. For more detailed information, see E-mail notifications (p. 131).

5. [Optional] Click **Send test e-mail message** to check if the settings are correct.

## 4.7.8    Error handling

These options are effective for Windows and Linux operating systems and bootable media.

These options enable you to specify how to handle errors that might occur during backup.

### Do not show messages and dialogs while processing (silent mode)

The preset is: **Disabled**.

With the silent mode enabled, the program will automatically handle situations requiring user interaction (except for handling bad sectors, which is defined as a separate option). If an operation cannot continue without user interaction, it will fail. Details of the operation, including errors, if any, can be found in the operation log.

### Re-attempt, if an error occurs

The preset is: **Enabled**. **Number of attempts: 30. Interval between attempts: 30 seconds.**

When a recoverable error occurs, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts is performed, depending on which comes first.

For example, if the backup destination on the network becomes unavailable or not reachable, the program will attempt to reach the destination every 30 seconds, but no more than 5 times. The attempts will be stopped as soon as the connection is resumed OR the specified number of attempts is performed, depending on which comes first.

### Ignore bad sectors

The preset is: **Disabled**.

When the option is disabled, the program will display a pop-up window each time it comes across a bad sector and ask for a user decision as to whether to continue or stop the backup procedure. In order to back up the valid information on a rapidly dying disk, enable ignoring bad sectors. The rest of the data will be backed up and you will be able to mount the resulting disk backup and extract valid files to another disk.

## 4.7.9    Event tracing

It is possible to duplicate log events of the backup operations, performed on the managed machine, in the Application Event Log of Windows; or send the events to the specified SNMP managers.

### 4.7.9.1    SNMP notifications

This option is effective for both Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option defines whether the agent(s) operating on the managed machine have to send the log events of the backup operations to the specified Simple Network Management Protocol (SNMP) managers. You can choose the types of events to be sent.

For detailed information about using SNMP with Acronis Backup & Recovery 11, please see "Support for SNMP (p. 28)".

The preset is: **Use the setting set in the Machine options.**

***To select whether to send the backup operations events to the SNMP managers:***

Choose one of the following:

- **Use the setting set in the Machine options** – to use the setting specified for the machine. For more information refer to Machine options.

- **Send SNMP notifications individually for backup operation events** – to send the events of the backup operations to the specified SNMP managers.

    - **Types of events to send** – choose the types of events to be sent: **All events**, **Errors and warnings**, or **Errors only**.

    - **Server name/IP** – type the name or IP address of the host running the SNMP management application, the messages will be sent to.

    - **Community** – type the name of the SNMP community to which both the host running the SNMP management application and the sending machine belong. The typical community is "public".

    Click **Send test message** to check if the settings are correct.

- **Do not send SNMP notifications** – to disable sending the log events of the backup operations to SNMP managers.

### 4.7.9.2    Windows event log

This option is effective only in Windows operating systems.

This option is not available when operating under the bootable media.

This option defines whether the agent(s) operating on the managed machine have to log events of the backup operations in the Application Event Log of Windows (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

The preset is: **Use the setting set in the Machine options.**

***To select whether to log the backup operations events in the Application Event Log of Windows:***

Choose one of the following:

- **Use the setting set in the Machine options** – to use the setting specified for the machine. For more information refer to Machine options.
- **Log the following event types** – to log events of the backup operations in the Application Event Log. Specify the types of events to be logged:
    - **All events** – log all events (information, warnings and errors)
    - **Errors and warnings**
    - **Errors only**
- **Do not log** - to disable logging events of the backup operations in the Application Event Log.

## 4.7.10 Fast incremental/differential backup

The option is effective in Windows and Linux operating systems and bootable media.

This option is effective for incremental and differential disk-level backup.

This option defines whether a file change is detected using the file size and time stamp or by comparing the file contents to those stored in the archive.

The preset is: **Enabled.**

Incremental or differential backup captures only data changes. To speed up the backup process, the program determines whether a file has changed or not by the file size and the date/time when the file was last modified. Disabling this feature will make the program compare the entire file contents to those stored in the archive.

## 4.7.11 File-level backup snapshot

This option is effective only for file-level backup in Windows and Linux operating systems.

This option defines whether to back up files one by one or by taking an instant data snapshot.

*Note: Files that are stored on network shares are always backed up one by one.*

The preset is: **Create snapshot if it is possible.**

Select one of the following:

- **Always create a snapshot**

    The snapshot enables backing up of all files including files opened for exclusive access. The files will be backed up at the same point in time. Choose this setting only if these factors are critical, that is, backing up files without a snapshot does not make sense. To use a snapshot, the backup plan has to run under the account with the Administrator or Backup Operator privileges. If a snapshot cannot be taken, the backup will fail.

- **Create a snapshot if it is possible**

    Back up files directly if taking a snapshot is not possible.

- **Do not create a snapshot**

Always back up files directly. Administrator or Backup Operator privileges are not required. Trying to back up files that are opened for exclusive access will result in a read error. Files in the backup may be not time-consistent.

## 4.7.12 File-level security

These options are effective only for file-level backup in Windows operating systems.

**In archives, store encrypted files in a decrypted state**

This option defines whether to decrypt files before saving them to a backup archive.

The preset is: **Disabled**.

Simply ignore this option if you do not use the encryption. Enable the option if encrypted files are included in the backup and you want them to be accessed by any user after recovery. Otherwise, only the user who encrypted the files/folders will be able to read them. Decryption may also be useful if you are going to recover encrypted files on a different machine.

*File encryption is available in Windows using the NTFS file system with the Encrypting File System (EFS). To access a file or folder encryption setting, select **Properties > General > Advanced Attributes > Encrypt contents to secure data.***

**Preserve file security settings in archives**

This option defines whether to back up NTFS permissions for files along with the files.

The preset is: **Enabled.**

When the option is enabled, files and folders are saved in the archive with the original permissions to read, write or execute the files for each user or user group. If you recover a secured file/folder on a machine without the user account specified in the permissions, you may not be able to read or modify this file.

To completely eliminate this kind of problem, disable preserving file security settings in archives. The recovered files and folders will always inherit the permissions from the folder to which they are recovered or from the disk, if recovered to the root.

Alternatively, you can disable recovery (p. 130) of the security settings, even if they are available in the archive. The result will be the same - the files will inherit the permissions from the parent folder.

*To access file or folder NTFS permissions, select **Properties > Security**.*

## 4.7.13 Media components

This option is effective for both Windows and Linux operating systems, when the backup destination is removable media.

When backing up to removable media, you can make this media work as regular Linux-based bootable media (p. 229) by writing additional components to it. As a result, you will not need a separate rescue disc.

The preset is: **No bootable components.**

Choose one of the following components you want to put on the bootable media:

- **Acronis Bootable Agent** is a bootable rescue utility (based on Linux kernel) that includes most of the functionality of the Acronis Backup & Recovery 11 agent. Put this component on the media if you want more functionality during recovery. You will be able to configure the recovery operation in the same way as under regular bootable media; use Active Restore or Universal Restore. If the media is being created in Windows, the disk management functionality will also be available.

- **Acronis Bootable Agent and One-Click Restore**. The One-Click Restore is the minimal addition to a disk backup stored on removable media, allowing for easy recovery from this backup. If you boot a machine from the media and click **Run Acronis One-click Restore,** the disk will be immediately recovered from the backup contained on the same media.

*Caution: Because the one-click approach does not presume user selections, such as selecting volumes to recover, Acronis One-Click Restore always recovers the entire disk. If your disk contains several volumes and you are planning to use Acronis One-Click Restore, include all the volumes in the backup. Any volumes missing from the backup will be lost.*

## 4.7.14 Mount points

This option is effective only in Windows for a file-level backup of a data source that includes mounted volumes or cluster shared volumes.

This option is effective only when you select for backup a folder that is higher in the folder hierarchy than the mount point. (A mount point is a folder on which an additional volume is logically attached.)

- If such folder (a parent folder) is selected for backup, and the **Mount points** option is enabled, all files located on the mounted volume will be included in the backup. If the **Mount points** option is disabled, the mount point in the backup will be empty.

  During recovery of a parent folder, the mount point content will or will not be recovered, depending on whether the **Mount points** option for recovery (p. 130) is enabled or disabled.

- If you select the mount point directly, or select any folder within the mounted volume, the selected folders will be considered as ordinary folders. They will be backed up regardless of the state of the **Mount points** option and recovered regardless of the state of the **Mount points** option for recovery (p. 130).

The preset is: **Disabled**.

*Tip. You can back up Hyper-V virtual machines residing on a cluster shared volume by backing up the required files or the entire volume with file-level backup. Just power off the virtual machines to be sure that they are backed up in a consistent state.*

**Example**

Let's assume that the **C:\Data1\** folder is a mount point for the mounted volume. The volume contains folders **Folder1** and **Folder2**. You create a backup plan for file-level backup of your data.

If you select the check box for volume C and enable the **Mount points** option, the **C:\Data1\** folder in your backup will contain **Folder1** and **Folder2**. When recovering the backed up data, be aware of proper using the **Mount points** option for recovery (p. 130).

If you select the check box for volume C, and disable the **Mount points** option, the **C:\Data1\** folder in your backup will be empty.

If you select the check box for the **Data1**, **Folder1** or **Folder2** folder, the checked folders will be included in the backup as ordinary folders, regardless of the state of the **Mount points** option.

## 4.7.15 Multi-volume snapshot

This option is effective only for Windows operating systems.

This option applies to disk-level backup. This option also applies to file-level backup when the file-level backup is performed by taking a snapshot. (The File-level backup snapshot (p. 85) option determines whether a snapshot will be taken during file-level backup).

The option determines whether to take snapshots of multiple volumes at the same time or one by one.

The preset is: **Enable.**

When this option is set to **Enable**, snapshots of all volumes being backed up will be created simultaneously. Use this option to create a time-consistent backup of data spanned across multiple volumes, for instance for an Oracle database.

When this option is set to **Disable**, the volumes' snapshots will be taken one after the other. As a result, if the data spans across several volumes, the resulting backup may be not consistent.

## 4.7.16 Notifications

Acronis Backup & Recovery 11 provides the ability of notifying users about backup completion through e-mail or the messaging service.

### 4.7.16.1 E-mail

This option is effective for Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option enables you to receive e-mail notifications about the backup task's successful completion, failure or need for interaction along with the full log of the task.

The preset is: **Disabled**.

***To configure e-mail notification***

1.  Select the **Send e-mail notifications** check box to activate notifications.
2.  Under **Send e-mail notifications**, select the appropriate check boxes as follows:
    *   **When backup completes successfully** – to send notification when the backup task has completed successfully
    *   **When backup fails** – to send a notification when the backup task has failed
    *   **When user interaction is required** – to send to send notification during the operation when user interaction is required.
3.  For the e-mail message to include the log entries related to the backup, select the **Add full log to the notification** check box.
4.  In the **E-mail addresses** field, type the e-mail address to which notifications will be sent. You can enter several addresses separated by semicolons.
5.  In the **Subject** field, type the notification subject or leave the default value.
6.  In the **SMTP server** field**,** enter the name of the SMTP server.
7.  In the **Port** field– set the port of the SMTP server. By default, the port is set to **25**.

8. In the **User name** field, enter the user name.

9. In the **Password** field, enter the password.

10. Click **Additional e-mail parameters...** to configure the additional e-mail parameters as follows:

    a. **From** - type the e-mail address of the user from whom the message will be sent. If you leave this field empty, messages will be constructed as if they are from the destination address.

    b. **Use encryption** – you can opt for encrypted connection to the mail server. SSL and TLS encryption types are available for selection.

    c. Some Internet service providers require authentication on the incoming mail server before being allowed to send something. If this is your case, select the **Log on to incoming mail server** check box to enable a POP server and to set up its settings:

       ▪ **Incoming mail server (POP)** – enter the name of the POP server.

       ▪ **Port** – set the port of the POP server. By default, the port is set to **110**.

       ▪ **User name** – enter the user name.

       ▪ **Password** – enter the password.

    d. Click **OK**.

11. Click **Send test e-mail message** to check if the settings are correct.

## 4.7.16.2    Messenger service (WinPopup)

This option is effective for Windows and Linux operating systems on the sending machine and only for Windows on the receiving machine.

This option is not available when operating under bootable media.

The option enables you to receive WinPopup notifications about the backup task's successful completion, failure or need for interaction.

The preset is: **Disabled.**
Before configuring WinPopup notifications, make sure the Messenger service is started on both the machine executing the task and the machine that will receive messages.

The Messenger service is not started by default in the Microsoft Windows Server 2003 family. Change the service Startup mode to Automatic and start the service.

*To configure WinPopup notifications:*

1. Select the **Send WinPopup notifications** check box.

2. In the **Machine name** field, enter the name of the machine to which notifications will be sent. Multiple names are not supported.

Under **Send notifications**, select the appropriate check boxes as follows:

   ▪ **When backup completes successfully** – to send notification when the backup operation is completed successfully

   ▪ **When backup fails** – to send notification when the backup operation is failed

   ▪ **When user interaction is required** – to send notification during the operation when user interaction is required.

Click **Send test WinPopup message** to check if the settings are correct.

# 4.7.17 Pre/Post commands

This option is effective for Windows and Linux operating systems and PE-based bootable media.

The option enables you to define the commands to be automatically executed before and after the backup procedure.

The following scheme illustrates when pre/post commands are executed.

| Pre-backup command | Backup | Post-backup command |
|---|---|---|

Examples of how you can use the pre/post commands:

- Delete some temporary files from the disk before starting backup
- Configure a third-party antivirus product to be started each time before the backup starts
- Copy an archive to another location after the backup ends.

The program does not support interactive commands, i.e. commands that require user input (for example, "pause").

### To specify pre/post commands

1. Enable pre/post commands execution by checking the following options:
   - **Execute before the backup**
   - **Execute after the backup**
2. Do any of the following:
   - Click **Edit** to specify a new command or a batch file
   - Select the existing command or the batch file from the drop-down list
3. Click **OK**.

## 4.7.17.1 Pre-backup command

### To specify a command/batch file to be executed before the backup process starts

1. In the **Command** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)
2. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
3. In the **Arguments** field specify the command's execution arguments, if required.
4. Depending on the result you want to obtain, select the appropriate options as described in the table below.
5. Click **Test command** to check if the command is correct.

| Check box | Selection | | | |
|---|---|---|---|---|
| **Fail the task if the command execution fails\*** | Selected | Cleared | Selected | Cleared |
| **Do not back up until the command execution is complete** | Selected | Selected | Cleared | Cleared |
| **Result** | | | | |
| | **Preset** Perform the backup | Perform the backup after the | N/A | Perform the backup concurrently with |

| | only after the command is successfully executed. Fail the task if the command execution fails. | command is executed despite execution failure or success. | | the command execution and irrespective of the command execution result. |
|---|---|---|---|---|

* A command is considered failed if its exit code is not equal to zero.

## 4.7.17.2   Post-backup command

***To specify a command/executable file to be executed after the backup is completed***

1.  In the **Command** field, type a command or browse to a batch file.

2.  In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.

3.  In the **Arguments** field, specify the command execution arguments, if required.

4.  Select the **Fail the task if the command execution fails** check box if successful execution of the command is critical for you. The command is considered failed if its exit code is not equal to zero. If the command execution fails, the program will remove the resulting TIB file and temporary files if possible, and the task run result will be set to Failed.

    When the check box is not selected, the command execution result does not affect the task execution failure or success. You can track the command execution result by exploring the log or the errors and warnings displayed in the **Log** view.

5.  Click **Test Command** to check if the command is correct.

## 4.7.18   Pre/Post data capture commands

This option is effective for both Windows and Linux operating systems.

The option enables you to define the commands to be automatically executed before and after data capture (that is, taking the data snapshot). Data capture is performed by Acronis Backup & Recovery 11 at the beginning of the backup procedure.

The following scheme illustrates when the pre/post data capture commands are executed.

| | <--------------- | **Backup** | ---------------> | | |
|---|---|---|---|---|---|
| Pre-backup command | Pre-data capture command | Data capture | Post-data capture command | | Post-backup command |

If the Volume Shadow Copy Service (p. 95) option is enabled, the commands' execution and the Microsoft VSS actions will be sequenced as follows:

"Before data capture" commands -> VSS Suspend -> Data capture -> VSS Resume -> "After data capture" commands.

Using the pre/post data capture commands, you can suspend and resume a database or application that is not compatible with VSS. As opposed to the Pre/Post commands (p. 90), the pre/post data capture commands will be executed before and after the data capture process. This takes seconds. The entire backup procedure may take much longer, depending on the amount of data to be backed up. Therefore, the database or application idle time will be minimal.

***To specify pre/post data capture commands***

1. Enable pre/post data capture commands execution by checking the following options:
   - **Execute before the data capture**
   - **Execute after the data capture**

2. Do any of the following:
   - Click **Edit** to specify a new command or a batch file
   - Select the existing command or the batch file from the drop-down list

3. Click **OK**.

## 4.7.18.1   Pre-data capture command

***To specify a command/batch file to be executed before data capture***

1. In the **Command** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)

2. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.

3. In the **Arguments** field specify the command's execution arguments, if required.

4. Depending on the result you want to obtain, select the appropriate options as described in the table below.

5. Click **Test command** to check if the command is correct.

| Check box | Selection | | | |
|---|---|---|---|---|
| **Fail the backup task if the command execution fails\*** | Selected | Cleared | Selected | Cleared |
| **Do not perform the data capture until the command execution is complete** | Selected | Selected | Cleared | Cleared |
| **Result** | | | | |
| | **Preset**<br><br>Perform the data capture only after the command is successfully executed. Fail the task if the command execution fails. | Perform the data capture after the command is executed despite execution failure or success. | N/A | Perform the data capture concurrently with the command and irrespective of the command execution result. |

\* A command is considered failed if its exit code is not equal to zero.

## 4.7.18.2   Post-data capture command

***To specify a command/batch file to be executed after data capture***

1. In the **Command** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)

2. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.

3. In the **Arguments** field specify the command's execution arguments, if required.

4. Depending on the result you want to obtain, select the appropriate options as described in the table below.

5. Click **Test command** to check if the command is correct.

| Check box | Selection | | | |
|---|---|---|---|---|
| **Fail the task if the command execution fails*** | Selected | Cleared | Selected | Cleared |
| **Do not back up until the command execution is complete** | Selected | Selected | Cleared | Cleared |
| Result | | | | |
| | **Preset**<br><br>Continue the backup only after the command is successfully executed. Delete the TIB file and temporary files and fail the task if the command execution fails. | Continue the backup after the command is executed despite command execution failure or success. | N/A | Continue the backup concurrently with the command execution and irrespective of the command execution result. |

* A command is considered failed if its exit code is not equal to zero.

## 4.7.19  Replication/cleanup inactivity time

This option is effective only if you set up replication or retention rules (p. 66) for the backups.

This option defines a time period when starting replication or applying retention rules is not allowed. The operations will be performed when the inactivity time ends, if the machine is powered on at that moment. The operations that had started before the inactivity time began continue without interruption.

The inactivity time affects all locations, including the primary one.

The preset is: **Disabled.**

To specify the inactivity time, select the **Do not start replication/cleanup within the following time** check box, and then select the days and the time period during the day.

**Usage example**

You may want to use this option to separate the backup process from replication or cleanup. For example, suppose that you back up machines locally during the day and replicate the backups to a network folder. Make the inactivity time contain the working hours. Replication will be performed after the working hours, when network load is lower.

## 4.7.20  Sector-by-sector backup

The option is effective only for disk-level backup.

To create an exact copy of a disk or volume on a physical level, select the **Back up sector-by-sector** check box. The resulting backup will be equal in size to the disk being backed up (if the **Compression level** (p. 82) option is set to **None**). Use the sector-by-sector backup for backing up drives with unrecognized or unsupported file systems and other proprietary data formats.

## 4.7.21 Task failure handling

This option is effective for Windows and Linux operating systems.

This option is not available when operating under the bootable media.

This option determines the program behavior when any of the backup plan's tasks fails.

The preset is **not to restart a failed task.**

The program will try to execute the failed task again if you select the **Restart a failed task** check box and specify the number of attempts and the time interval between the attempts. The program stops trying as soon as an attempt completes successfully OR the specified number of attempts is performed, depending on which comes first.



*If the task fails because of a mistake in the backup plan, you can edit the plan while the task is in the Idle state. While the task is running, you have to stop it prior to editing the backup plan.*

## 4.7.22 Task start conditions

This option is effective in Windows and Linux operating systems.

This option is not available when operating under bootable media.

This option determines the program behavior in case a backup task is about to start (the scheduled time comes or the event specified in the schedule occurs), but the condition (or any of multiple conditions) is not met. For more information on conditions please see Scheduling (p. 53) and Conditions (p. 63).

The preset is: **Wait until the conditions are met.**

**Wait until the conditions are met**

With this setting, the scheduler starts monitoring the conditions and launches the task as soon as the conditions are met. If the conditions are never met, the task will never start.

To handle the situation when the conditions are not met for too long and further delaying the backup is becoming risky, you can set the time interval after which the task will run irrespective of the condition. Select the **Run the task anyway after** check box and specify the time interval. The task will start as soon as the conditions are met OR the maximum time delay lapses, depending on which comes first.

**Time diagram: Wait until conditions are met**



## Skip the task execution

Delaying a backup might be unacceptable, for example, when you need to back up data strictly at the specified time. Then it makes sense to skip the backup rather than wait for the conditions, especially if the events occur relatively often.

# 4.7.23  Volume Shadow Copy Service

This option is effective only for Windows operating systems.

The option defines whether a Volume Shadow Copy Service (VSS) provider has to notify VSS-aware applications that the backup is about to start. This ensures the consistent state of all data used by the applications; in particular, completion of all database transactions at the moment of taking the data snapshot by Acronis Backup & Recovery 11. Data consistency, in turn, ensures that the application will be recovered in the correct state and become operational immediately after recovery.

The preset is: **Use Volume Shadow Copy Service**.

## Use VSS

When **Use Volume Shadow Copy Service** is selected, choose a snapshot provider from the following list:

- **Hardware/software - Select automatically** (selected by default)

    VSS will use the hardware-based provider that supports the source volume. If none is found, VSS will use the Microsoft Software Shadow Copy provider.

- **Software - Select automatically**

    In most cases, VSS will use the Microsoft Software Shadow Copy provider.

- **Software - Acronis VSS Provider**

    VSS will use Acronis VSS Provider for taking snapshots.

- **Software - System provider**

    VSS will use the system provider (Microsoft Software Shadow Copy provider) for taking snapshots.

- **Software - A software provider**

    VSS will use any available software-based provider that supports the source volume.

- **Hardware - Select automatically**

    VSS will use the hardware-based provider that supports the source volume. If no hardware-based provider is found, snapshots will be created by means of Acronis Backup & Recovery 11.

## Do not use VSS

If you select **Do not use VSS**,    the data snapshot will be taken by Acronis Backup & Recovery 11.

Choose **Do not use VSS** if your database is incompatible with VSS. Backup process is fastest, but data consistency of the applications whose transactions are not completed at the time of taking a snapshot cannot be guaranteed. You may use Pre/Post data capture commands (p. 91) to indicate which commands should be performed before and after taking the snapshot. This ensures that the data is being backed up in a consistent state. For instance, specify pre-data capture commands that will suspend the database and flush all caches to ensure that all transactions are completed; and specify post-data capture commands that will resume the database operations after the snapshot is taken.

## Volume shadow copy writers

Before backing up the data of VSS-aware applications, make sure that the volume shadow copy writers for those applications are turned on by examining the list of writers that are present in the operating system. To view this list, run the following command:

```
vssadmin list writers
```

*Note: In Microsoft Windows Small Business Server 2003, the writer for Microsoft Exchange Server 2003 is turned off by default. For instructions on how to turn it on, see the corresponding Microsoft Help and Support article (http://support.microsoft.com/kb/838183/en).*

# 5 Recovery

When it comes to data recovery, first consider the most functional method: connect the console to the managed **machine running the operating system** and create the recovery task.

If the machine's **operating system fails to start** or you need to **recover data to bare metal**, boot the machine from the bootable media (p. 229) or using Acronis Startup Recovery Manager and configure recovery.

Acronis Universal Restore lets you recover and boot up operating systems on **dissimilar hardware** or a virtual machine.

Acronis Backup & Recovery 11 enables you to transfer Windows operating systems between BIOS-based hardware and hardware that supports Unified Extensible Firmware Interface (UEFI). See the "Recovering BIOS-based systems to UEFI-based and back" (p. 114) section for more details.

A **Windows system can be brought online in seconds** while it is still being recovered. Using the proprietary Acronis Active Restore (p. 118) technology, Acronis Backup & Recovery 11 will boot the machine into the operating system found in the backup as if the system were on the physical disk. The system becomes operational and ready to provide necessary services. Thus, the system downtime will be minimal.

A dynamic volume can be recovered over an existing volume, to unallocated space of a disk group, or to unallocated space of a basic disk. To learn more about recovering dynamic volumes, please turn to the "Backup and recovery of dynamic volumes (Windows)" (p. 25) section.

Acronis Backup & Recovery 11 Agent for Windows has the ability to recover a disk (volume) backup to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat Kernel-based Virtual Machine (KVM). The virtual appliance can then be imported to XenServer. You can move the files of the Kernel-based Virtual Machine to a machine running Linux; from there, you can run this virtual machine by using the Virtual Machine Manager program. The VMware Workstation machine can be converted to the open virtualization format (OVF) using the VMware OVF tool. With Acronis Backup & Recovery 11 Agent for Hyper-V or Agent for ESX(i), you can create a new virtual machine on the respective virtualization server.

You might need to prepare target disks before recovery. Acronis Backup & Recovery 11 includes a handy disk management utility which enables you to create or delete volumes, change a disk partitioning style, create a disk group and perform other disk management operations on the target hardware, both under the operating system and on bare metal. To find out more about Acronis Disk Director LV, see "Disk management" (p. 172).

## 5.1 Creating a recovery task

***To create a recovery task, perform the following steps***

**What to recover**

> **Select data** (p. 99)
>> Select data to recover.
> **Access credentials** (p. 102)

[Optional] Provide credentials for the archive location if the task account does not have the right to access it. To access this option, click **Show access credentials**.

## Where to recover

This section appears after the required backup is selected and the type of data to recover is defined. The parameters you specify here depend on the type of data being recovered.

**Disks** (p. 103)

**Volumes** (p. 105)

**Files** (p. 109)

### Acronis Active Restore

[Optional] The **Acronis Active Restore** check box is available when recovering Windows starting from Windows 2000. Acronis Active Restore brings a system online immediately after the recovery is started. The operating system boots from the backup image and the machine becomes operational and ready to provide necessary services. The data required to serve incoming requests is recovered with the highest priority; everything else is recovered in the background.

For details, see "Acronis Active Restore" (p. 118).

### Access credentials (p. 103)

[Optional] Provide credentials for the destination if the task credentials do not enable recovery of the selected data. To access this setting, click **Show access credentials**.

## When to recover

### Recover    (p. 110)

Select when to start recovery. The task can start immediately after its creation, be scheduled for a specified date and time in the future or simply saved for manual execution.

## Task parameters

### Task name

[Optional] Enter a unique name for the recovery task. A conscious name lets you quickly identify the task among the others.

### Recovery options

[Optional] Customize the recovery operation by configuring the recovery options, such as pre/post recovery commands, recovery priority, error handling or notification options. If you do nothing in this section, the default values (p. 126) will be used.

After any of the settings are changed against the default value, a new line that displays the newly set value appears. The setting status changes from **Default** to **Custom**. Should you modify the setting again, the line will display the new value unless the new value is the default one. When the default value is set, the line disappears. Therefore, in this section you always see only the settings that differ from the default values.

Clicking **Reset to default** resets all the settings to default values.

### Task credentials

[Optional] The task will run on behalf of the user who is creating the task. You can change the task account credentials if necessary. To access this setting, click **Show task credentials**.

## [Optional] Universal Restore for Windows/Linux

*Applies to: system disk or volume recovery. Using Acronis Universal Restore requires a separate license.*

**Universal Restore for Windows/Linux** (p. 112)

Use Acronis Universal Restore when you need to recover and boot up an operating systems on dissimilar hardware.

After you complete all the required steps, click **OK** to create the commit creating of the recovery task.

# 5.1.1   What to recover

## 1. Specifying the archive location

In the **Data path** field, specify the archive location path or click **Browse** and select the required location as described in "Selecting archive location" (p. 100).

In the advanced editions of Acronis Backup & Recovery 11, you can select either to specify the archive location path as described above, or use the centralized data catalog.

## 2. Selecting data

The backed up data can be selected using the **Data view** tab, or the **Archive view** tab. The **Data view** tab displays all the backed up data by versions (the date and time of backup creation) within the selected archive location. The **Archive view** tab displays the backed up data by the archives.

*Note*: *File-level recovery with Agent for ESX(i) or Agent for Hyper-V is not possible.*

### Selecting data using the Data view

Since the **Data view** tab shares the same functionality with the data catalog, selecting data on the **Data view** tab is performed in the same way as in the catalog. For more information about selecting data, see "Data catalog" (p. 101).

### Selecting data using the Archive view

1. Expand the required archive and select one of the successive backups by its creation date and time. Thus, you can revert the disk data to a certain moment in time.

   If the list of archives is too long, you can filter the archives by selecting only the required type of archives to display. To do this, select the required archive type in the **Show** list.

2. For disk or volume backups only: in the **Backup contents**, select the type of data to display from the drop-down box:

   - **Disks** - to recover disks as a whole (with all their volumes).
   - **Volumes** - to recover individual basic and/or dynamic volumes.
   - **Files** - to recover individual files and folders.

3. In the **Backup contents**, select the check boxes for the items you need to recover.

4. Click **OK**.

### Selecting MBR
When recovering a system volume, you will usually select the disk's MBR if:

- The operating system cannot boot.
- The disk is new and does not have MBR.
- You are recovering custom or non-Windows boot loaders (such as LILO and GRUB).
- The disk geometry is different to that stored in the backup.

There are probably other times when you may need to recover the MBR, but the above are the most common.

When recovering the MBR of one disk to another Acronis Backup & Recovery 11 recovers Track 0, which does not affect the target disk's partition table and partition layout. Acronis Backup & Recovery 11 automatically updates Windows loaders after recovery, so there is no need to recover the MBR and Track 0 for Windows systems, unless the MBR is damaged.

## 5.1.1.1 Selecting archive location

| Location | Details |
|---|---|
| Online backup storage | If the archive is stored in Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then, expand the **Online backup storage** group and select the account. |
| | *Exporting and mounting are not supported for backups stored in Acronis Online Backup Storage.* |
| Personal | If the archive is stored in a personal vault, expand the **Personal** group and click the required vault. |
| Centralized | If the archive is stored in a centralized vault, expand the **Centralized** group and click the appropriate vault. |
| Machine name | This is the local machine name. |
| Local folders | If the archive is stored in a local folder on the machine, expand the **<Machine name>** group and select the required folder. |
| CD, DVD, etc. | If the archive is stored on optical media such as CD or DVD, expand the **<Machine name>** group, then select the required drive. First insert the last DVD. Then insert the discs in order starting from the first one when the program prompts. |
| Tape device | If the archive is stored on a locally attached tape device, expand the **Tape drives** group, then click the required device. |
| | Tape devices are available only if you have upgraded from Acronis Backup & Recovery 10. For information about using tapes, see the "Tape devices" section of the product Help. |
| Network folders | If the archive is stored on a network share, expand the **Network folders** group, then select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them. |
| FTP, SFTP | If the archive is stored on an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows: |
| | **ftp://ftp_server:port _number or sftp://sftp_server:port number** |
| | If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP. |
| | After entering access credentials, the folders on the server become available. Click the appropriate folder on the server. |
| | You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials. |
| | *According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.* |

## 5.1.1.2    Data catalog

Data catalog lets you easily find the required version of data and select it for recovery. On a managed machine, the data catalog functionality is available through the **Data view** tab for any vault accessible from this machine. On the management server, the catalog functionality is available through both **Data view** and the centralized **Data catalog**. The centralized data catalog displays in a single place all the data stored in the centralized managed vaults.

### Selecting the backed up data for recovery

1.  Do any of the following:
    -   To access the **Data view** tab, connect the console to a machine or to the management server, navigate to **Vaults** view, and click the required vault.
    -   To access the **Data catalog**, connect the console to the management server and select **Data catalog** in the **Navigation** tree.
2.  In the **Show** field, select the type of data to display:
    -   Select **Machines/disks/volumes** to browse and search for entire disks and volumes in disk-level backups.
    -   Select **Folders/files** to browse and search for files and folders in both file-level and disk-level backups.
3.  In the **Display data backed up for** field, specify the time period for which the backed up data will be displayed.
4.  Do any of the following:
    -   Select the data to recover in the catalog tree, or in the table to the right of the catalog tree.
    -   In the search string, type the information that helps to identify the required data items (this can be a machine name, a file or folder name, or a disk label) and then click **Search**. You can use the asterisks (*) and question marks (?) wildcards.

        As a result, in the **Search** window, you will see the list of backed up data items whose names fully or partially coincide with the entered value. If the list of matches is too long, you can refine the search criteria by specifying the date or time range of backup creation, and the size range of backed up items. When the required data is found, select it, and click **OK** to return to the **Data catalog/Data view**.
5.  Use the **Versions** list to select the point of time to revert the data to. By default, the data will be reverted to latest point of time available for the time period selected in step 3.
6.  Having selected the required data, click **Recover** and configure the parameters of the recovery operation.

### What if the data does not appear in the catalog or data view

The probable reasons of the issue are as follows.

### Wrong time period is set

The required data was not backed up during the time period set by the **Display data backed up for** control.

**Solution:** Try to increase the time period.

### Cataloging is turned off

If the data is displayed partially or is not displayed at all, most likely the backup cataloging option (p. 79) was disabled during backup.

**Solutions:**

- Run the cataloging manually by clicking **Catalog now**. For the **Data catalog**, all backups stored in the managed vaults will be cataloged. For the **Data view**, only the backups stored on the selected vault will be cataloged. The backups that have already been cataloged, will not be cataloged again.

- Since cataloging a large number of backed up data may take a long time, you may prefer to use the **Archive view** of the respective vault. For more information about using the Archive view, see "Browsing the vault contents and data selection" in the Working with vaults (p. 136) section.

**The data is not supported by the catalog**

The following data cannot be displayed in the catalog or data view:

- Data from the encrypted and password-protected archives.
- Data backed up to removable media, such as CD, DVD, BD, Iomega REV.
- Data backed up to Acronis Online Backup Storage.
- Data backed up using Acronis True Image Echo or earlier product versions.
- Data backed up using the simplified backup naming.

**Solution:** To be able to browse such data, use the **Archive view** tab of the respective vault.

**The data is not included in the centralized catalog**

Data from personal vaults (p. 136) is not displayed in the centralized catalog.

**Solution:** To be able to browse such data, connect directly to a machine, select the required personal vault and then select **Data view**.

## 5.1.2    Access credentials for location

Specify the credentials required for access to the location where the backup is stored.

***To specify credentials***

1. Select one of the following:

   - **Use the task credentials**

     The software will access the location using the credentials of the task account specified in the **Task parameters** section.

   - **Use the following credentials**

     The software will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

     Specify:

     - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).
     - **Password**. The password for the account.

2. Click **OK**.

*According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

## 5.1.3 Access credentials for destination

***To specify credentials***

1. Select one of the following:

   - **Use the task credentials**

     The program will access the destination using the credentials of the task account specified in the **Task parameters** section.

   - **Use the following credentials**

     The program will access the destination using the credentials you specify. Use this option if the task account does not have access permissions to the destination.

     Specify:

     - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).
     - **Password**. The password for the account.

2. Click **OK**.

## 5.1.4 Where to recover

Specify the destination the selected data will be recovered to.

## 5.1.4.1 Selecting target disks

Available disk or volume destinations depend on the agents operating on the machine.

**Recover to:**

**Physical machine**

*Available when the Acronis Backup & Recovery 11 Agent for Windows or Agent for Linux is installed.*

The selected disks will be recovered to the physical disks of the machine the console is connected to. On selecting this, you proceed to the regular disk mapping procedure described below.

**New virtual machine**

- *If Acronis Backup & Recovery 11 Agent for Windows is installed.*

  The selected disks will be recovered to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation or Citrix XenServer Open Virtual Appliance (OVA). The virtual machine files will be saved to the destination you specify.

- *If Acronis Backup & Recovery 11 Agent for Hyper-V or Agent for ESX(i) is installed.*

  These agents enable creating a new virtual machine on the virtualization server you specify.

  The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **Virtual Machine Settings** (p. 122) section. Check the settings and make changes if necessary.

  Then you proceed to the regular disk mapping procedure described below.

**Existing virtual machine**

*Available when the Acronis Backup & Recovery 11 Agent for Hyper-V or Agent for ESX(i) is installed.*

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular disk mapping procedure described below.

*Please be aware that the target machine will be powered off automatically before the recovery. If you prefer to power it off manually, modify the VM power management option.*

### Disks/volumes

#### Map automatically

Acronis Backup & Recovery 11 attempts to map the selected disks to the target disks as described in the "How the automatic mapping works" (p. 105) section. If you are unsatisfied with the mapping result, you can re-map disks manually. To do this, you have to unmap the disks in a reverse order; that is, the last mapped disk should be unmapped first. Then, map the disks manually as described below.

### Disk #:

#### Disk # (MODEL) (p. 104)

Select the destination disk for each of the source disks.

#### NT signature (p. 104)

Select the way the recovered disk's signature will be handled. The disk signature is used by Windows and the Linux kernel version 2.6 and later.

## Disk destination

***To specify a destination disk:***

1. Select a disk where you want the selected disk to recover to. The destination disk's space should be at least the same size as the uncompressed image data.

2. Click **OK**.

*All the data stored on the target disk will be replaced by the backed up data, so be careful and watch out for non-backed-up data that you might need.*

## NT signature

The NT signature is a record that is kept in the MBR. It uniquely identifies the disk for the operating system.

When recovering a disk containing a system volume, you can choose what to do with the NT signature of the target disk. Specify any of the following parameters:

- **Select automatically**

  The software will keep the NT signature of the target disk if it is the same as the NT signature stored in the backup. (In other words, if you recover the disk to the same disk that was backed up.) Otherwise, the software will generate a new NT signature for the target disk.

  This is the default selection recommended in most cases. Use the following settings only if you absolutely need to.

- **Create new**

  Acronis Backup & Recovery 11 will generate a new NT signature for the target hard disk.

- **Recover from backup**

  Acronis Backup & Recovery 11 will replace the NT signature of the target hard disk with one from the disk backup.

*Note: You should be absolutely sure that none of the existing disks on this machine has the same NT signature. Otherwise, the operating system runs from the first disk at the startup; discovers the same signature on the second one, automatically generates a new unique NT signature and assigns it to the second disk. As a result, all the volumes on the second disk will lose their letters, all paths will be invalid on the disk, and programs won't find their files. The operating system on that disk will be unbootable.*

Recovering the disk signature may be desirable due to the following reasons:

- Acronis Backup & Recovery 11 schedules tasks using the signature of the source hard disk. If you recover the same disk signature, you don't need to re-create or edit the tasks created previously.
- Some installed applications use disk signature for licensing and other purposes.
- This enables you to keep all the Windows Restore Points on the recovered disk.
- To recover VSS snapshots used by Windows Vista's "Previous Versions" feature.

- **Keep existing**

  The program will leave the NT signature of the target hard disk untouched.

## How the automatic mapping works

Acronis Backup & Recovery 11 automatically maps the disks or volumes to the target disks only if the system bootability can be preserved. Otherwise, the automatic mapping is canceled and you have to map the disks or volumes manually.

Also, you have to map the volumes manually if they are Linux logical volumes, or Linux software RAID (MD devices). For more information on recovering logical volumes and MD devices, see Recovering MD devices and logical volumes.

The automatic mapping is performed as follows.

1. If the disk or volume is recovered to its original location, the mapping process reproduces the original disk/volume layout.

   *The original location for a disk or volume means exactly the same disk or volume that has been backed up. A volume will not be considered original if its size, location or other physical parameters have been changed after backup. Changing the volume letter or label does not prevent the software from recognizing the volume.*

2. If the disk or volume is recovered to a different location:

   - **When recovering disks**: The software checks the target disks for size and volumes. A target disk must contain no volumes and its size must be large enough to place the disk being recovered. Not initialized target disks will be initialized automatically.

     If the required disks cannot be found, you have to map the disks manually.

   - **When recovering volumes**: The software checks the target disks for unallocated space.

     If there is enough unallocated space, the volumes will be recovered "as is".

     If unallocated space on the target disks is less than the size of the volumes being recovered, the volumes will be proportionally shrunk (by decreasing their free space) in order to fit the unallocated space. If the shrunk volumes still cannot fit the unallocated space, you have to map the volumes manually.

## 5.1.4.2    Selecting target volumes

Available volume destinations depend on the agents operating on the machine.

**Recover to:**

**Physical machine**

*Available when the Acronis Backup & Recovery 11 Agent for Windows or Agent for Linux is installed.*

The selected volumes will be recovered to the physical disks of the machine the console is connected to. On selecting this, you proceed to the regular volume mapping procedure described below.

**New virtual machine**

▪ *If Acronis Backup & Recovery 11 Agent for Windows is installed.*

The selected volumes will be recovered to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation or Citrix XenServer Open Virtual Appliance (OVA). The virtual machine files will be saved to the destination you specify.

▪ *If Acronis Backup & Recovery 11 Agent for Hyper-V or Agent for ESX(i) is installed.*

These agents enable creating a new virtual machine on the virtualization server you specify.

The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **Virtual Machine Settings** (p. 122) section. Check the settings and make changes if necessary.

Then you proceed to the regular volume mapping procedure described below.

**Existing virtual machine**

*Available when the Acronis Backup & Recovery 11 Agent for Hyper-V or Agent for ESX(i) is installed.*

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular volume mapping procedure described below.

*Please be aware that the target machine will be powered off automatically before recovery. If you prefer to power it off manually, modify the VM power management option.*

### Disks/volumes

**Map automatically**

Acronis Backup & Recovery 11 attempts to map the selected volumes to the target disks as described in the "How the automatic mapping works" (p. 105) section. If you are unsatisfied with the mapping result, you can re-map volumes manually. To do this, you have to unmap the volumes in a reverse order; that is, the last mapped volume should be unmapped first. Then, map the volumes manually as described below.

**Recover [Disk #] MBR to: [If the Master Boot Record is selected for recovery]**

**Disk #** (p. 107)

Choose the disk to recover the Master Boot Record to.

**NT signature:** (p. 104)

Select the way the disk's signature contained in the MBR will be handled. The disk signature is used by Windows and the Linux kernel version 2.6 and later.

**Recover [Volume] [Letter] to:**

**Disk # /Volume**

Sequentially map each of the source volumes to a volume or an unallocated space on the destination disk.

**Size**: (p. 107)

[Optional] Change the recovered volume size, location and other properties.

## MBR destination

***To specify a destination disk:***

1. Select the disk to recover the MBR to.
2. Click **OK**.

## Volume destination

***To specify a target volume or unallocated space***

1. Select a volume or unallocated space where you want the selected volume to be recovered to. The destination volume/unallocated space should be at least the same size as the uncompressed image data.
2. Click **OK**.

*All the data stored on the target volume will be replaced by the backed up data, so be careful and watch out for non-backed-up data that you might need.*

### When using bootable media

Disk letters seen under Windows-style bootable media might differ from the way Windows identifies drives. For example, the D: drive in the rescue utility might correspond to the E: drive in Windows.

*Be careful! To be on the safe side, it is advisable to assign unique names to the volumes.*

The Linux-style bootable media shows local disks and volumes as unmounted (sda1, sda2...).

## Changing volume properties

## Size and location

When recovering a volume to a basic MBR disk, you can resize and relocate the volume by dragging it or its borders with a mouse or by entering corresponding values in the appropriate fields. Using this feature, you can redistribute the disk space between the volumes being recovered. In this case, you will have to recover the volume to be reduced first.

*Note: Volumes backed up using the sector-by-sector option cannot be resized.*

*Tip: A volume cannot be resized when being recovered from a backup split into multiple removable media. To be able to resize the volume, copy all parts of the backup to a single location on a hard disk.*

## Type

A basic MBR disk can contain up to four primary volumes or up to three primary volumes and multiple logical drives. By default, the program selects the original volume's type. You can change this setting, if required.

- **Primary**. Information about primary volumes is contained in the MBR partition table. Most operating systems can boot only from the primary volume of the first hard disk, but the number of primary volumes is limited.

  If you are going to recover a system volume to a basic MBR disk, select the Active check box. Active volume is used for loading an operating system. Choosing active for a volume without an

installed operating system could prevent the machine from booting. You cannot set a logical drive or dynamic volume active.

▪ **Logical**. Information about logical volumes is located not in the MBR, but in the extended partition table. The number of logical volumes on a disk is unlimited. A logical volume cannot be set as active. If you recover a system volume to another hard disk with its own volumes and operating system, you will most likely need only the data. In this case, you can recover the volume as logical to access the data only.

## File system

By default, the recovered volume will have the same file system as the original volume has. You can change the volume's file system during recovery, if required.

Acronis Backup & Recovery 11 can make the following file system conversions: FAT 16 -> FAT 32 and Ext2 -> Ext3. For volumes with other native file systems, this option is not available.

Assume you are going to recover a volume from an old, low-capacity FAT16 disk to a newer disk. FAT16 would not be effective and might even be impossible to set on the high-capacity hard disk. That's because FAT16 supports volumes up to 4 GB, so you will not be able to recover a 4 GB FAT16 volume to a volume that exceeds that limit, without changing the file system. It would make sense here to change the file system from FAT16 to FAT32.

Older operating systems (MS-DOS, Windows 95 and Windows NT 3.x, 4.x) do not support FAT32 and will not be operable after you recover a volume and change its file system. These can be normally recovered on a FAT16 volume only.

## Volume (partition) alignment

Acronis Backup & Recovery 11 automatically eliminates volume misalignment – a situation, when volume clusters are not aligned with disk sectors. The misalignment occurs when recovering volumes created with the Cylinder/Head/Sector (CHS) addressing scheme to ahard disk drive (HDD) or solid-state drive (SSD) drive that has a 4-KB sector size. The CHS addressing scheme is used, for example, in all Windows operating systems earlier than Windows Vista.

If volumes are misaligned, the cluster overlaps more physical sectors than it would have occupied if aligned. As a result, more physical sectors need to be erased and rewritten each time the data changes. The redundant read/write operations noticeably slow down the disk speed and overall system performance. SSD drive misalignment decreases not only system performance, but drive lifetime. Since SSD memory cells are designed for a certain amount of read/write operations, redundant read/write operations lead to early degradation of the SSD drive.

When recovering dynamic volumes and logical volumes created in Linux with Logical Volume Manager (LVM), the appropriate alignment is set up automatically.

When recovering basic MBR and GPT volumes, you can select the alignment method manually if the automatic alignment does not satisfy you for some reason. The following options are available:

▪ **Select automatically** - (Default) recommended. The software will automatically set the appropriate alignment based on the source and target disk/volume properties.

Use the following options only if you absolutely need to.

▪ **CHS (63 sectors)** - select this option if the recovered volume will be used under Microsoft Windows XP and Windows Server 2003 (or earlier) on disks having 512 bytes per physical sector.

- **VMWare VMFS (64 KB)** - select this option when recovering the volume as a VMware Virtual Machine File System partition.

- **Vista alignment (1 MB)** - select this option if the recovered volume will be used under Windows operating systems starting from Windows Vista, or when recovering volumes to an HDD or SSD drive that has a 4-KB sector size.

- **Custom** - Specify the volume alignment manually. It is recommended that the value be a multiple of the physical sector size.

## Logical drive letter (for Windows only)

By default, the first unused letter will be assigned to the volume. To assign other letter, select the desired letter from a drop-down list.

If you select the empty value, no letter will be assigned to the recovered volume, hiding it from the OS. You should not assign letters to volumes that are inaccessible to Windows, such as to those other than FAT and NTFS.

## 5.1.4.3 Selecting target location for files and folders

**Where to recover**

**Destination**

Select a location to recover the backed up files to:

- **Original location**

  Files and folders will be recovered to the same path(s) as they are in the backup. For example, if you have backed up all files and folders in *C:\Documents\Finance\Reports\*, the files will be recovered to the same path. If the folder does not exist, it will be created automatically.

- **New location**

  Files will be recovered to the location that you specify in the tree. The files and folders will be recovered without recreating a full path, unless you clear the **Recover without full path** check box.

**Recovery agent**

Select Acronis Agent that will perform file recovery. The agent selection is available only when the software cannot detect the agent on the machine the files will be recovered to.

**Overwriting**

Choose what to do if the program finds in the target folder a file with the same name as in the archive:

- **Overwrite existing file** - this will give the file in the backup priority over the file on the hard disk.

- **Overwrite existing file if it is older** - this will give priority to the most recent file modification, whether it be in the backup or on the disk.

- **Do not overwrite existing file** - this will give the file on the hard disk priority over the file in the backup.

If you allow files to be overwritten, you still have an option to prevent overwriting of specific files by excluding them from the recovery operation.

**Recovery exclusions** (p. 110)

Specify files and folders you do not wish to be recovered.

## Recovery exclusions

Set up exclusions for the specific files you do not wish to recover.

Use the **Add**, **Edit**, **Remove** and **Remove All** buttons to create the list of file masks. Files whose names match any of the masks will be skipped during recovery.

You can use one or more wildcard characters * and ? in a file mask:

- The asterisk (*) substitutes for zero or more characters in a file name; for example, the file mask Doc*.txt yields files such as Doc.txt and Document.txt
- The question mark (?) substitutes for exactly one character in a file name; for example, the file mask Doc?.txt yields files such as Doc1.txt and Docs.txt — but not the files Doc.txt or Doc11.txt

**Exclusion examples**

| Criterion | Example | Description |
|---|---|---|
| **Windows and Linux** | | |
| By name | F.log | Excludes all files named "F.log" |
| | F | Excludes all folders named "F" |
| By mask (*) | *.log | Excludes all files with the .log extension |
| | F* | Excludes all files and folders with names starting with "F" (such as folders F, F1 and files F.log, F1.log) |
| By mask (?) | F???.log | Excludes all .log files with names consisting of four symbols and starting with "F" |
| **Windows** | | |
| By file path | Finance\F.log | Excludes files named "F.log" from all folders with the name "Finance" |
| By folder path | Finance\F\ or Finance\F | Excludes folders named "F" from all folders with the name "Finance" |
| **Linux** | | |
| By file path | /home/user/Finance/F.log | Excludes the file named "F.log" located in the folder /home/user/Finance |

The above settings are not effective for the files or folders that were explicitly selected for recovery. For example, assume that you selected the folder MyFolder and the file MyFile.tmp outside that folder, and selected to skip all .tmp files. In this case, all .tmp files in the folder MyFolder will be skipped during the recovery process, but the file MyFile.tmp will not be skipped.

## 5.1.5    When to recover

Select when to start the recovery task:

- **Now** - the recovery task will be started immediately after you click **OK** on the **Recover data** page.

- **Later** - the recovery task will be started manually afterwards. If you need to schedule the task, clear the **Task will be started manually** check box, and specify the required date and time.

## 5.1.6    Task credentials

Provide credentials for the account under which the task will run.

***To specify credentials***

1. Select one of the following:

   - **Run under the current user**

     The task will run under the credentials with which the user who starts the tasks is logged on. If the task has to run on schedule, you will be asked for the current user's password on completing the task creation.

   - **Use the following credentials**

     The task will always run under the credentials you specify, whether started manually or executed on schedule.

     Specify:

     - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).

     - **Password**. The password for the account.

2. Click **OK**.

To learn more about using credentials in Acronis Backup & Recovery 11, see the Owners and credentials (p. 21) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine section.

## 5.2    Acronis Universal Restore

Acronis Universal Restore is the Acronis proprietary technology that helps recover and boot up an operating system on dissimilar hardware or a virtual machine. Universal Restore handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

Universal Restore is extremely useful in the following scenarios:

1. Instant recovery of a failed system on different hardware.
2. Hardware-independent cloning and deployment of operating systems.
3. Physical-to-physical, physical-to-virtual and virtual-to-physical machine migration.

## 5.2.1    Getting Universal Restore

Universal Restore is always available when recovering a system from the online storage.

Universal Restore comes free with the Acronis Backup & Recovery 11 Advanced Server SBS Edition and Virtual Edition.

Universal Restore for the other product editions is purchased separately. It has its own license.

To enable Universal Restore on a managed machine, do any of the following:

- Install Universal Restore from the product installation package (in addition to Agent for Windows, Agent for Linux or Bootable Media Builder).

- If the agent is already installed, you can connect the management console to the machine, click **Help > Change license** and specify the license key or the license server from where to take the Universal Restore license.

You need to re-create bootable media to make the newly installed add-on operational in the bootable environment.

## 5.2.2 Using Universal Restore

**During recovery**

Universal Restore is available when configuring a disk or volume recovery, if a Windows or Linux operating system is present in your selection of disks or volumes. If there are more than one operating systems in your selection, you can apply Universal Restore to all Windows systems, all Linux systems or to both Windows and Linux systems.

If the software cannot detect whether an operating system is present in the backup, it suggests using Universal Restore on the off-chance of the system presence. These cases are as follows:

- the backup is split into several files
- the backup is located in a deduplicating vault, in Acronis Online Backup Storage, on an FTP/SFTP server, tape, CD or DVD.

Sometimes Universal Restore is applied in the background because the software knows what drivers or modules are required for the supported virtual machines. These cases are as follows:

- recovering a system to a new virtual machine
- recovering a system to any virtual machine by means of Agent for ESX(i) or Agent for Hyper-V.

Universal Restore is not available when:

- the backup is located in Acronis Secure Zone
- you have chosen to use Acronis Active Restore (p. 227)

This is because these features are primarily meant for instant data recovery on the same machine.

**Without recovery**

Under bootable media, you can also use Universal Restore without recovery by clicking **Apply Universal Restore** in the media welcome screen. Universal Restore will be applied to the operating system that already exists on the machine. If there are multiple operating systems, you are prompted to choose the one to apply Universal Restore to.

## 5.2.2.1 Universal Restore in Windows

**Preparation**

**Prepare drivers**

Before applying Universal Restore to a Windows operating system, make sure that you have the drivers for the new HDD controller and the chipset. These drivers are critical to start the operating system. Use the CD or DVD supplied by the hardware vendor or download the drivers from the vendor's Web site. The driver files should have the *.inf, *.sys or *.oem extensions. If you download the drivers in the *.exe, *.cab or *.zip format, extract them using a third-party application.

The best practice is to store drivers for all the hardware used in your organization in a single repository sorted by device type or by the hardware configurations. You can keep a copy of the repository on a DVD or a flash drive; pick some drivers and add them to the bootable media; create the custom bootable media with the necessary drivers (and the necessary network configuration) for each of your servers. Or you can simply specify the path to the repository every time Universal Restore is used.

### Check access to the drivers in bootable environment

Make sure you have access to the device with drivers when working under bootable media. Even if you configure system disk recovery in a Windows environment, the machine will reboot and recovery will proceed in the Linux-based environment. Use WinPE-based media if the device is available in Windows but Linux-based media does not detect it.

### What if you do not have drivers

Windows 7 includes more drivers than the older Windows operating systems. There is a great chance that Universal Restore finds all necessary drivers in the Windows 7 driver folder. So, you may not necessarily have to specify the external path to the drivers. Nevertheless, performing Universal Restore is critical so the system uses the correct drivers.

*The Windows default driver storage folder is determined in the registry value **DevicePath**, which can be found in the registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion**. This storage folder is usually WINDOWS/inf.*

## Universal Restore settings

### Automatic driver search

Specify where the program will search for the Hardware Abstraction Layer (HAL), HDD controller driver and network adapter driver(s):

- If the drivers are on a vendor's disc or other removable media, turn on the **Search removable media**.
- If the drivers are located in a networked folder or on the bootable media, specify the path to the folder by clicking **Add folder**.

During recovery, Universal Restore will perform the recursive search in all the sub-folders of the specified folder, find the most suitable HAL and HDD controller drivers of all those available, and install them into the recovered system. Universal Restore also searches for the network adapter driver; the path to the found driver is then transmitted by Universal Restore to the operating system. If the hardware has multiple network interface cards, Universal Restore will try to configure all the cards' drivers.

### Mass storage drivers to install anyway

To access this setting, expand **Show mass storage drivers to install anyway**.

You need this setting if:

- The target hardware has a specific mass storage controller such as RAID (especially NVIDIA RAID) or a fibre channel adapter.
- You recover a system to a virtual machine that uses a SCSI hard drive controller and is booted into bootable media. Use SCSI drivers bundled with your virtualization software or download the latest drivers versions from the software manufacturer Web site.
- If the automatic drivers search does not help to boot the system.

Specify the appropriate drivers by clicking **Add driver**. The drivers defined here will be installed, with appropriate warnings, even if the program finds a better driver.

**The recovery process**

If Universal Restore cannot find a compatible driver in the specified locations, it will display a prompt with the problem device. Do any of the following:

- Add the driver to any of the previously specified locations and click **Retry**.
- If you do not remember the location, continue the recovery. If the result is not satisfactory, start Universal Restore without recovery by clicking **Apply Universal Restore** in the media welcome screen. When configuring the operation, specify the necessary driver.

Once Windows boots, it will initialize the standard procedure for installing new hardware. The network adapter driver will be installed silently if the driver has the Microsoft Windows signature. Otherwise, Windows will ask for confirmation on whether to install the unsigned driver.

After that, you will be able to configure the network connection and specify drivers for the video adapter, USB and other devices.

## 5.2.2.2 Applying Universal Restore to multiple operating systems

During recovery, you can use Universal Restore for operating systems of a certain type: all Windows systems, all Linux systems, or both.

If your selection of volumes to recover contains multiple Windows systems, you can specify all drivers for them in a single list. Each driver will be installed in the operating system for which it is intended.

# 5.3 Recovering BIOS-based systems to UEFI-based and vice versa

Acronis Backup & Recovery 11 supports transferring Windows operating systems between BIOS-based hardware and hardware that supports Unified Extensible Firmware Interface (UEFI).

*Important*: *When migrating systems to UEFI-based hardware, use Linux-based bootable media, since PE-based bootable media and Acronis PXE Server do not support UEFI.*

## 5.3.1 Recovering volumes

Let's assume you backed up the system and boot volumes (or the entire machine) and want to recover these volumes to a different platform. The ability of the recovered system to boot up depends on the following factors:

- **Source operating system**: convertible or non-convertible OS. The following Microsoft Windows operating systems are convertible, that is, allow changing the boot mode from BIOS to UEFI and back:
  - Windows Vista SP1 x64 and later.
  - Windows Server 2008 x64 SP1 and later.
  - Windows 7 x64.
  - Windows Server 2008 R2 x64.

  All other operating systems are non-convertible.

- **Source and target disk partition style**: MBR or GPT. System and boot volumes of BIOS platforms use MBR disks. System and boot volumes of UEFI platforms use GPT disks.

  When selecting not initialized target disk for recovery, this disk will be automatically initialized either to GPT or to MBR depending on the original disk partitioning style, the current boot mode (UEFI or BIOS) and the type of operating systems (convertible or non-convertible) that are located on this volume.

  If the initialization may result in bootability loss, the software takes the partitioning style from the source volume ignoring the target disk size. In such cases the software can select the MBR partitioning style for disks whose size is more than 2 TB; however, a user cannot use the space beyond 2 TB.

  If required, you can initialize the target disk manually by using the Disk management (p. 172) functionality.

The following table summarizes whether it is possible to retain the system bootability when recovering boot and system volumes of a BIOS-based system to UEFI-based and back.

- The plus (+) sign on a green background means that the system will be bootable. No user action is required.
- The plus (+) sign on a yellow background means you need to perform additional steps to make the system bootable.
- The minus (-) sign on a red background means the system will not be able to boot due to BIOS and UEFI platform limitations.

| Original system | Target hardware | | | |
|---|---|---|---|---|
| | BIOS<br><br>Disk: MBR | BIOS<br><br>Disk: GPT | UEFI<br><br>Disk: MBR | UEFI<br><br>Disk: GPT |
| BIOS<br><br>OS: convertible | + | -<br><br>**Solution**<br><br>Recover the operating system to an MBR disk. | +<br><br>**Additional steps**<br><br>1. Before recovery, turn off the UEFI mode in BIOS<br>2. Perform the recovery under the bootable media.<br><br>or<br><br>After recovery, turn off the UEFI mode in BIOS. Otherwise, the operating system cannot be booted. | +<br><br>The convertible OS will be automatically converted to support UEFI booting. |
| BIOS<br><br>OS: non-convertible | | | | -<br><br>**Solution**<br><br>Recover the operating system to an MBR disk. |
| UEFI<br><br>OS: convertible | +<br><br>The convertible OS will be automatically converted to support BIOS booting. | +<br><br>**Additional steps**<br><br>1. Before recovery, turn on the UEFI mode in BIOS. | - | + |

| Original system | Target hardware | | | |
|---|---|---|---|---|
| | BIOS Disk: MBR | BIOS Disk: GPT | UEFI Disk: MBR | UEFI Disk: GPT |
| UEFI OS: non-convertible | - **Solution** Recover the operating system to a GPT disk. | 2. Perform the recovery under the bootable media. or After recovery, turn on the UEFI mode in BIOS. Otherwise, the operating system cannot be booted. | **Solution** Recover the operating system to a GPT disk. | |

## 5.3.2 Recovering disks

Let's assume you backed up a whole disk (with all its volumes) and want to recover this disk to a different target platform.

The ability of the recovered system to boot up in different modes depends on the operating systems installed on the source disk. Operating systems can be **convertible** i.e. allow changing the boot mode from BIOS to UEFI and back, or be **non-convertible**. For the list of convertible operating systems, see Recovering volumes (p. 114).

- When a source disk contains one or more operating systems and *all* of them are convertible, the boot mode can be automatically changed. Depending on the current boot mode, the target disk may be initialized either to GPT or to MBR partitioning style.

- If *at least one* operating system on a source disk is non-convertible (or the source disk contains any boot volumes of the non-convertible OSes), the boot mode cannot be changed automatically and the software will initialize the target disk as the source one. To boot up the target machine, you have to turn on/off the UEFI mode in BIOS manually. Otherwise, the system will not boot after recovery.

The following table summarizes all cases of recovering disks of a BIOS-based system to UEFI-based and vice versa.

- The plus (+) sign on a green background means that the source disk will be recovered without any limitations.

- The plus (+) sign on a yellow background means that the source disk will be recovered with some limitations. To overcome such limitations, you might need to perform some additional steps.

- The plus/minus (+/-) signs on a red background means the disk will be recovered, but some issues might occur.

| Original system | Target hardware | | | |
|---|---|---|---|---|
| | Platform: BIOS Target disk <2 TB | Platform: BIOS Target disk >2 TB | Platform: UEFI Target disk <2 TB | Platform: UEFI Target disk >2 TB |

| Original system | Target hardware | | | |
|---|---|---|---|---|
| | Platform: BIOS<br><br>Target disk <2 TB | Platform: BIOS<br><br>Target disk >2 TB | Platform: UEFI<br><br>Target disk <2 TB | Platform: UEFI<br><br>Target disk >2 TB |
| BIOS<br><br>OS: convertible | +<br><br>The source disk will be recovered without any modification. | +<br><br>The target disk will be initialized as MBR.<br><br>**Limitation**: only 2 TB of the disk space will be available for use.<br><br>To overcome the limitation:<br><br>1. Turn on the UEFI mode in BIOS<br><br>2. Boot from a bootable media, and perform the recovery. As a result, the recovered OS will be converted to use UEFI for booting. | +<br><br>The target disk will be initialized as GPT.<br><br>The recovered disk OS will be automatically converted to support UEFI booting.<br><br>**Note**: If you want to recover the source disk "as is":<br><br>1. Turn off the UEFI mode in BIOS*<br><br>2. Boot from a bootable media, and perform the recovery. | +<br><br>The target disk will be initialized as GPT.<br><br>The recovered disk OS will be automatically converted to support UEFI booting. |
| BIOS<br><br>OS: non-convertible | +<br><br>The source disk will be recovered without any modification. | +<br><br>The target disk will be initialized as a source one (MBR).<br><br>**Limitation**: only 2 TB of the disk space will be available for use. | +/-<br><br>The target disk will be initialized as the source one (MBR).<br><br>**Additional steps**<br><br>Turn off the UEFI mode in BIOS after recovery*<br><br>**Possible issue**<br><br>If target machine does not support BIOS, the system will not boot after recovery. | +/-<br><br>The target disk will be initialized as the source one (MBR).<br><br>**Limitation**: only 2 TB of the disk space will be available for use.<br><br>**Additional steps**<br><br>Turn off the UEFI mode in BIOS after recovery*<br><br>**Possible issue**<br><br>If target machine does not support BIOS, the system will not boot after recovery. |

| Original system | Target hardware | | | |
|---|---|---|---|---|
| | Platform: BIOS Target disk <2 TB | Platform: BIOS Target disk >2 TB | Platform: UEFI Target disk <2 TB | Platform: UEFI Target disk >2 TB |
| UEFI OS: convertible | + The target disk will be initialized as MBR. The recovered disk OS will be automatically converted to support BIOS booting. **Note**: If you want to recover the source disk "as is": 1. Turn on the UEFI mode in BIOS. 2. Boot from a bootable media, and perform the recovery. | + The target disk will be initialized as MBR. The recovered disk OS will be automatically converted to support BIOS booting. **Limitation**: only 2 TB of the disk space will be available for use. | + The source disk will be recovered without any modification. | + The source disk will be recovered without any modification. |
| UEFI OS: non-convertible | +/- The target disk will be initialized as the source one (GPT). **Additional steps** 1. Turn on the UEFI mode in BIOS. 2. Boot from a bootable media, and perform the recovery. **Possible issue** If target machine does not support UEFI, the system will not boot after recovery. | +/- The target disk will be initialized as the source one (GPT). **Additional steps** 1. Turn on the UEFI mode in BIOS. 2. Boot from a bootable media, and perform the recovery. **Possible issue** If target machine does not support UEFI, the system will not boot after recovery. | + The source disk will be recovered without any modification. | + The source disk will be recovered without any modification. |

* In most of the current motherboards there is a BIOS compatibility mode. So, if the system does not find any UEFI boot loader it will try to boot system in BIOS mode.

# 5.4   Acronis Active Restore

Active Restore is the Acronis proprietary technology that brings a system online immediately after the system recovery is started.

Customers familiar with Acronis Recovery for Microsoft Exchange can note that this product uses Active Restore to achieve immediate availability of an Exchange information store after starting the recovery. While based on the same technology, recovery of the Information Store proceeds in quite a different way than the operating system recovery described in this section.

### Supported operating systems

Acronis Active Restore is available when recovering Windows starting from Windows 2000.

### Limitation

The only supported archive location is a local drive, or more precisely, any device available through the machine's BIOS. This may be Acronis Secure Zone, a USB hard drive, a flash drive or any internal hard drive.

### How it works

When configuring a recovery operation, you select disks or volumes to recover from a backup. Acronis Backup & Recovery 11 scans the selected disks or volumes in the backup. If this scan finds a supported operating system, the Acronis Active Restore option becomes available.

If you do not enable the option, the system recovery will proceed in the usual way and the machine will become operational after the recovery is completed.

If you enable the option, the sequence of actions will be set as follows.

Once the system recovery is started, the operating system boots from the backup. The machine becomes operational and ready to provide necessary services. The data required to serve incoming requests is recovered with the highest priority; everything else is recovered in the background.

Because serving requests is performed simultaneously with recovery, the system operation can slow down even if recovery priority in the recovery options is set to **Low**. This way, the system downtime is reduced to a minimum at the cost of a temporary performance downgrade.

### Usage scenarios

1. The system uptime is one of the efficiency criteria.

   Examples: Client-oriented online services, Web-retailers, polling stations.

2. The system/storage space ratio is heavily biased toward storage.

   Some machines are being used as storage facilities, where the operating system claims a small space segment and all other disk space is committed to storage, such as movies, sounds or other multimedia files. Some of these storage volumes can be extremely large as compared to the system and so practically all the recovery time will be dedicated to recovering the files, which might be used much later on, if in any near future at all.

   If you opt for Acronis Active Restore, the system will be operational in a short time. Users will be able to open the necessary files from the storage and use them while the rest of the files, which are not immediately necessary, are being recovered in the background.

   Examples: movie collection storage, music collection storage, multimedia storage.

### How to use

1. Back up the system disk or volume to a location accessible through the system's BIOS. This may be Acronis Secure Zone, a USB hard drive, a flash drive or any internal hard drive.

*If your operating system and its loader reside on different volumes, always include both volumes in the backup. The volumes must also be recovered together; otherwise there is a high risk that the operating system will not start.*

2.  Create bootable media.

3.  If a system failure occurs, boot the machine using the bootable media. Start the console and connect to the bootable agent.

4.  Configure the system recovery: select the system disk or volume and select the **Use Acronis Active Restore** check box.

    *Acronis Active Restore will choose for the boot-up and subsequent recovery the first operating system found during the backup scan. Do not try to recover more than one operating system using Active Restore if you want the result to be predictable. When recovering a multi-boot system, choose only one system volume and boot volume at a time.*

5.  Once the system recovery is started, the operating system boots from the backup. The Acronis Active Restore icon appears in the system tray. The machine becomes operational and ready to provide necessary services. The immediate user sees the drive tree and icons and can open files or launch applications even though they were not yet recovered.

    The Acronis Active Restore drivers intercept system queries and set the immediate priority for recovery of the files that are necessary to serve the incoming requests. While this on-the-fly recovery proceeds, the continuing recovery process is transferred to the background.

    *Please do not shut down or reboot the machine until the recovery is completed. If you switch off the machine, all the changes made to the system since the last boot up would be lost. The system will not be recovered, not even partially. The only possible solution in this case will be to restart the recovery process from a bootable media.*

6.  The background recovery continues until all the selected volumes are recovered, the log entry is made and the Acronis Active Restore icon disappears from the system tray.

# 5.5   How to convert a disk backup to a virtual machine

Rather than converting a TIB file to a virtual disk file, which requires additional operations to bring the virtual disk into use, Acronis Backup & Recovery 11 performs the conversion by recovery of a disk backup to a fully configured and operational new virtual machine. You have the ability to adapt the virtual machine configuration to your needs when configuring the recovery operation.

With **Acronis Backup & Recovery 11 Agent for Windows**, you can recover a disk (volume) backup to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat Kernel-based Virtual Machine (KVM).

Files of the new virtual machine will be placed in the folder you select. You can start the machine using the respective virtualization software or prepare the machine files for further usage. The Citrix XenServer Open Virtual Appliance (OVA) can be imported to a XenServer using Citrix XenCenter. You can move the files of the Kernel-based Virtual Machine to a machine running Linux; from there, you can run this virtual machine by using the Virtual Machine Manager program. The VMware Workstation machine can be converted to the open virtualization format (OVF) using the VMware OVF tool.

With **Acronis Backup & Recovery 11 Agent for Hyper-V** or **Agent for ESX(i)**, you can recover a disk (volume) backup to a new virtual machine on the respective virtualization server.

*Tip. Microsoft Virtual PC does not support disks that are larger than 127 GB. Acronis enables you to create a Virtual PC machine with larger disks so that you can attach the disks to a Microsoft Hyper-V virtual machine.*

### To convert a disk backup to a virtual machine:

1. Connect the console to the management server or to a machine where Agent for Windows, Agent for Hyper-V or Agent for ESX(i) is installed.

2. Click **Recover** to open the **Recover data** (p. 97) page.

3. Click **Select data** (p. 99). Use the **Data view** tab or the **Archive view** tab to select the disks or volumes to convert.

4. In **Recover to**, select **New virtual machine**.

5. [Only if the console is connected to the management server] Click **Browse**. Select the machine with agent that will perform conversion.

6. In **VM type**, select the resulting virtual machine type and location.

7. [Optional] In **Virtual machine settings** (p. 122) you can change the path to the new virtual machine, rename the machine, change the disk provisioning mode, the allocated memory, and other settings.

   *The same type of machines with the same name cannot be created in the same folder. Change either the VM name or the path if you get an error message caused by identical names.*

8. The destination disk for each of the source disks or source volumes and MBRs will be selected automatically. If required, you can change the destination disks.

   *On a Microsoft Virtual PC, be sure to recover the disk or volume where the operating system's loader resides to the Hard disk 1. Otherwise, the operating system will not boot. This cannot be fixed by changing the boot device order in BIOS, because a Virtual PC ignores these settings.*

9. In **When to recover**, specify when to start the recovery task.

10. [Optional] In **Task parameters,** review **Recovery options** and change the settings from the default ones, if need be. You can specify in **Recovery options > VM power management** whether to start the new virtual machine automatically, after the recovery is completed. This option is available only when the new machine is created on a virtualization server.

11. Click **OK**. If the recovery task is scheduled for the future, specify the credentials under which the task will run.

You will be taken to the **Backup plans and tasks** view, where you can examine the state and progress of the recovery task.

### Post-conversion operations

The resulting machine always has SCSI disk interface and basic MBR volumes. If the machine uses a custom boot loader, you might need to configure the loader to point to the new devices and reactivate it. Configuring GRUB is described in "How to reactivate GRUB and change its configuration" (p. 124).

*Tip. If you want to preserve logical (LVM) volumes on a Linux machine, consider the alternative method of conversion. Create a new virtual machine, boot it using bootable media and perform recovery just like you do on a physical machine. The LVM structure can be automatically recreated during recovery.*

# 5.5.1 Virtual machine type / virtualization server selection

The new virtual machine can be created either on a virtualization server (this requires Acronis Backup & Recovery 11 Agent for Hyper-V or Agent for ESX(i) to be installed) or in any accessible local or networked folder.

***To select the virtualization server the new virtual machine will be created on***

1. Choose the **Place on the virtualization server that I select** option.
2. In the left part of the window, select the virtualization server. Use the right part of the window to review details on the selected server.
3. Click **OK** to return to the **Data recovery** page.

***To select the type of virtual machine***

1. Choose the **Save as files of the VM type that I select to the folder that I specify** option.
2. In the left part of the window, select the virtual machine type. Use the right part of the window to review details on the selected virtual machine type.
3. Click **OK** to return to the **Data recovery** page.

# 5.5.2 Virtual machine settings

The following virtual machine settings can be configured.

## Name

**Default name**: New virtual machine

Enter the name for the new virtual machine.

## Storage

**Initial setting:** the default storage of the virtualization server if the new machine is created on the virtualization server. Otherwise the current user's documents folder.

This is the place where the new virtual machine will be created. Whether you can change the storage on the virtualization server or not, depends on the virtualization product brand and settings. VMware ESX may have multiple storages. A Microsoft Hyper-V server enables creating a new virtual machine in any local folder.

## Memory

**Initial setting:** if not contained in the backup, it is the default setting of the virtualization server.

This is the amount of memory allocated to the new virtual machine. The memory adjustment range depends on the host hardware, the host operating system and the virtualization product settings. For example, virtual machines may be allowed to use no more than 30% of memory.

## Disks

**Initial setting:** the number and size of the source machine's disks.

The number of disks is generally equal to that of the source machine. It might be different if the program has to add more disks to accommodate the source machine volumes because of limitations set by the virtualization product. You can add virtual disks to the machine configuration or, in some cases, delete the proposed disks.

When adding a new virtual disk, along with interface and capacity, you can specify its format.

- **Thin format.** The disk occupies as much storage space as the data it stores. This saves the storage space. To enable thin format, select the **Thin provisioning** (for ESX), or **Dynamically expanding disk** (for Hyper-V) check box.

- **Thick format.** The disk occupies all the provisioned storage space. This improves the virtual machine performance. To use thick format, clear the **Thin provisioning** (for ESX), or **Dynamically expanding disk** (for Hyper-V) check box.

The default setting is thick format if a physical machine was backed up. When recovering from a virtual machine backup, the software tries to reproduce the format of the original machine's disks. If this is not possible, thick format is used.

*Implementation of Xen machines is based on Microsoft Virtual PC and inherits its limitations: up to 3 IDE disks and 1 processor. SCSI disks are not supported.*

**Processors**

**Initial setting:** if not contained in the backup or if the backed up setting is not supported by the virtualization server, it is the default server's setting.

This is the number of processors of the new virtual machine. In most cases, it is set to one. The result of assignment of more than one processor to the machine is not guaranteed. The number of virtual processors may be limited by the host CPU configuration, the virtualization product and the guest operating system. Multiple virtual processors are generally available on multi-processor hosts. A multicore host CPU or hyperthreading may enable multiple virtual processors on a single-processor host.

# 5.6   Bootability troubleshooting

If a system was bootable at the time of backup, you expect that it will boot after recovery. However, the information the operating system stores and uses for booting up may become outdated during recovery, especially if you change volume sizes, locations or destination drives. Acronis Backup & Recovery 11 automatically updates Windows loaders after recovery. Other loaders might also be fixed, but there are cases when you have to re-activate the loaders. Specifically when you recover Linux volumes, it is sometimes necessary to apply fixes or make booting changes so that Linux can boot and load correctly.

Below is a summary of typical situations that require additional user actions.

Why a recovered operating system may be unbootable

- **The machine BIOS is configured to boot from another HDD.**

    **Solution**: Configure the BIOS to boot from the HDD where the operating system resides.

- **The system was recovered on dissimilar hardware and the new hardware is incompatible with the most critical drivers included in the backup**

    **Solution**: Boot the machine using bootable media and apply Acronis Universal Restore (p. 112) to install the appropriate drivers and modules.

- **Windows was recovered to a dynamic volume that cannot be bootable**

    **Solution**: Recover Windows to a basic, simple or mirrored volume.

- **A    system volume was recovered to a disk that does not have an MBR**

When you configure recovery of a system volume to a disk that does not have an MBR, the program prompts whether you want to recover the MBR along with the system volume. Opt for not recovering, only if you do not want the system to be bootable.

**Solution**: Recover the volume once again along with the MBR of the corresponding disk.

- **The system uses Acronis OS Selector**

  Because the Master Boot Record (MBR) can be changed during the system recovery, Acronis OS Selector, which uses the MBR, might become inoperable. If this happens, reactivate Acronis OS Selector as follows.

  **Solution**: Boot the machine from the Acronis Disk Director's bootable media and select in the menu **Tools -> Activate OS Selector**.

- **The system uses GRand Unified Bootloader (GRUB) and was recovered from a normal (not from a raw, that is, sector-by-sector) backup**

  One part of the GRUB loader resides either in the first several sectors of the disk or in the first several sectors of the volume. The rest is on the file system of one of the volumes. System bootability can be recovered automatically only when the GRUB resides in the first several sectors of the disk and on the file system to which direct access is possible. In other cases, the user has to manually reactivate the boot loader.

  **Solution**: Reactivate the boot loader. You might also need to fix the configuration file.

- **The system uses Linux Loader (LILO) and was recovered from a normal (not from a raw, that is, sector-by-sector) backup**

  LILO contains numerous references to absolute sector numbers and so cannot be repaired automatically except for the case when all data is recovered to the sectors that have the same absolute numbers as on the source disk.

  **Solution**: Reactivate the boot loader. You might also need to fix the loader configuration file for the reason described in the previous item.

- **The system loader points to the wrong volume**

  This may happen when system or boot volumes are not recovered to their original location.

  **Solution**: Modification of the boot.ini or the boot\bcd files fixes this for Windows loaders. Acronis Backup & Recovery 11 does this automatically and so you are not likely to experience the problem.

  For the GRUB and LILO loaders, you will need to correct the GRUB configuration files. If the number of the Linux root partition has changed, it is also recommended that you change /etc/fstab so that the SWAP volume can be accessed correctly.

- **Linux was recovered from an LVM volume backup to a basic MBR disk**

  Such system cannot boot because its kernel tries to mount the root file system at the LVM volume.

  **Solution**: Change the loader configuration and /etc/fstab so that the LVM is not used and reactivate the boot loader.

## 5.6.1   How to reactivate GRUB and change its configuration

Generally, you should refer to the boot loader manual pages for the appropriate procedure. There is also the corresponding Knowledge Base article on the Acronis Web site.

The following is an example of how to reactivate GRUB in case the system disk (volume) is recovered to identical hardware.

1. Start Linux or boot from the bootable media, and then press CTRL+ALT+F2.

2. Mount the system you are recovering:

```
mkdir /mnt/system/
mount -t ext3 /dev/sda2 /mnt/system/ # root partition
mount -t ext3 /dev/sda1 /mnt/system/boot/ # boot partition
```

3. Mount the **proc** and **dev** file systems to the system you are recovering:

```
mount -t proc none /mnt/system/proc/
mount -o bind /dev/ /mnt/system/dev/
```

4. Save a copy of the GRUB menu file, by running one of the following commands:

```
cp /mnt/system/boot/grub/menu.lst /mnt/system/boot/grub/menu.lst.backup
```

   or

```
cp /mnt/system/boot/grub/grub.conf /mnt/system/boot/grub/grub.conf.backup
```

5. Edit the **/mnt/system/boot/grub/menu.lst** file (for Debian, Ubuntu, and SUSE Linux distributions) or the **/mnt/system/boot/grub/grub.conf** file (for Fedora and Red Hat Enterprise Linux distributions)—for example, as follows:

```
vi /mnt/system/boot/grub/menu.lst
```

6. In the **menu.lst** file (respectively **grub.conf**), find the menu item that corresponds to the system you are recovering. This menu items have the following form:

```
title Red Hat Enterprise Linux Server (2.6.24.4)
      root (hd0,0)
      kernel /vmlinuz-2.6.24.4 ro root=/dev/sda2 rhgb quiet
      initrd /initrd-2.6.24.4.img
```

The lines starting with **title**, **root**, **kernel**, and **initrd** respectively determine:

- The title of the menu item.

- The device on which the Linux kernel is located—typically, this is the boot partition or the root partition, such as **root (hd0,0)** in this example.

- The path to the kernel on that device and the root partition—in this example, the path is **/vmlinuz-2.6.24.4** and the root partition is **/dev/sda2**. You can specify the root partition by label (such as **root=LABEL=/**), identifier (in the form **root=UUID=**_some_uuid_), or device name (such as **root=/dev/sda2**).

- The path to the **initrd** service on that device.

7. Edit the file **/mnt/system/etc/fstab** to correct the names of any devices that have changed as a result of the recovery.

8. Start the GRUB shell by running one of the following commands:

```
chroot /mnt/system/ /sbin/grub
```

   or

```
chroot /mnt/system/ /usr/sbin/grub
```

9. Specify the disk on which GRUB is located—typically, the boot or root partition:

```
root (hd0,0)
```

10. Install GRUB. For example, to install GRUB in the master boot record (MBR) of the first disk, run the following command:

```
setup (hd0)
```

11. Exit the GRUB shell:

```
quit
```

12. Unmount the mounted file systems and then reboot:

```
umount /mnt/system/dev/
umount /mnt/system/proc/
umount /mnt/system/boot/
umount /mnt/system/
reboot
```

13. Reconfigure the bootloader by using tools and documentation from the Linux distribution that you use. For example, in Debian and Ubuntu, you may need to edit some commented lines in the **/boot/grub/menu.lst** file and then run the **update-grub** script; otherwise, the changes might not take effect.

## 5.6.2    About Windows loaders

**Windows NT/2000/XP/2003**

A part of the loader resides in the partition boot sector, the rest is in the files ntldr, boot.ini, ntdetect.com, ntbootdd.sys. boot.ini is a text file that contains the loader configuration. Example:

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Microsoft Windows XP Professional"
/noexecute=optin /fastdetect
```

**Windows Vista/2008**

A part of the loader resides in the partition boot sector, the rest is in the files bootmgr, boot\bcd. At starting Windows, boot\bcd is mounted to the registry key HKLM \BCD00000000.

# 5.7    Default recovery options

Each Acronis agent has its own default recovery options. Once an agent is installed, the default options have pre-defined values, which are referred to as **presets** in the documentation. When creating a recovery task, you can either use a default option, or override the default option with the custom value that will be specific for this task only.

You can also customize a default option itself by changing its value against the pre-defined one. The new value will be used by default in all recovery tasks you will create later on this machine.

To view and change the default recovery options, connect the console to the managed machine and then select **Options > Default backup and recovery options > Default recovery options** from the top menu.

**Availability of the recovery options**

The set of available recovery options depends on:

- The environment the agent operates in (Windows, bootable media).
- The type of data being recovered (disk, file).
- The operating system being recovered from the disk backup.

The following table summarizes the availability of the recovery options.

| | Agent for Windows | | Bootable media (Linux-based or PE-based) | |
|---|---|---|---|---|
| | Disk recovery | File recovery (also from a disk backup) | Disk recovery | File recovery (also from a disk backup) |
| **Additional settings (p. 128):** | | | | |
| Validate backup archive before recovery | + | + | + | + |
| Use FTP in Active mode | + | + | + | + |
| Restart the machine automatically if it is required for recovery | + | + | - | - |
| Restart the machine automatically after recovery is finished | - | - | + | + |
| Check file system after recovery | + | - | + | - |
| Change SID after recovery | Windows recovery | - | Windows recovery | - |
| Set current date and time for recovered files | - | + | - | + |
| **Error handling (p. 129):** | | | | |
| Do not show messages and dialogs while processing (silent mode) | + | + | + | + |
| Re-attempt if an error occurs | + | + | + | + |
| **Event tracing:** | | | | |
| Windows events log (p. 130) | + | + | - | - |
| SNMP (p. 129) | + | + | - | - |
| **File-level security (p. 130):** | | | | |
| Recover files with their security settings | - | + | - | + |
| Mount points (p. 130) | - | + | - | - |
| **Notifications:** | | | | |
| E-mail (p. 131) | + | + | - | - |
| Win Pop-up (p. 132) | + | + | - | - |
| Pre/Post recovery commands (p. 132) | + | + | PE only | PE only |
| Recovery priority (p. 134) | + | + | - | - |

## 5.7.1 Additional settings

Specify the additional settings for the recovery operation by selecting or clearing the following check boxes.

### Validate backup archive before recovery

The preset is **Disabled**.

This option defines whether to validate a backup to ensure that the backup is not corrupted, before data is recovered from it.

### Use FTP in Active mode

The preset is: **Disabled**.

Enable this option if the FTP server supports active mode and you want this mode to be used for file transfers.

### Restart machine automatically if it is required for recovery

This option is effective when recovery takes place on a machine running an operating system.

The preset is **Disabled**.

The option defines whether to reboot the machine automatically if it is required for recovery. Such might be the case when a volume locked by the operating system has to be recovered.

### Restart machine automatically after recovery is finished

This option is effective when operating under bootable media.

The preset is **Disabled**.

This option enables booting the machine into the recovered operating system without user interaction.

### Check file system after recovery

This option is effective only when recovering disks or volumes.

When operating under bootable media, this option is not effective for the NTFS file system.

The preset is **Disabled**.

This option defines whether to check the integrity of the file system after a disk or volume recovery.

### Set current date and time for recovered files

This option is effective only when recovering files.

The preset is **Enabled**.

This option defines whether to recover the files' date and time from the archive or assign the files the current date and time.

## 5.7.2 Error handling

These options are effective for Windows and Linux operating systems and bootable media.

These options enable you to specify how to handle errors that might occur during recovery.

**Do not show messages and dialogs while processing (silent mode)**

The preset is: **Disabled**.

With the silent mode enabled, the program will automatically handle situations requiring user interaction where possible. If an operation cannot continue without user interaction, it will fail. Details of the operation, including errors, if any, can be found in the operation log.

**Re-attempt, if an error occurs**

The preset is: **Enabled**. **Number of attempts: 30. Interval between attempts: 30 seconds.**

When a recoverable error occurs, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts is performed, depending on which comes first.

For example, if the network location becomes unavailable or not reachable, the program will attempt to reach the location every 30 seconds, but no more than 5 times. The attempts will be stopped as soon as the connection is resumed OR the specified number of attempts is performed, depending on which comes first.

## 5.7.3 Event tracing

It is possible to duplicate log events of the recovery operations, performed on the managed machine, in the Application Event Log of Windows; or send the events to the specified SNMP managers.

### 5.7.3.1 SNMP notifications

This option is effective for both Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option defines whether the agent(s) operating on the managed machine have to send the log events of the recovery operations to the specified Simple Network Management Protocol (SNMP) managers. You can choose the types of events to be sent.

For detailed information about using SNMP with Acronis Backup & Recovery 11, please see "Support for SNMP (p. 28)".

The preset is: **Use the setting set in the Machine options.**

***To select whether to send the recovery operations events to the SNMP managers:***

Choose one of the following:

- **Use the setting set in the Machine options** – to use the setting specified for the machine. For more information refer to Machine options.
- **Send SNMP notifications individually for recovery operation events** – to send the events of the recovery operations to the specified SNMP managers.

- **Types of events to send** – choose the types of events to be sent: **All events**, **Errors and warnings**, or **Errors only**.

- **Server name/IP** – type the name or IP address of the host running the SNMP management application, the messages will be sent to.

- **Community** – type the name of SNMP community to which both the host running SNMP management application and the sending machine belong. The typical community is "public".

  Click **Send test message** to check if the settings are correct.

- **Do not send SNMP notifications** – to disable sending the log events of the recovery operations to SNMP managers.

### 5.7.3.2    Windows event log

This option is effective only in Windows operating systems.

This option is not available when operating under the bootable media.

This option defines whether the agent(s) operating on the managed machine have to log events of the recovery operations in the Application Event Log of Windows (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

The preset is: **Use the setting set in the Machine options.**

*To select whether to log the recovery operations events in the Application Event Log of Windows:*

Select one of the following:

- **Use the setting set in the Machine options** – to use the setting specified for the machine. For more information refer to Machine options.

- **Log the following event types** – to log events of the recovery operations in the Application Event Log. Specify the types of events to be logged:

  - **All events** – log all events (information, warnings and errors)

  - **Errors and warnings**

  - **Errors only**

- **Do not log** - to disable logging events of the recovery operations in the Application Event Log.

### 5.7.4    File-level security

This option is effective only for recovery from file-level backup of Windows files.

This option defines whether to recover NTFS permissions for files along with the files.

The preset is: **Recover files with their security settings**.

If the file NTFS permissions were preserved during backup (p. 86), you can choose whether to recover the permissions or let the files inherit the NTFS permissions from the folder to which they are recovered.

### 5.7.5    Mount points

This option is effective only in Windows for recovering data from a file-level backup.

Enable the **Mount points** option to recover files and folders that were stored on the mounted volumes and were backed up with the enabled **Mount points** option. For details of backing up mounted volumes or cluster shared volumes, see Mount points (p. 87).

The preset is: **Disabled**.

This option is effective only when you select for backup a folder that is higher in the folder hierarchy than the mount point. If you select for recovery folders within the mount point or the mount point itself, the selected items will be recovered regardless of the **Mount points** option value.

*Please be aware that if the volume is not mounted at the moment of recovery, the data will be recovered directly to the folder that has been the mount point at the time of backing up.*

## 5.7.6    Notifications

Acronis Backup & Recovery 11 provides the ability of notifying users about recovery completion through e-mail or the messaging service.

### 5.7.6.1    E-mail

This option is effective for Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option enables you to receive e-mail notifications about the recovery task's successful completion, failure or need for interaction along with the full log of the task.

The preset is: **Disabled**.

***To configure e-mail notification***

1. Select the **Send e-mail notifications** check box to activate notifications.
2. Under **Send e-mail notifications**, select the appropriate check boxes as follows:
    - **When recovery completes successfully** – to send notification when the recovery task has completed successfully.
    - **When recovery fails** – to send notification when the recovery task has failed.
    - **When user interaction is required** – to send to send notification during the operation when user interaction is required.
3. In the **E-mail addresses** field, type the e-mail address to which notifications will be sent. You can enter several addresses separated by semicolons.
4. In the **Subject** field, type the notification subject or leave the default value.
5. In the **SMTP server** field**,** enter the name of the SMTP server.
6. In the **Port** field– set the port of the SMTP server. By default, the port is set to **25**.
7. In the **User name** field, enter the user name.
8. In the **Password** field, enter the password.
9. Click **Additional e-mail parameters...** to configure the additional e-mail parameters as follows:
    a. **From** - type the e-mail address of the user from whom the message will be sent. If you leave this field empty, messages will be constructed as if they are from the destination address.
    b. **Use encryption** – you can opt for encrypted connection to the mail server. SSL and TLS encryption types are available for selection.

Copyright © Acronis, Inc., 2000-2011

c. Some Internet service providers require authentication on the incoming mail server before being allowed to send something. If this is your case, select the **Log on to incoming mail server** check box to enable a POP server and to set up its settings:

- **Incoming mail server (POP)** – enter the name of the POP server.
- **Port** – set the port of the POP server. By default, the port is set to **110**.
- **User name** – enter the user name.
- **Password** – enter the password.

d. Click **OK**.

10. Click **Send test e-mail message** to check if the settings are correct.

## 5.7.6.2 Messenger service (WinPopup)

This option is effective for Windows and Linux operating systems.

This option is not available when operating under bootable media.

The option enables you to receive WinPopup notifications about about the recovery task's successful completion, failure or need for interaction.

The preset is: **Disabled.**

Before configuring WinPopup notifications, make sure the Messenger service is started on both the machine executing the task and the machine that will receive messages.

The Messenger service is not started by default in the Microsoft Windows Server 2003 family. Change the service Startup mode to Automatic and start the service.

***To configure WinPopup notifications:***

1. Select the **Send WinPopup notifications** check box.
2. In the **Machine name** field, enter the name of the machine to which notifications will be sent. Multiple names are not supported.
3. Under **Send notifications**, select the appropriate check boxes as follows:

- **When recovery completes successfully** – to send notification when the recovery task has completed successfully
- **When recovery fails** – to send notification when the recovery task has failed.
- **When user interaction is required** check box – to send notification during the operation when user interaction is required.

4. Click **Send Test WinPopup Message** to check if the settings are correct.

## 5.7.7 Pre/Post commands

This option is effective for Windows and Linux operating systems and PE-based bootable media.

The option enables you to define the commands to be automatically executed before and after the data recovery.

Example of how you can use the pre/post commands:

- Launch the **Checkdisk** command in order to find and fix logical file system errors, physical errors or bad sectors to be started before the recovery starts or after the recovery ends.

The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)

A post-recovery command will not be executed if the recovery proceeds with reboot.

***To specify pre/post commands***

1. Enable pre/post commands execution by checking the following options:
   - **Execute before the recovery**
   - **Execute after the recovery**
2. Do any of the following:
   - Click **Edit** to specify a new command or a batch file
   - Select the existing command or the batch file from the drop-down list
3. Click **OK**.

## 5.7.7.1    Pre-recovery command

***To specify a command/batch file to be executed before the recovery process starts***

1. In the **Command** field, type a command or browse to a batch file. The program does not support interactive commands, i.e. commands that require user input (for example, "pause".)
2. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
3. In the **Arguments** field specify the command's execution arguments, if required.
4. Depending on the result you want to obtain, select the appropriate options as described in the table below.
5. Click **Test command** to check if the command is correct.

| Check box | Selection | | | |
|---|---|---|---|---|
| **Fail the task if the command execution fails\*** | Selected | Cleared | Selected | Cleared |
| **Do not recover until the command execution is complete** | Selected | Selected | Cleared | Cleared |
| Result | | | | |
| | **Preset**<br><br>Perform the recovery only after the command is successfully executed. Fail the task if the command execution failed. | Perform the recovery after the command is executed despite execution failure or success. | N/A | Perform the recovery concurrently with the command execution and irrespective of the command execution result. |

\* A command is considered failed if its exit code is not equal to zero.

## 5.7.7.2    Post-recovery command

***To specify a command/executable file to be executed after the recovery is completed***

1. In the **Command** field, type a command or browse to a batch file.

2. In the **Working** directory field, specify a path to a directory where the command/batch file will be executed.

3. In the **Arguments** field, specify the command execution arguments, if required.

4. Select the **Fail the task if the command execution fails** check box if successful execution of the command is critical for you. The command is considered failed if its exit code is not equal to zero. If the command execution fails, the task run result will be set to Failed.

   When the check box is not selected, the command execution result does not affect the task execution failure or success. You can track the command execution result by exploring the **Log** view.

5. Click **Test command** to check if the command is correct.

*A post-recovery command will not be executed if the recovery proceeds with reboot.*

## 5.7.8    Recovery priority

This option is effective for both Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The priority of a process running in a system determines the amount of CPU and system resources allocated to that process. Decreasing the recovery priority will free more resources for other applications. Increasing the recovery priority might speed up the recovery process by requesting the operating system to allocate more resources to the application that will perform the recovery. However, the resulting effect will depend on the overall CPU usage and other factors like disk I/O speed or network traffic.

The preset is: **Normal.**

### *To specify the recovery process priority*

Select one of the following:

- **Low** – to minimize resources taken by the recovery process, leaving more resources to other processes running on the machine

- **Normal** – to run the recovery process with normal speed, allocating resources on a par with other processes

- **High** – to maximize the recovery process speed by taking resources from the other processes.

# 6   Storing the backed up data

## 6.1   Vaults

A vault is a location for storing backup archives. For ease of use and administration, a vault is associated with the archives' metadata. Referring to this metadata makes for fast and convenient operations with archives and backups stored in the vault.

A vault can be organized on a local or networked drive, detachable media or a tape device attached to the Acronis Backup & Recovery 11 Storage Node.

There are no settings for limiting a vault size or number of backups in a vault. You can limit the size of each archive using cleanup, but the total size of archives stored in the vault is limited by the storage size only.

**Why create vaults?**

We recommend that you create a vault in each destination where you are going to store backup archives. This will ease your work as follows.

**Quick access to the vault**

You will not have to remember paths to the folders where the archives are stored. When creating a backup plan or a task that requires selection of an archive or an archive destination place, the list of vaults will be available for quick access without drilling down through the folders tree.

**Easy archive management**

A vault is available for access from the **Navigation** pane. Having selected the vault, you can browse the archives stored there and perform the following archive management operations:

- Get a list of backups included in each archive
- Recover data from a backup
- Examine backup content
- Validate all archives in the vault or individual archives or backups
- Mount a volume backup to copy files from the backup to a physical disk
- Safely delete archives and backups from the archives.

Creating vaults is highly recommended but is not obligatory. You may choose not to use the shortcuts and always specify the location path.

Creating a vault results in adding the vault name to the **Vaults** section of the **Navigation** pane.

**'Vaults' view**

**Vaults** (on the navigation pane) - top item of the vaults tree. Click this item to display centralized and personal vaults. To perform actions on any vault, use the toolbar that is located at the top of the **Vaults** view. See the Actions on personal vaults (p. 137) section.

> **Personal vaults.** These vaults available when the console is connected to a managed machine. Click any vault in the vaults tree to open the detailed view of this vault (p. 136) and to take actions on archives (p. 155) and backups (p. 156) stored in there.

## 6.1.1 Working with vaults

This section briefly describes the main GUI elements of the selected vault, and suggests ways to work with them.

### Examining information on a vault

Information about the selected vault is located at the top pane of the selected vault. Using the stacked bar, you can estimate the vault's load. The vault's load is   the proportion of the vault's free space and occupied space (not available if the vault is located on a tape library). Free space is a space on the storage device where the vault is located. For example, if the vault is located on a hard disk, the vault free space is the free space of the respective volume. Occupied space is the total size of backup archives and their metadata, if it is located in the vault.

You can obtain the total number of archives and backups stored in the vault and full path to the vault.

For managed vaults only, you can examine the name of the storage node that manages the vault, encryption and deduplication states.

### Browsing the vault contents and data selection

You can browse the vault content and select data to recover by using the **Data view** tab, or the **Archive view** tab.

**Data view**

The **Data view** tab lets you browse and select the backed up data by versions (backup date and time). The **Data view** tab shares the same searching and cataloging functionality with the data catalog (p. 101).

**Archive view**

The **Archive view** tab displays the backed up data by archives. Use the **Archive view** to perform operations with archives and backups stored in the vault. For more information about these operations, see the following sections:

- Operations with archives stored in a vault (p. 155).
- Operations with backups (p. 156).
- Sorting, filtering and configuring table items (p. 16).

## 6.1.2 Personal vaults

A vault is called personal if it was created using direct connection of the console to a managed machine. Personal vaults are specific for each managed machine. Personal vaults are visible to any user that can log on to the system. A user's right to back up to a personal vault is defined by the user's permission for the folder or device where the vault is located.

A personal vault can be organized on a network share, FTP server, detachable media or removable, Acronis Online Backup Storage, tape device, or on a hard drive local to the machine. Acronis Secure Zone is considered as a personal vault available to all users that can log on the system. Personal vaults are created automatically when backing up any of the above locations.

Personal vaults can be used by local backup plans or local tasks. Centralized backup plans cannot use personal vaults except for Acronis Secure Zone.

**Sharing a personal vault**

Multiple machines can refer to the same physical location; for example, to the same shared folder. However, each of the machines has its own shortcut in the **Vaults** tree. Users that back up to a shared folder can see and manage each other's archives according to their access permissions for that folder. To ease archive identification, the **Personal vault** view has the **Owner** column that displays the owner of each archive. To find out more about the owner concept see Owners and credentials (p. 21).

**Metadata**

The **.meta** folder is created during backup in every personal vault. This folder contains additional information about archives and backups stored in the vault, such as archive owners or the machine name. If you accidentally delete the .meta folder, it will be automatically recreated next time you access the vault. But some information, like owner names and machine names, may be lost.

## 6.1.2.1    Actions on personal vaults

***To access actions***

1.  Connect the console to the management server.
2.  In the **Navigation** pane, click **Vaults** > **Personal**.

All the operations described here are performed by clicking the corresponding buttons on the vaults toolbar. These operations can be also accessed from the **[Vault name] actions** item of the main menu.

The following is a guideline for you to perform operations with personal vaults.

| To | Do |
|---|---|
| Create a personal vault | Click  **Create**. <br><br> The procedure of creating personal vaults is described in-depth in the Creating a personal vault (p. 138) section. |
| Edit a vault | 1.  Select the vault. <br><br> 2.  Click  **Edit**. <br><br> The **Edit personal vault** page lets you edit the vault's name and information in the **Comments** field. |
| Change user account for accessing a vault | Click  **Change user**. <br><br> In the appearing dialog box, provide the credentials required for accessing the vault. |
| Create Acronis Secure Zone | Click  **Create Acronis Secure Zone**. <br><br> The procedure of creating the Acronis Secure Zone is described in-depth in the Creating Acronis Secure Zone (p. 139) section. |
| Explore a vault's content | Click  **Explore**. <br><br> In the appearing Explorer window, examine the selected vault's content. |
| Validate a vault | Click  **Validate**. <br><br> You will be taken to the Validation (p. 143) page, where this vault is already pre-selected as a source. The vault validation checks all the archives stored in the vault. |

| Delete a vault | Click ❌ **Delete**.<br><br>The deleting operation actually removes only a shortcut to the folder from the **Vaults** view. The folder itself remains untouched. You have the option to keep or delete archives contained in the folder. |
|---|---|
| Refresh vault table information | Click 🔄 **Refresh**.<br><br>While you are reviewing the vault content, archives can be added to the vault, deleted or modified. Click **Refresh** to update the vault information with the most recent changes. |

## Creating a personal vault

*To create a personal vault*

1. In the **Name** field, type a name for the vault being created.

2. [Optional] In the **Comments** field, add a description of the vault.

3. Click **Path** and specify a path to the folder that will be used as the vault. A personal vault can be organized on a network share, FTP server, detachable media, Acronis Online Backup Storage, tape device, or on a hard drive local to the machine.

4. [Optional] If the vault is created on a tape device:

    a. Click **Drives** to specify the tape drive(s) to be used when backing up to the vault. By default, all available drives will be used. Click **Use the following drives only** and select or clear required check boxes;

    b. Click **Tape pool** and specify the pool whose tapes will be used by the vault. By default, the **Acronis** pool is selected.

5. Click **OK**. As a result, the created vault appears in the **Personal** group of the vaults tree.

## Merging and moving personal vaults

### What if I need to move the existing vault from one place to another?

Proceed as follows

1. Make sure that none of the backup plans uses the existing vault while moving files, or disable the given plans. See Actions on backup plans and tasks (p. 188).

2. Move the vault folder with all its content to a new place manually by means of a third-party file manager.

3. Create a new vault.

4. Edit the backup plans and tasks: redirect their destination to the new vault.

5. Delete the old vault.

### How can I merge two vaults?

Suppose you have two vaults *A* and *B* in use. Both vaults are used by backup plans. You decide to leave only vault *B*, moving all the archives from vault *A* there.

To do this, proceed as follows

1. Make sure that none of the backup plans uses vault *A* while merging, or disable the given plans. See Actions on backup plans and tasks (p. 188).

2. Move the content of vault *A* folder to vault *B* manually by means of a third-party file manager.

3. Edit the backup plans that use vault *A*: redirect their destination to vault *B*.

4.  In the vaults tree, select vault *B* to check whether the archives are displayed. If not, click **Refresh**.

5.  Delete vault *A*.

# 6.2    Acronis Secure Zone

Acronis Secure Zone is a secure partition that enables keeping backup archives on a managed machine disk space and therefore recovery of a disk to the same disk where the backup resides.

Should the disk experience a physical failure, the zone and the archives located there will be lost. That's why Acronis Secure Zone should not be the only location where a backup is stored. In enterprise environments, Acronis Secure Zone can be thought of as an intermediate location used for backup when an ordinary location is temporarily unavailable or connected through a slow or busy channel.

**Advantages**

Acronis Secure Zone:

- Enables recovery of a disk to the same disk where the disk's backup resides.
- Offers a cost-effective and handy method for protecting data from software malfunction, virus attack, operator error.
- Since it is internal archive storage, it eliminates the need for a separate media or network connection to back up or recover the data. This is especially useful for mobile users.
- Can serve as a primary destination when using replication of backups (p. 68).

**Limitations**

- The Acronis Secure Zone cannot be organized on a dynamic disk.

# 6.2.1    Creating Acronis Secure Zone

You can create Acronis Secure Zone while the operating system is running or using bootable media.

***To create Acronis Secure Zone, perform the following steps.***

**Location and size**

Disk (p. 140)

> Choose a hard disk (if several) on which to create the zone. Acronis Secure Zone is created using unallocated space, if available, or at the expense of the volume's free space.

Size (p. 140)

> Specify the exact size of the zone. Moving or resizing of locked volumes, such as the volume containing the currently active operating system, requires a reboot.

**Security**

Password (p. 140)

> [Optional] Protect the Acronis Secure Zone from unauthorized access with a password. The prompt for the password appear at any operation relating to the zone.

After you configure the required settings, click OK. In the Result confirmation (p. 141) window, review the expected layout and click OK to start creating the zone.

## 6.2.1.1 Acronis Secure Zone Disk

The Acronis Secure Zone can be located on any fixed hard drive. Acronis Secure Zone is always created at the end of the hard disk. A machine can have only one Acronis Secure Zone. Acronis Secure Zone is created using unallocated space, if available, or at the expense of the volumes' free space.

*The Acronis Secure Zone cannot be organized on a dynamic disk.*

**To allocate space for Acronis Secure Zone**

1. Choose a hard disk (if several) on which to create the zone. The unallocated space and free space from all volumes of the first enumerated disk are selected by default. The program displays the total space available for the Acronis Secure Zone.

2. If you need to allocate more space for the zone, you can select volumes from which free space can be taken. Again, the program displays the total space available for the Acronis Secure Zone depending on your selection. You will be able to set the exact zone size in the **Acronis Secure Zone Size** (p. 140) window.

3. Click **OK**.

## 6.2.1.2 Acronis Secure Zone Size

Enter the Acronis Secure Zone size or drag the slider to select any size between the minimum and the maximum ones. The minimum size is approximately 50MB, depending on the geometry of the hard disk. The maximum size is equal to the disk's unallocated space plus the total free space on all the volumes you have selected in the previous step.

If you have to take space from the boot or the system volume, please bear the following in mind:

- Moving or resizing of the volume from which the system is currently booted will require a reboot.
- Taking all free space from a system volume may cause the operating system to work unstably and even fail to start. Do not set the maximum zone size if the boot or the system volume is selected.

## 6.2.1.3 Password for Acronis Secure Zone

Setting up a password protects the Acronis Secure Zone from unauthorized access. The program will ask for the password at any operation relating to the zone and the archives located there, such as data backup and recovery, validating archives, resizing and deleting the zone.

**To set up a password**

1. Choose **Use password**.
2. In the **Enter the password** field, type a new password.
3. In the **Confirm the password** field, re-type the password.
4. Click **OK**.

**To disable password**

1. Choose **Do not use**.
2. Click **OK**.

## 6.2.1.4  Result confirmation

The **Result confirmation** window displays the expected partition layout according to the settings you have chosen. Click **OK**, if you are satisfied with the layout and the Acronis Secure Zone creation will start.

**How the settings you make will be processed**

This helps you to understand how creating the Acronis Secure Zone will transform a disk containing multiple volumes.

- Acronis Secure Zone is always created at the end of the hard disk. When calculating the final layout of the volumes, the program will first use unallocated space at the end.

- If there is no or not enough unallocated space at the end of the disk, but there is unallocated space between volumes, the volumes will be moved to add more unallocated space to the end.

- When all unallocated space is collected but it is still not enough, the program will take free space from the volumes you select, proportionally reducing the volumes' size. Resizing of locked volumes requires a reboot.

- However, there should be free space on a volume, so that the operating system and applications can operate; for example, for creating temporary files. The program will not decrease a volume where free space is or becomes less than 25% of the total volume size. Only when all volumes on the disk have 25% or less free space, will the program continue decreasing the volumes proportionally.

As is apparent from the above, setting the maximum possible zone size is not advisable. You will end up with no free space on any volume which might cause the operating system or applications to work unstably and even fail to start.

# 6.2.2  Managing Acronis Secure Zone

Acronis Secure Zone is considered as a personal vault (p. 240). Once created on a managed machine, the zone is always present in the list of **Personal vaults**. Centralized backup plans can use Acronis Secure Zone as well as local plans.

All the archive management operations available in vaults are also applicable for Acronis Secure Zone. To learn more about archive management operations, see Operations with archives and backups (p. 155).

## 6.2.2.1  Increasing Acronis Secure Zone

***To increase Acronis Secure Zone***

1. On the **Manage Acronis Secure Zone** page, click **Increase**.
2. Select volumes from which free space will be used to increase the Acronis Secure Zone.
3. Specify the new size of the zone by:
   - dragging the slider and selecting any size between the current and maximum values. The maximum size is equal to the disk's unallocated space plus the total free space of all selected partitions;
   - typing an exact value in the Acronis Secure Zone Size field.

   When increasing the size of the zone, the program will act as follows:
   - first, it will use the unallocated space. Volumes will be moved, if necessary, but not resized. Moving of locked volumes requires a reboot.

- If there is not enough unallocated space, the program will take free space from the selected volumes, proportionally reducing the volumes' size. Resizing of locked partitions requires a reboot.

*Reducing a system volume to the minimum size might prevent the machine's operating system from booting.*

4. Click **OK**.

## 6.2.2.2  Decreasing Acronis Secure Zone

***To decrease Acronis Secure Zone***

1. On the **Manage Acronis Secure Zone** page, click **Decrease**.
2. Select volumes that will receive free space after the zone is decreased.
3. Specify the new size of the zone by:
    - dragging the slider and selecting any size between the current and minimum values. The minimum size is approximately 50MB, depending on the geometry of the hard disk;
    - typing an exact value in the **Acronis Secure Zone Size** field.
4. Click **OK**.

## 6.2.2.3  Deleting Acronis Secure Zone

***To delete Acronis Secure Zone:***

1. On the **Manage Acronis Secure Zone** page, click **Delete**.
2. In the **Delete Acronis Secure Zone** window, select volumes to which you want to add the space freed from the zone and then click **OK**.

    If you select several volumes, the space will be distributed proportionally to each partition. If you do not select any volume, the freed space becomes unallocated.

After you click **OK**, Acronis Backup & Recovery 11 will start deleting the zone.

# 7 Operations with archives and backups

## 7.1 Validating archives and backups

Validation is an operation that checks the possibility of data recovery from a backup.

Validation of a file backup imitates recovery of all files from the backup to a dummy destination. Validation of a disk or volume backup calculates a checksum for every data block saved in the backup. Both procedures are resource-intensive.

Validation of an archive will validate all the archive's backups. A vault (or a location) validation will validate all archives stored in this vault (location).

While successful validation means high probability of successful recovery, it does not check all factors that influence the recovery process. If you back up the operating system, only a test recovery in a bootable environment to a spare hard drive can guarantee success of the recovery. At least ensure that the backup can be successfully validated using the bootable media.

### Different ways to create a validation task

Using the **Validation** page is the most general way to create a validation task. Here you can validate immediately or set up a validation schedule for any backup, archive or vault you have permission to access.

Validation of an archive or of the latest backup in the archive can be scheduled as part of the backup plan. For more information, see Creating a backup plan (p. 29).

To access the **Validation** page first select a validation object: a vault, an archive, or a backup.

- To select a vault, click the **Vaults** icon in the **Navigation** pane and select the vault by expanding the vaults tree in the **Vaults** view or directly in the **Navigation** pane.

- To select an archive, select a vault, and then in the **Vault** view select the **Archive view** tab and click the archive name.

- To select a backup, select an archive in the **Archive view**, expand the archive by clicking the expand button to the left of the archive name, and then click the backup.

After selecting the validation object, select **Validate** from the context menu. The **Validation** page will be opened with the pre-selected object as a source. All you need to do is to select when to validate and (optionally) provide a name for the task.

***To create a validation task, perform the following steps.***

### What to validate

**Validate**

Choose an object to validate:

**Archive** (p. 149) - in this case, you need to specify the archive.

**Backup** (p. 144) - specify the archive first. Then, select the desired backup in this archive.

**Vault** (p. 144) - select a vault (or other location), to validate archives from.

**Credentials** (p. 145)

[Optional] Provide credentials for accessing the source if the task account does not have enough privileges to access it.

**When to validate**

> **Start validation** (p. 145)
>
> > Specify when and how often to perform validation.

**Task parameters**

> **Task name**
>
> > [Optional] Enter a unique name for the validation task. A conscious name lets you quickly identify the task among the others.
>
> **Task's credentials** (p. 146)
>
> > [Optional] The validation task will run on behalf of the user who is creating the task. You can change the task credentials if necessary.
>
> **Comments**
>
> > [Optional] Enter comments on the task.

After you configure all the required settings, click **OK** to create the validation task.

## 7.1.1    Archive selection

**Selecting the archive**

1.  Enter the full path to the archive location in the **Path** field, or select the required location in the tree (p. 100).

    When operating on a machine booted with bootable media:

    ▪ To access a managed vault, type the following string in the **Path** field:

    **bsp://node_address/vault_name/**

    ▪ To access an unmanaged centralized vault, type the full path to the vault's folder.

2.  In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each location you select.

    While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3.  Click **OK**.

## 7.1.2    Backup selection

*To specify a backup to validate*

1.  In the upper pane, select a backup by its creation date/time.

    The bottom part of the window displays the selected backup content, assisting you to find the right backup.

2.  Click **OK**.

## 7.1.3    Vault selection

*To select a vault or a location*

1.  Enter the full path to the vault (location) in the **Path** field or select the desired location in the tree.

- To select a centralized vault, expand the **Centralized** group and click the appropriate vault.
- To select a personal vault, expand the **Personal** group and click the appropriate vault.
- To select a local folder    (CD/DVD drive, or locally attached tape device), expand the **Local folders** group and click the required folder.
- To select a network share, expand the **Network folders** group, select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.
- To select a folder stored on NFS share, expand the **NFS drives** group and click the folder.
- To select **FTP** or **SFTP** server, expand the corresponding group and click the appropriate folder on the server.

  *According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

  To assist you with choosing the right vault, the table displays the names of the archives contained in each vault you select. While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

1. Click **OK**.

## 7.1.4   Access credentials for source

Specify the credentials required for access to the location where the backup archive is stored.

***To specify credentials***

1. Select one of the following:

   - **Use the task credentials**

     The software will access the location using the credentials of the task account specified in the **Task parameters** section.

   - **Use the following credentials**

     The software will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

     Specify:
     - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).
     - **Password**. The password for the account.

2. Click **OK**.

*According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

## 7.1.5   When to validate

As validation is a resource-intensive operation, it makes sense to schedule validation to the managed machine's off-peak period. On the other hand, if you prefer to be immediately informed whether the data is not corrupted and can be successfully recovered, consider starting validation right after the task creation.

*Choose one of the following:*

- **Now** - to start the validation task right after its creation, that is, after clicking OK on the Validation page.
- **Later** - to start the one-time validation task, at the date and time you specify.

  Specify the appropriate parameters as follows:

  - **Date and time** - the date and time when to start the task.
  - **The task will be started manually (do not schedule the task)** - select this check box, if you wish to start the task manually later.

- **On schedule** - to schedule the task. To learn more about how to configure the scheduling parameters, please see the Scheduling (p. 53) section.

## 7.1.6    Task credentials

Provide credentials for the account under which the task will run.

*To specify credentials*

1.  Select one of the following:

    - **Run under the current user**

      The task will run under the credentials with which the user who starts the tasks is logged on. If the task has to run on schedule, you will be asked for the current user's password on completing the task creation.

    - **Use the following credentials**

      The task will always run under the credentials you specify, whether started manually or executed on schedule.

      Specify:

      - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).
      - **Password**. The password for the account.

2.  Click **OK**.

To learn more about using credentials in Acronis Backup & Recovery 11, see the Owners and credentials (p. 21) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine section.

# 7.2    Exporting archives and backups

The export operation creates a copy of an archive or a self-sufficient part copy of an archive in the location you specify. The original archive remains untouched.

The export operation can be applied to:

- **A single archive** - an exact archive copy will be created.
- **A single backup** - an archive consisting of a single full backup will be created. The export of an incremental or a differential backup is performed using consolidation of the preceding backups up to the nearest full backup.

- **Your choice of backups** belonging to the same archive - the resulting archive will contain only the specified backups. Consolidation is performed as required, so the resulting archive may contain full, incremental and differential backups.

**Usage scenarios**

Export enables you to separate a specific backup from a chain of incremental backups for fast recovery, writing onto removable or detachable media or other purposes.

**Example.** When backing up data to a remote location through an unstable or low-bandwidth network connection (such as backing up through WAN using VPN access), you may want to save the initial full backup to a detachable media. Then, send the media to the remote location. There, the backup will be exported from the media to the target storage. Subsequent incremental backups, which are usually much smaller, can be transferred over the network.

By exporting a managed vault to a detachable media, you obtain a portable unmanaged vault that can be used in the following scenarios:

- Keeping an off-site copy of your vault or of the most important archives.
- Physical transportation of a vault to a distant branch office.
- Recovery without access to the storage node in case of networking problems or failure of the storage node.
- Recovery of the storage node itself.

Export from an HDD-based vault to a tape device can be considered as simple on-demand archive staging.

**The resulting archive's name**

By default, the exported archive inherits the name of the original archive. Because having multiple archives of the same names in the same location is not advisable, the following actions are disabled with the default archive name:

- Exporting part of an archive to the same location.
- Exporting an archive or part of an archive to a location where an archive of the same name exists.
- Exporting an archive or part of an archive to the same location twice.

In any of the above cases, provide an archive name that is unique to the destination folder or vault. If you need to redo the export using the same archive name, first delete the archive that resulted from the previous export operation.

**The resulting archive's options**

The exported archive inherits the options of the original archive, including encryption and the password. When exporting a password-protected archive, you are prompted for the password. If the original archive is encrypted, the password is used to encrypt the resulting archive.

**Source and destination locations**

When the console is connected to a **managed machine**, you can export an archive or part of an archive to and from any location accessible to the agent residing on the machine. These include personal vaults, locally attached tape devices, removable media and, in the advanced product versions, managed and unmanaged centralized vaults.

When the console is connected to a **management server**, two export methods are available:

- Export from a **managed vault**. The export is performed by the storage node that manages the vault. The destination can be a network share or a local folder of the storage node.
- Export from an **unmanaged centralized vault**. The export is performed by the agent installed on the managed machine you specify. The destination can be any location accessible to the agent, including a managed vault.

*Tip. When configuring export to a deduplicating managed vault, choose a machine where the deduplication add-on to the agent is installed. Otherwise the export task will fail.*

**Operations with an export task**

An export task starts immediately after you complete its configuration. An export task can be stopped or deleted in the same way as any other task.

Once the export task is completed, you can run it again at any time. Before doing so, delete the archive that resulted from the previous task run if the archive still exists in the destination vault. Otherwise the task will fail. You cannot edit an export task to specify another name for the destination archive (this is a limitation).

*Tip. You can implement the staging scenario manually, by regularly running the archive deletion task followed by the export task.*

**Different ways to create an export task**

Using the **Export** page is the most general way to create an export task. Here, you can export any backup, or archive you have permission to access.

You can access the **Export** page from the **Vaults** view. Right-click the object to export (archive or backup) and select **Export** from the context menu.

To access the **Export** page first select a validation object: an archive or a backup.

1. Select a vault. For this click the **Vaults** icon in the **Navigation** pane and select the vault expanding the vaults tree in the **Vaults** view or directly in the **Navigation** pane.
2. To select an archive, select a vault, and then in the **Vault** view select the **Archive view** tab and click the archive name.
3. To select a backup, select an archive in the **Archive view**, expand the archive by clicking the expand button to the left of archive name, and then click the backup.

After selecting the validation object, select **Export** from the context menu. The **Export** page will be opened with the pre-selected object as a source. All you need to do is to select a destination and (optionally) provide a name for the task.

***To export an archive or a backup perform the following steps.***

### What to export

#### Export

Select the type of objects to export:

**Archive** - in this case, you need to specify the archive only.

**Backups** - you need to specify the archive first, and then select the desired backup(s) in this archive.

#### Browse

Select the **Archive** (p. 149) or the **Backups** (p. 149).

**Show access credentials** (p. 149)

[Optional] Provide credentials for accessing the source if the task account does not have enough privileges to access it.

**Where to export**

**Browse** (p. 150)

Specify the path to the location where the new archive will be created.

Be sure to provide a distinct name and comment for the new archive.

**Show access credentials** (p. 151)

[Optional] Provide credentials for the destination if the task credentials do not have enough privileges to access it.

After you have performed all the required steps, click **OK** to start the export task.

As a result, the program shows the **Execution state** of the task in the **Backup plans and tasks** view. When the task ends the **Task Information** window shows the final state of the task execution.

## 7.2.1    Archive selection

**Selecting the archive**

1.  Enter the full path to the archive location in the **Path** field, or select the required location in the tree (p. 100).

2.  In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each location you select.

    While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3.  Click **OK**.

## 7.2.2    Backup selection

*To specify a backup(s) to export*

1.  At the top of the window, select the respective check box(es).

    To ensure that you choose the right backup, click on the backup and look at the bottom table that displays the volumes contained in the selected backup.

    To obtain information on a volume, right-click it and then select **Information**.

2.  Click **OK**.

## 7.2.3    Access credentials for source

Specify credentials required for access to the location where the source archive, or the backup is stored.

*To specify credentials*

1.  Select one of the following:

    ▪ **Use the current user credentials**

       The software will access the location using the credentials of the current user.

- **Use the following credentials**

  The program will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

  Specify:

  - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).

  - **Password**. The password for the account.

2. Click **OK**.

*According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

# 7.2.4 Destination selection

Specify a destination where the exported object will be stored. Exporting backups to the same archive is not allowed.

## 1. Selecting the export destination

Enter the full path to the destination in the **Path** field, or select the desired destination in the tree.

- To export data to a centralized unmanaged vault, expand the **Centralized vaults** group and click the vault.

- To export data to a personal vault, expand the **Personal vaults** group and click the vault.

- To export data to a local folder on the machine, expand the **Local folders** group and click the required folder.

- To export data to a network share, expand the **Network folders** group, select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

  *Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.*

- To export data to an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

  **ftp://ftp_server:port _number** or **sftp://sftp_server:port number**

  If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

  After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

  You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

  *According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

- To export data to a locally attached tape device, expand the **Tape drives** group, then click the required device. In stand-alone editions of Acronis Backup & Recovery 11, tape devices are available only if you have upgraded from Acronis Backup & Recovery 10. For information about using tapes, see the "Tape devices" section.

### 2. Using the archives table

To assist you with choosing the right destination, the table on the right displays the names of the archives contained in each location you select in the tree.

While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

### 3. Naming the new archive

By default, the exported archive inherits the name of the original archive. Because having multiple archives of the same names in the same location is not advisable, the following actions are disabled with the default archive name:

- Exporting part of an archive to the same location.
- Exporting an archive or part of an archive to a location where an archive of the same name exists.
- Exporting an archive or part of an archive to the same location twice.

In any of the above cases, provide an archive name that is unique to the destination folder or vault. If you need to redo the export using the same archive name, first delete the archive that resulted from the previous export operation.

## 7.2.5    Access credentials for destination

Specify credentials required for access to the location where the resulting archive will be stored. The user whose name is specified will be considered as the archive owner.

***To specify credentials***

1.  Select one of the following:

    - **Use the current user credentials**

      The software will access the destination using the credentials of the current user.

    - **Use the following credentials**

      The software will access the destination using the credentials you specify. Use this option if the task account does not have access permissions to the destination.

      Specify:

      - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).
      - **Password**. The password for the account.

2.  Click **OK**.

*According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

## 7.3    Mounting an image

Mounting volumes from a disk backup (image) lets you access the volumes as though they were physical disks. Multiple volumes contained in the same backup can be mounted within a single

mount operation. The mount operation is available when the console is connected to a managed machine running either Windows or Linux.

Mounting volumes in the read/write mode enables you to modify the backup content, that is, save, move, create, delete files or folders, and run executables consisting of one file.

*You can mount volumes if the disk backup is stored in a local folder (excepting removable media), Acronis Secure Zone or on a network share.*

## Usage scenarios

- **Sharing**: mounted images can be easily shared to networked users.
- **"Band aid" database recovery solution**: mount up an image that contains an SQL database from a recently failed machine. This will provide access to the database until the failed machine is recovered.
- **Offline virus clean**: if a machine is attacked, the administrator shuts it down, boots with bootable media and creates an image. Then, the administrator mounts this image in read/write mode, scans and cleans it with an antivirus program, and finally recovers the machine.
- **Error check**: if recovery failed due to a disk error, mount the image in the read/write mode. Then, check the mounted disk for errors with the **chkdsk /r** command.

***To mount an image, perform the following steps.***

## Source

**Archive** (p. 152)

Specify the path to the archive location and select the archive containing disk backups.

**Backup** (p. 153)

Select the backup.

**Access credentials** (p. 154)

[Optional] Provide credentials for the archive location.

## Mount settings

**Volumes** (p. 154)

Select volumes to mount and configure the mount settings for every volume: assign a letter or enter the mount point, choose the read/write or read only access mode.

When you complete all the required steps, click **OK** to mount the volumes.

# 7.3.1 Archive selection

## Selecting the archive

1. Enter the full path to the location in the **Path** field, or select the desired folder in the folders tree.
   - If the archive is stored in Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then expand the **Online backup storage** group and select the account.

   *Exporting and mounting are not supported for backups stored in Acronis Online Backup Storage.*

   - If the archive is stored in a centralized vault, expand the **Centralized** group and click the vault.
   - If the archive is stored in a personal vault, expand the **Personal** group and click the vault.

- If the archive is stored in a local folder on the machine, expand the **Local folders** group and click the required folder.

*If the archive is located on removable media, e.g. DVDs, first insert the last DVD and then insert the discs in order starting from the first one when the program prompts.*

- If the archive is stored on a network share, expand the **Network folders** group, then select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

*Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.*

- If the archive is stored on an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

  **ftp://ftp_server:port _number or sftp://sftp_server:port number**

  If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

  After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

  You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

*According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

- If the archive is stored on a locally attached tape device, expand the **Tape drives** group, then click the required device.

When operating on a machine booted with bootable media:

- To access a managed vault, type the following string in the **Path** field:

  **bsp://node_address/vault_name/**

- To access an unmanaged centralized vault, type the full path to the vault's folder.

2. In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each vault/folder you select.

   While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3. Click **OK**.

## 7.3.2    Backup selection

***To select a backup:***

1. Select one of the backups by its creation date/time.

2. To assist you with choosing the right backup, the bottom table displays the volumes contained in the selected backup.

   To obtain information on a volume, right-click it and then click **Information**.

3. Click **OK**.

## 7.3.3 Access credentials

***To specify credentials***

1. Select one of the following:

   - **Use the current user credentials**

     The program will access the location using the credentials of the current user.

   - **Use the following credentials**

     The program will access the location using the credentials you specify. Use this option if the current user account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

     Specify:

     - **User name**. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).

     - **Password**. The password for the account.

2. Click **OK**.

*According to the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that the user name and password can be intercepted by an eavesdropper using a packet sniffer.*

## 7.3.4 Volume selection

Select the volumes to mount and configure the mounting parameters for each of the selected volumes as follows:

1. Select the check box for each volume you need to mount.

2. Click on the selected volume to set its mounting parameters.

   - **Access mode** - choose the mode you want the volume to be mounted in:

     - **Read only** - enables exploring and opening files within the backup without committing any changes.

     - **Read/write** - with this mode, the program assumes that the backup content will be modified, and creates an incremental backup to capture the changes.

   - **Assign letter** (in Windows) - Acronis Backup & Recovery 11 will assign an unused letter to the mounted volume. If required, select another letter to assign from the drop-down list.

   - **Mount point** (in Linux) - specify the directory where you want the volume to be mounted.

3. If several volumes are selected for mounting, click on every volume to set its mounting parameters, described in the previous step.

4. Click **OK**.

## 7.3.5 Managing mounted images

Once a volume is mounted, you can browse files and folders contained in the backup using a file manager and copy the desired files to any destination. Thus, if you need to take out only a few files and folders from a volume backup, you do not have to perform the recovery procedure.

### Exploring images

Exploring mounted volumes lets you view and modify (if mounted in the read/write mode) the volume's content.

To explore a mounted volume select it in the table and click 🔍 **Explore**. The default file manager window opens, allowing the user to examine the mounted volume contents.

**Unmounting images**

Maintaining the mounted volumes takes considerable system resources. It is recommended that you unmount the volumes after the necessary operations are completed. If not unmounted manually, a volume will remain mounted until the operating system restarts.

To unmount an image, select it in the table and click 🖼️ **Unmount**.

To unmount all the mounted volumes, click 🖼️ **Unmount all**.

# 7.4    Operations available in vaults

By using vaults, you can easily access archives and backups and perform archive management operations.

***To perform operations with archives and backups***

1.  In the **Navigation** pane, select the vault whose archives you need to manage.
2.  In the vault view, select the **Archive view** tab. This tab displays all archives stored in the selected vault.
3.  Proceed as described in:
    -  Operations with archives (p. 155)
    -  Operations with backups (p. 156)

## 7.4.1    Operations with archives

***To perform any operation with an archive***

1.  In the **Navigation** pane, select the vault that contains archives.
2.  On the **Archive view** tab of the vault, select the archive. If the archive is protected with a password, you will be asked to provide it.
3.  Perform operations by clicking the corresponding buttons on the toolbar. These operations can also be accessed from the **'[Archive name]' actions** item of the main menu.

The following is a guideline for you to perform operations with archives stored in a vault.

| To | Do |
|---|---|
| Validate an archive | Click ✅ **Validate**.<br><br>The **Validation (p. 143)** page will be opened with the pre-selected archive as a source.<br><br>Validation of an archive will check all the archive's backups. |
| Export an archive | Click 🗄️ **Export**.<br><br>The **Export** (p. 146) page will be opened with the pre-selected archive as a source. The export of an archive creates a duplicate of the archive with all its backups in the location you specify. |
| Delete a single archive or multiple archives | 1.  Select one or more archives you want to delete.<br><br>2.  Click ❌ **Delete**. |

| | The program duplicates your selection in the **Backups deletion** (p. 157) window that has check boxes for each archive and each backup. Review the selection and make corrections if need be (select the check boxes for the desired archives), then confirm the deletion. |
|---|---|
| Delete all archives in the vault | Please be aware that if filters have been applied to the vaults list, you see only a part of the vault content. Be sure that the vault does not contain archives you need to retain before starting the operation. Click ![icon] **Delete all**. The program duplicates your selection in the new window that has check boxes for each archive and each backup. Review the selection and correct if need be, then confirm the deletion. |

## 7.4.2   Operations with backups

***To perform any operation with an archive***

1. In the **Navigation** pane, select the vault that contains archives.

2. On the **Archive view** tab of the vault, select the archive. Then, expand the archive and click the backup to select it. If the archive is protected with a password, you will be asked to provide it.

3. Perform operations by clicking the corresponding buttons on the toolbar. These operations can also be accessed from the **'[Backup name]' actions** item of the main menu.

The following is a guideline for you to perform operations with backups.

| To | Do |
|---|---|
| View backup content in a separate window | Click ![icon] **View content**. In the **Backup Content** window, examine the backup content. |
| Recover | Click ![icon] **Recover**. The **Recover data** (p. 97) page will be opened with the pre-selected backup as a source. |
| Convert a disk/volume backup to a virtual machine | Right-click the disk backup, then select **Convert to VM**. The **Recover data** (p. 97) page will be opened with the pre-selected backup as a source. Select the location and the type of new virtual machine and then proceed as with regular disk or volume recovery. |
| Validate a backup | Click ![icon] **Validate**. The **Validation (p. 143)** page will be opened with the pre-selected backup as a source. Validation of a file backup imitates recovering of all files from the backup to a dummy destination. Validation of a disk backup calculates a checksum for every data block saved in the backup. |
| Export a backup | Click ![icon] **Export**. The **Export** (p. 146) page will be opened with the pre-selected backup as a source. The export of a backup creates a new archive with a self-sufficient copy of the backup in the location you specify. |
| Convert a backup to full | Click ![icon] **Convert to full backup** to replace the incremental or differential backup with a full backup for the same point in time. See "Converting a backup to full" (p. 157) for more information. |

| Delete a single or multiple backups | Select one of the backups you want to delete, then click ❌ **Delete**.<br><br>The program duplicates your selection in the **Backups deletion** (p. 157) window that has check boxes for each archive and each backup. Review the selection and correct if need be (select the check boxes for the desired backups), then confirm the deletion. |
|---|---|
| Delete all archives and backups in the vault | Please be aware that if filters have been applied to the vaults list, you see only a part of the vault content. Be sure that the vault does not contain archives you need to retain before starting the operation.<br><br>Click ❌❌ **Delete all**.<br><br>The program duplicates your selection in the **Backups deletion** (p. 157) window that has check boxes for each archive and each backup. Review the selection and correct if need be, then confirm the deletion. |

## 7.4.3   Converting a backup to full

When the chain of incremental backups in an archive becomes long, conversion of an incremental backup to a full one increases the reliability of your archive. You may also want to convert a differential backup if there are incremental backups that depend on it.

During the conversion, the selected incremental or differential backup is replaced with a full backup for the same point in time. The previous backups in the chain are not changed. All subsequent incremental and differential backups up to the nearest full backup are also updated. The new backup versions are created first and only after that are the old ones deleted. Therefore, the location must have enough space to temporarily store both the old and the new versions.

Conversion does not create a copy of a backup. To obtain a self-sufficient copy of the backup on a flash drive or removable media, use the export (p. 146) operation.

**Example**

You have the following backup chain in your archive:

`F1 I2 I3 I4 D5 I6 I7 I8 F9 I10 I11 D12 F13`

Here **F** means full backup, **I** -incremental, **D** - differential.

You convert to full the **I4** backup. The **I4, D5, I6, I7, I8** backups will be updated, while **I10 I11 D12** will remain unchanged, because they depend on **F9**.

*Limitation: The **Convert to full backup** operation is not allowed for backups on tapes and CD/DVD.*

## 7.4.4   Deleting archives and backups

The **Backups deletion** window displays the same tab as for the vaults view, but with check boxes for each archive and backup. The archive or backup you have chosen to delete has the check mark. Review the archive or backup that you have selected to delete. If you need to delete other archives and backups select the respective check boxes, then click **Delete selected** and confirm the deletion.

**What happens if I delete a backup that is a base of an incremental or differential backup?**

To preserve archive consistency, the program will consolidate the two backups. For example, you delete a full backup but retain the next incremental one. The backups will be combined into a single

full backup which will be dated the incremental backup date. When you delete an incremental or differential backup from the middle of the chain, the resulting backup type will be incremental.

Please be aware that consolidation is just a method of deletion but not an alternative to deletion. The resulting backup will not contain data that was present in the deleted backup and was absent from the retained incremental or differential backup.

There should be enough space in the vault for temporary files created during consolidation. Backups resulting from consolidation always have maximum compression.

# 8 Bootable media

**Bootable media**

Bootable media is physical media (CD, DVD, USB drive or other media supported by a machine BIOS as a boot device) that boots on any PC-compatible machine and enables you to run Acronis Backup & Recovery 11 Agent either in a Linux-based environment or Windows Preinstallation Environment (WinPE), without the help of an operating system. Bootable media is most often used to:

- recover an operating system that cannot start
- access and back up the data that has survived in a corrupted system
- deploy an operating system on bare metal
- create basic or dynamic volumes on bare metal
- back up sector-by-sector a disk with an unsupported file system
- back up offline any data that cannot be backed up online because of restricted access, being permanently locked by the running applications or for any other reason.

A machine can be booted into the above environments either with physical media, or using the network boot from Acronis PXE Server, Windows Deployment Services (WDS) or Remote Installation Services (RIS). These servers with uploaded bootable components can be thought of as a kind of bootable media too. You can create bootable media or configure the PXE server or WDS/RIS using the same wizard.

## Linux-based bootable media

Linux-based media contains Acronis Backup & Recovery 11 Bootable Agent based on Linux kernel. The agent can boot and perform operations on any PC-compatible hardware, including bare metal and machines with corrupted or non-supported file systems. The operations can be configured and controlled either locally or remotely using the management console.

## PE-based bootable media

PE-based bootable media contains a minimal Windows system called Windows Preinstallation Environment (WinPE) and Acronis Plug-in for WinPE, that is, a modification of Acronis Backup & Recovery 11 Agent that can run in the preinstallation environment.

WinPE proved to be the most convenient bootable solution in large environments with heterogeneous hardware.

**Advantages:**

- Using Acronis Backup & Recovery 11 in Windows Preinstallation Environment provides more functionality than using Linux-based bootable media. Having booted PC-compatible hardware into WinPE, you can use not only Acronis Backup & Recovery 11 Agent, but also PE commands and scripts and other plug-ins you've added to the PE.
- PE-based bootable media helps overcome some Linux-related bootable media issues such as support for certain RAID controllers or certain levels of RAID arrays only. Media based on PE 2.x, that is, Windows Vista or Windows Server 2008 kernel, allows for dynamic loading of the necessary device drivers.

**Limitation:**

PE-based bootable media does not support UEFI.

# 8.1 How to create bootable media

To enable creating physical media, the machine must have a CD/DVD recording drive or allow a flash drive to be attached. To enable PXE or WDS/RIS configuration, the machine must have a network connection. Bootable Media Builder can also create an ISO image of a bootable disk to burn it later on a blank disk.

## Linux-based bootable media

Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component.

Select the way volumes and network resources will be handled—called the media style:

- A media with Linux-style volume handling displays the volumes as, for example, hda1 and sdb2. It tries to reconstruct MD devices and logical (LVM) volumes before starting a recovery.
- A media with Windows-style volume handling displays the volumes as, for example, C: and D:. It provides access to dynamic (LDM) volumes.

The wizard will guide you through the necessary operations. Please refer to Linux-based bootable media (p. 161) for details.

## PE-based bootable media

Acronis Plug-in for WinPE can be added to WinPE distributions based on any of the following kernels:

- Windows XP Professional with Service Pack 2 (PE 1.5)
- Windows Server 2003 with Service Pack 1 (PE 1.6)
- Windows Vista (PE 2.0)
- Windows Vista SP1 and Windows Server 2008 (PE 2.1)
- Windows 7 (PE 3.0)

If you already have media with PE1.x distribution, unpack the media ISO to a local folder and start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component. The wizard will guide you through the necessary operations. Please refer to Adding the Acronis Plug-in to WinPE 1.x (p. 165) for details.

To be able to create or modify PE 2.x or 3.0 images, install Bootable Media Builder on a machine where Windows Automated Installation Kit (AIK) is installed. The further operations are described in the Adding the Acronis Plug-in to WinPE 2.x or 3.0 (p. 166) section.

If you do not have a machine with WAIK, prepare as follows:

1. Download and install Windows Automated Installation Kit (WAIK).

   Automated Installation Kit (AIK) for Windows Vista (PE 2.0):

   http://www.microsoft.com/Downloads/details.aspx?familyid=C7D4BC6D-15F3-4284-9123-67983 0D629F2&displaylang=en

   Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008 (PE 2.1):

   http://www.microsoft.com/downloads/details.aspx?FamilyID=94bb6e34-d890-4932-81a5-5b50c 657de08&DisplayLang=en

   Automated Installation Kit (AIK) for Windows 7 (PE 3.0):

http://www.microsoft.com/downloads/details.aspx?familyid=696DD665-9F76-4177-A811-39C26 D3B3B34&displaylang=en

You can find system requirements for installation by following the above links.

2. [optional] Burn the WAIK to DVD or copy to a flash drive.

3. Install the Microsoft .NET Framework v.2.0 from this kit (NETFXx86 or NETFXx64, depending on your hardware.)

4. Install Microsoft Core XML (MSXML) 5.0 or 6.0 Parser from this kit.

5. Install Windows AIK from this kit.

6. Install Bootable Media Builder on the same machine.

It is recommended that you familiarize yourself with the help documentation supplied with Windows AIK. To access the documentation, select **Microsoft Windows AIK -> Documentation** from the start menu.

### Using Bart PE

You can create a Bart PE image with Acronis Plug-in using the Bart PE Builder. Please refer to Building Bart PE with Acronis Plug-in from Windows distribution (p. 167) for details.

## 8.1.1   Linux-based bootable media

***When using the media builder, you have to specify:***

1. [optional] The parameters of the Linux kernel. Separate multiple parameters with spaces.

    For example, to be able to select a display mode for the bootable agent each time the media starts, type: **vga=ask**

    For a list of parameters, see Kernel parameters (p. 162).

2. The Acronis bootable components to be placed on the media.

    Universal Restore will be enabled if Acronis Backup & Recovery 11 Universal Restore is installed on the machine where the media is created.

3. [optional] The timeout interval for the boot menu plus the component that will automatically start on timeout.

    - If not configured, the Acronis loader waits for someone to select whether to boot the operating system (if present) or the Acronis component.

    - If you set, say, **10 sec.** for the bootable agent, the agent will launch 10 seconds after the menu is displayed. This enables unattended onsite operation when booting from a PXE server or WDS/RIS.

4. [optional] Remote logon settings:

    - user name and password to be entered on the console side at connection to the agent. If you leave these fields empty, the connection will be enabled on typing any symbols in the prompt window.

5. [optional] Network settings (p. 163):

    - TCP/IP settings to be assigned to the machine network adapters.

6. [optional] Network port (p. 164):

    - the TCP port that the bootable agent listens for incoming connection.

7. The type of media to create. You can:

    - create CD, DVD or other bootable media such as removable USB flash drives if the hardware BIOS allows for boot from such media

- build an ISO image of a bootable disc to burn it later on a blank disc
- upload the selected components to Acronis PXE Server
- upload the selected components to a WDS/RIS.

8. [optional] Windows system drivers to be used by Acronis Universal Restore (p. 164). This window appears only if the Acronis Universal Restore add-on is installed and a media other than PXE or WDS/RIS is selected.

9. Path to the media ISO file or the name or IP and credentials for PXE or WDS/RIS.

## 8.1.1.1    Kernel parameters

This window lets you specify one or more parameters of the Linux kernel. They will be automatically applied when the bootable media starts.

These parameters are typically used when experiencing problems while working with the bootable media. Normally, you can leave this field empty.

You also can specify any of these parameters by pressing F11 while in the boot menu.

**Parameters**

When specifying multiple parameters, separate them with spaces.

**acpi=off**

Disables Advanced Configuration and Power Interface (ACPI). You may want to use this parameter when experiencing problems with a particular hardware configuration.

**noapic**

Disables Advanced Programmable Interrupt Controller (APIC). You may want to use this parameter when experiencing problems with a particular hardware configuration.

**vga=ask**

Prompts for the video mode to be used by the bootable media's graphical user interface. Without the **vga** parameter, the video mode is detected automatically.

**vga=**_mode_number_

Specifies the video mode to be used by the bootable media's graphical user interface. The mode number is given by _mode_number_ in the hexadecimal format—for example: **vga=0x318**

Screen resolution and the number of colors corresponding to a mode number may be different on different machines. We recommend using the **vga=ask** parameter first to choose a value for _mode_number_.

**quiet**

Disables displaying of startup messages when the Linux kernel is loading, and starts the management console after the kernel is loaded.

This parameter is implicitly specified when creating the bootable media, but you can remove this parameter while in the boot menu.

Without this parameter, all startup messages will be displayed, followed by a command prompt. To start the management console from the command prompt, run the command: **/bin/product**

**nousb**

Disables loading of the USB (Universal Serial Bus) subsystem.

**nousb2**

Disables USB 2.0 support. USB 1.1 devices still work with this parameter. This parameter allows you to use some USB drives in the USB 1.1 mode if they do not work in the USB 2.0 mode.

**nodma**

Disables direct memory access (DMA) for all IDE hard disk drives. Prevents the kernel from freezing on some hardware.

**nofw**

Disables the FireWire (IEEE1394) interface support.

**nopcmcia**

Disables detection of PCMCIA hardware.

**nomouse**

Disables mouse support.

*module_name***=off**

Disables the module whose name is given by *module_name*. For example, to disable the use of the SATA module, specify: **sata_sis=off**

**pci=bios**

Forces the use of PCI BIOS instead of accessing the hardware device directly. You may want to use this parameter if the machine has a non-standard PCI host bridge.

**pci=nobios**

Disables the use of PCI BIOS; only direct hardware access methods will be allowed. You may want to use this parameter when the bootable media fails to start, which may be caused by the BIOS.

**pci=biosirq**

Uses PCI BIOS calls to get the interrupt routing table. You may want to use this parameter if the kernel is unable to allocate interrupt requests (IRQs) or discover secondary PCI buses on the motherboard.

These calls might not work properly on some machines. But this may be the only way to get the interrupt routing table.

## 8.1.1.2 Network settings

While creating Acronis bootable media, you have an option to pre-configure network connections that will be used by the bootable agent. The following parameters can be pre-configured:

- IP address
- Subnet mask
- Gateway
- DNS server
- WINS server.

Once the bootable agent starts on a machine, the configuration is applied to the machine's network interface card (NIC.) If the settings have not been pre-configured, the agent uses DHCP auto configuration. You also have the ability to configure the network settings manually when the bootable agent is running on the machine.

**Pre-configuring multiple network connections**

You can pre-configure TCP/IP settings for up to ten network interface cards. To ensure that each NIC will be assigned the appropriate settings, create the media on the server for which the media is customized. When you select an existing NIC in the wizard window, its settings are selected for saving on the media. The MAC address of each existing NIC is also saved on the media.

You can change the settings, except for the MAC address; or configure the settings for a non-existent NIC, if need be.

Once the bootable agent starts on the server, it retrieves the list of available NICs. This list is sorted by the slots the NICs occupy: the closest to the processor on top.

The bootable agent assigns each known NIC the appropriate settings, identifying the NICs by their MAC addresses. After the NICs with known MAC addresses are configured, the remaining NICs are assigned the settings that you have made for non-existent NICs, starting from the upper non-assigned NIC.

You can customize bootable media for any machine, and not only for the machine where the media is created. To do so, configure the NICs according to their slot order on that machine: NIC1 occupies the slot closest to the processor, NIC2 is in the next slot and so on. When the bootable agent starts on that machine, it will find no NICs with known MAC addresses and will configure the NICs in the same order as you did.

**Example**

The bootable agent could use one of the network adapters for communication with the management console through the production network. Automatic configuration could be done for this connection. Sizeable data for recovery could be transferred through the second NIC, included in the dedicated backup network by means of static TCP/IP settings.

## 8.1.1.3    Network port

While creating bootable media, you have an option to pre-configure the network port that the bootable agent listens for incoming connection. The choice is available between:

- the default port
- the currently used port
- the new port (enter the port number).

If the port has not been pre-configured, the agent uses the default port number (9876.) This port is also used as default by the Acronis Backup & Recovery 11 Management Console.

## 8.1.1.4    Drivers for Universal Restore

While creating bootable media, you have an option to add Windows drivers to the media. The drivers will be used by Universal Restore when recovering Windows on a machine with a dissimilar processor, different motherboard or different mass storage device than in the backed up system.

You will be able to configure the Universal Restore:

- to search the media for the drivers that best fit the target hardware

- to get the mass-storage drivers that you explicitly specify from the media. This is necessary when the target hardware has a specific mass storage controller (such as a SCSI, RAID, or Fiber Channel adapter) for the hard disk.

For more information please refer to Universal Restore.

The drivers will be placed in the visible Drivers folder on the bootable media. The drivers are not loaded into the target machine RAM, therefore, the media must stay inserted or connected throughout the Universal Restore operation.

Adding drivers to bootable media is available on the condition that:

1. The Acronis Backup & Recovery 11 Universal Restore add-on is installed on the machine where the bootable media is created AND
2. You are creating a removable media or its ISO or detachable media, such as a flash drive. Drivers cannot be uploaded on a PXE server or WDS/RIS.

The drivers can be added to the list only in groups, by adding the INF files or folders containing such files. Selecting individual drivers from the INF files is not possible, but the media builder shows the file content for your information.

***To add drivers:***

1. Click **Add** and browse to the INF file or a folder that contains INF files.
2. Select the INF file or the folder.
3. Click **OK**.

The drivers can be removed from the list only in groups, by removing INF files.

***To remove drivers:***

1. Select the INF file.
2. Click **Remove**.

## 8.1.2   Adding the Acronis Plug-in to WinPE 1.x

Acronis Plug-in for WinPE can be added to:

- Windows PE 2004 (1.5) (Windows XP Professional with Service Pack 2)
- Windows PE 2005 (1.6) (Windows Server 2003 with Service Pack 1).

***To add Acronis Plug-in to WinPE 1.x:***

1. Unpack all files of your WinPE 1.x ISO to a separate folder on the hard disk.
2. Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component.
3. Select **Bootable media type: Windows PE**.
   - Select **Use WinPE files located in the folder I specify**
4. Specify path to the folder where the WinPE files are located.
5. Specify network settings (p. 163) for the machine network adapters or choose DHCP auto configuration.
6. Specify the full path to the resulting ISO file including the file name.
7. Check your settings in the summary screen and click **Proceed**.
8. Burn the .ISO to CD or DVD using a third-party tool or copy to a flash drive.

Once a machine boots into the WinPE, Acronis Backup & Recovery 11 starts automatically.

## 8.1.3   Adding the Acronis Plug-in to WinPE 2.x or 3.0

Bootable Media Builder provides three methods of integrating Acronis Backup & Recovery 11 with WinPE 2.x or 3.0:

- Adding the Acronis Plug-in to the existing PE ISO. This comes in handy when you have to add the plug-in to the previously configured PE ISO that is already in use.
- Creating the PE ISO with the plug-in from scratch.
- Adding the Acronis Plug-in to a WIM file for any future purpose (manual ISO building, adding other tools to the image and so on).

To be able to perform any of the above operations, install Bootable Media Builder on a machine where Windows Automated Installation Kit (WAIK) is installed. If you do not have such machine, prepare as described in How to create bootable media (p. 160).

*Bootable Media Builder supports only x86 WinPE 2.x or 3.0. These WinPE distributions can also work on x64 hardware.*

*PE images based on Win PE 2.x or 3.0 require at least 640 MB RAM to work.*

### Adding Acronis Plug-in to WinPE 2.x or 3.0 ISO

***To add Acronis Plug-in to WinPE 2.x or 3.0 ISO:***

1.  When adding the plug-in to the existing Win PE ISO, unpack all files of your Win PE ISO to a separate folder on the hard disk.
2.  Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component.
3.  Select **Bootable media type: Windows PE**.

    When creating a new PE ISO:

    - Select **Create Windows PE 2.x or 3.0 automatically**
    - The software runs the appropriate script and proceeds to the next window.

    When adding the plug-in to the existing PE ISO:

    - Select **Use WinPE files located in the folder I specify**
    - Specify path to the folder where the WinPE files are located.

4.  Specify network settings (p. 163) for the machine network adapters or choose DHCP auto configuration.
5.  [optional] Specify the Windows drivers to be added to Windows PE. Once you boot a machine into Windows PE, the drivers can help you access the device where the backup archive is located. Also, you will be able to point to these drivers when configuring Universal Restore. WinPE itself needs 32-bit drivers. For using Universal Restore on a 64-bit system, add 64-bit drivers as well.

    Click **Add** and specify the path to the necessary *.inf file for a corresponding SCSI, RAID, SATA controller, network adapter, tape drive or other device. You will have to repeat this procedure for each driver you want to be included in the resulting WinPE boot media.

6.  Choose whether you want to create ISO or WIM image or upload the media on a server (Acronis PXE Server, WDS or RIS).
7.  Specify the full path to the resulting image file including the file name, or specify the server and provide the user name and password to access it.
8.  Check your settings in the summary screen and click **Proceed**.
9.  Burn the .ISO to CD or DVD using a third-party tool or copy to a flash drive.

Once a machine boots into WinPE, Acronis Backup & Recovery 11 starts automatically.

***To create a PE image (ISO file) from the resulting WIM file:***

- replace the default boot.wim file in your Windows PE folder with the newly created WIM file. For the above example, type:

```
copy c:\AcronisMedia.wim c:\winpe_x86\ISO\sources\boot.wim
```

- use the **Oscdimg** tool. For the above example, type:

```
oscdimg -n –bc:\winpe_x86\etfsboot.com c:\winpe_x86\ISO
c:\winpe_x86\winpe_x86.iso
```

For more information on customizing Windows PE, see the Windows Preinstallation Environment User's Guide (Winpe.chm).

## 8.1.4 Building Bart PE with Acronis Plug-in from Windows distribution

1. Get the Bart PE builder.
2. Install Bootable Media Builder from the Acronis Backup & Recovery 11 setup file.
3. Change the current folder to the folder where the Acronis Plug-in for WinPE is installed—by default: C:\Program Files\Acronis\Bootable Components\WinPE.

   If the plug-in is installed in a folder other than the default folder, change the path accordingly (check the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\Bootable Components\Settings\WinPE for the plug-in location).
4. Unpack the WinPE.zip file to the current folder.
5. Run the following command:

   ```
   make_plugin.vbs
   ```

   As a result, a folder named **acronis** will be created in the current directory.
6. Copy the **acronis** folder to the %BartPE folder%\plugins folder.
7. Insert your Windows distribution CD if you do not have a copy of Windows installation files on the HDD.
8. Start the Bart PE builder.
9. Specify the path to the Windows installation files or Windows distribution CD.
10. Click **Plugins** and check whether the Acronis Backup & Recovery 11 plug-in is enabled. Enable if disabled.
11. Specify the output folder and the full path to the resulting ISO file including the file name or the media to create.
12. Build the Bart PE.
13. Burn the ISO to CD or DVD (if this has not been done yet) or copy to a flash drive.

Once the machine boots into the Bart PE and you configure the network connection, select **Go -> System -> Storage -> Acronis Backup & Recovery 11** to start.

## 8.2 Connecting to a machine booted from media

Once a machine boots from bootable media, the machine terminal displays a startup window with the IP address(es) obtained from DHCP or set according to the pre-configured values.

**Remote connection**

To connect to the machine remotely, select **Connect -> Manage a remote machine** in the console menu and specify one of the machine's IP addresses. Provide the user name and password if these have been configured when creating the bootable media.

**Local connection**

Acronis Backup & Recovery 11 Management Console is always present on the bootable media. Anyone who has physical access to the machine terminal can run the console and connect. Just click **Run management console** in the bootable agent startup window.

# 8.3    Working under bootable media

Operations on a machine booted with bootable media are very similar to backup and recovery under the operating system. The difference is as follows:

1.  Disk letters seen under Windows-style bootable media might differ from the way Windows identifies drives. For example, the D: drive under the rescue utility might correspond to the E: drive in Windows.

    *Be careful! To be on the safe side, it is advisable to assign unique names to the volumes.*

2.  The Linux-style bootable media shows local disks and volumes as unmounted (sda1, sda2...).
3.  Backups created using bootable media have simplified file names (p. 48). Standard names are assigned to the backups only if these are added to an existing archive with standard file naming, or if the destination does not support simplified file names.
4.  The Linux-style bootable media cannot write a backup to an NTFS-formatted volume. Switch to the Windows style if you need to do so.
5.  You can switch the bootable media between the Windows style and the Linux style by selecting **Tools** > **Change volume representation**.
6.  There is no **Navigation** tree in the media GUI. Use the **Navigation** menu item to navigate between views.
7.  Tasks cannot be scheduled; in fact, tasks are not created at all. If you need to repeat the operation, configure it from scratch.
8.  The log lifetime is limited to the current session. You can save the entire log or the filtered log entries to a file.
9.  Centralized vaults are not displayed in the folder tree of the **Archive** window.

    To access a managed vault, type the following string in the **Path** field:

    **bsp://node_address/vault_name/**

    To access an unmanaged centralized vault, type the full path to the vault's folder.

    After entering access credentials, you will see a list of archives located in the vault.

# 8.3.1    Setting up a display mode

For a machine booted from media, a display video mode is detected automatically based on the hardware configuration (monitor and graphics card specifications). If, for some reason, the video mode is detected incorrectly, do the following:

1.  In the boot menu, press F11.
2.  Add to the command prompt the following command: **vga=ask**, and then proceed with booting.

3. From the list of supported video modes, choose the appropriate one by typing its number (for example, **318**), and then press ENTER.

If you do not wish to follow this procedure every time you boot from media on a given hardware configuration, re-create the bootable media with the appropriate mode number (in our example, **vga=0x318**) typed in the **Kernel parameters** window (see the Bootable Media Builder (p. 161) section for details).

## 8.3.2 Configuring iSCSI and NDAS devices

This section describes how to configure Internet Small Computer System Interface (iSCSI) devices and Network Direct Attached Storage (NDAS) devices when working under bootable media.

These devices are connected to the machine through a network interface and appear as if they were locally-attached devices. On the network, an iSCSI device is identified by its IP address, and an NDAS device is identified by its device ID.

An iSCSI device is sometimes called an iSCSI target. A hardware or software component that provides interaction between the machine and the iSCSI target is called the iSCSI initiator. The name of the iSCSI initiator is usually defined by an administrator of the server that hosts the device.

### *To add an iSCSI device*

1. In a bootable media (Linux-based or PE-based), run the management console.
2. Click **Configure iSCSI/NDAS devices** (in a Linux-based media) or **Run the iSCSI Setup** (in a PE-based media).
3. Specify the IP address and port of the iSCSI device's host, and the name of the iSCSI initiator.
4. If the host requires authentication, specify the user name and password for it.
5. Click **OK**.
6. Select the iSCSI device from the list, and then click **Connect**.
7. If prompted, specify the user name and password to access the iSCSI device.

### *To add an NDAS device*

1. In a Linux-based bootable media, run the management console.
2. Click **Configure iSCSI/NDAS devices**.
3. In **NDAS devices**, click **Add device**.
4. Specify the 20-character device ID.
5. If you want to allow writing data onto the device, specify the five-character write key. Without this key, the device will be available in the read-only mode.
6. Click **OK**.

## 8.4 List of commands and utilities available in Linux-based bootable media

Linux-based bootable media contains the following commands and command line utilities, which you can use when running a command shell. To start the command shell, press CTRL+ALT+F2 while in the bootable media's management console.

**Acronis command-line utilities**

- `acrocmd`

- acronis
- asamba
- lash

## Linux commands and utilities

| | | |
|---|---|---|
| busybox | ifconfig | rm |
| cat | init | rmmod |
| cdrecord | insmod | route |
| chmod | iscsiadm | scp |
| chown | kill | scsi_id |
| chroot | kpartx | sed |
| cp | ln | sg_map26 |
| dd | ls | sh |
| df | lspci | sleep |
| dmesg | lvm | ssh |
| dmraid | mdadm | sshd |
| e2fsck | mkdir | strace |
| e2label | mke2fs | swapoff |
| echo | mknod | swapon |
| egrep | mkswap | sysinfo |
| fdisk | more | tar |
| fsck | mount | tune2fs |
| fxload | mtx | udev |
| gawk | mv | udevinfo |
| gpm | pccardctl | udevstart |
| grep | ping | umount |
| growisofs | pktsetup | uuidgen |
| grub | poweroff | vconfig |
| gunzip | ps | vi |
| halt | raidautorun | zcat |
| hexdump | readcd | |
| hotplug | reboot | |

# 8.5 Acronis Startup Recovery Manager

Acronis Startup Recovery Manager is a modification of the bootable agent (p. 229), residing on the system disk in Windows, or on the /boot partition in Linux and configured to start at boot time on pressing F11. It eliminates the need for a separate media or network connection to start the bootable rescue utility.

Acronis Startup Recovery Manager is especially useful for mobile users. If a failure occurs, reboot the machine, wait for the prompt "Press F11 for Acronis Startup Recovery Manager…" to appear, and hit F11. The program will start and you can perform recovery.

You can also back up using Acronis Startup Recovery Manager, while on the move.

On machines with the GRUB boot loader installed, you select the Acronis Startup Recovery Manager from the boot menu instead of pressing F11.

## Activate

Activation enables the boot time prompt "Press F11 for Acronis Startup Recovery Manager…" (if you do not have the GRUB boot loader) or adds the "Acronis Startup Recovery Manager" item to GRUB's menu (if you have GRUB).

*The system disk (or, the /boot partition in Linux) should have at least 100 MB of free space to activate Acronis Startup Recovery Manager.*

Unless you use the GRUB boot loader and it is installed in the Master Boot Record (MBR), Acronis Startup Recovery Manager activation overwrites the MBR with its own boot code. Thus, you may need to reactivate third-party boot loaders if they are installed.

Under Linux, when using a boot loader other than GRUB (such as LILO), consider installing it to a Linux root (or boot) partition boot record instead of the MBR before activating Acronis Startup Recovery Manager. Otherwise, reconfigure the boot loader manually after the activation.

## Do not activate

Disables boot time prompt "Press F11 for Acronis Startup Recovery Manager…" (or, the menu item in GRUB). If Acronis Startup Recovery Manager is not activated, you will need one of the following to recover the system when it fails to boot:

- boot the machine from a separate bootable rescue media
- use network boot from Acronis PXE Server or Microsoft Remote Installation Services (RIS).

# 9 Disk management

Acronis Disk Director Lite is a tool for preparing a machine disk/volume configuration for recovering the volume images saved by the Acronis Backup & Recovery 11 software.

Sometimes after the volume has been backed up and its image placed into a safe storage, the machine disk configuration might change due to a HDD replacement or hardware loss. In such case with the help of Acronis Disk Director Lite, the user has the possibility to recreate the necessary disk configuration so that the volume image can be recovered exactly "as it was" or with any alteration of the disk or volume structure the user might consider necessary.

*All operations on disks and volumes involve a certain risk of data damage. Operations on system, bootable or data volumes must be carried out very carefully to avoid potential problems with the booting process or hard disk data storage.*

*Operations with hard disks and volumes take a certain amount of time, and any power loss, unintentional turning off of the machine or accidental pressing of the Reset button during the procedure could result in volume damage and data loss.*

*All operations on volumes of dynamic disks in Windows XP and Windows 2000 require Acronis Managed Machine Service to be run under an account with administrator's rights.*

Please take all necessary precautions to avoid possible data loss.

## 9.1 Supported file systems

Acronis Disk Director Lite supports the following file systems:

- FAT 16/32
- NTFS

If it is necessary to perform an operation on a volume with a different file system, use the full version of Acronis Disk Director. It provides more tools and utilities to manage disks and volumes with the following file systems:

- FAT 16/32
- NTFS
- Ext2
- Ext3
- HFS+
- HFSX
- ReiserFS
- JFS
- Linux SWAP

## 9.2 Basic precautions

To avoid any possible disk and volume structure damage or data loss, please take all necessary precautions and follow these simple rules:

1. Back up the disk on which volumes will be created or managed. Having your most important data backed up to another hard disk, network share or removable media will allow you to work on disk volumes being reassured that your data is safe.

2. Test your disk to make sure it is fully functional and does not contain bad sectors or file system errors.

3. Do not perform any disk/volume operations while running other software that has low-level disk access. Close these programs before running Acronis Disk Director Lite.

With these simple precautions, you will protect yourself against accidental data loss.

# 9.3 Running Acronis Disk Director Lite

You can run Acronis Disk Director Lite under Windows or start it from a bootable media.

**Running Acronis Disk Director Lite under Windows**

If you run Acronis Backup & Recovery 11 Management Console, and connect it to a managed machine, the **Disk management** view will be available in the **Navigation** tree of the console, with which you can start Acronis Disk Director Lite.

**Running Acronis Disk Director Lite from a bootable media**

You can run Acronis Disk Director Lite on a bare metal, on a machine that cannot boot or on a non-Windows machine. To do so, boot the machine from a bootable media (p. 229) created with the Acronis Bootable Media Builder; run the management console and then click **Disk Management**.

# 9.4 Choosing the operating system for disk management

On a machine with two or more operating systems, representation of disks and volumes depends on which operating system is currently running.

A volume may have a different letter in different Windows operating systems. For example, volume E: might appear as D: or L: when you boot another Windows operating system installed on the same machine. (It is also possible that this volume will have the same letter E: under any Windows OS installed on the machine.)

A dynamic disk created in one Windows operating system is considered as a **Foreign Disk** in another Windows operating system or might be unsupported by this operating system.

When you need to perform a disk management operation on such machine, it is necessary to specify for which operating system the disk layout will be displayed and the disk management operation will be performed.

The name of the currently selected operating system is shown on the console toolbar after "**The current disk layout is for:**". Click the OS name to select another operating system in the **Operating System Selection** window. Under bootable media, this window appears after clicking **Disk management**. The disk layout will be displayed according to the operating system you select.

# 9.5 "Disk management" view

Acronis Disk Director Lite is controlled through the **Disk management** view of the console.

The top part of the view contains a disks and volumes table enabling data sorting and columns customization and toolbar. The table presents the numbers of the disks, as well as assigned letter, label, type, capacity, free space size, used space size, file system, and status for each volume. The toolbar comprises of icons to launch the **Undo**, **Redo** and **Commit** actions intended for pending operations (p. 187).

The graphic panel at the bottom of the view also graphically depicts all the disks and their volumes as rectangles with basic data on them (label, letter, size, status, type and file system).

Both parts of the view also depict all unallocated disk space that can be used in volume creation.

**Starting the operations**

Any operation can be launched:

▪ From the volume or disk context menu (both in the table and the graphic panel)

▪ From the **Disk management** menu of the console

▪ From the **Operations** bar on the **Actions and Tools** pane

*Note that the list of available operations in the context menu, the **Disk management** menu, and the **Operations** bar depends on the selected volume or disk type. The same is true for unallocated space as well.*

**Displaying operation results**

The results of any disk or volume operation, you have just planned, are immediately displayed in the **Disk management** view of the console. For example, if you create a volume, it will be immediately shown in the table, as well as in graphical form at the bottom of the view. Any volume changes, including changing the volume letter or label, are also immediately displayed in the view.

# 9.6    Disk operations

Acronis Disk Director Lite includes the following operations that can be performed on disks:

▪ Disk Initialization (p. 174) - initializes the new hardware added to the system

▪ Basic disk cloning (p. 175) - transfers complete data from the source basic MBR disk to the target

▪ Disk conversion: MBR to GPT    (p. 177) - converts an MBR partition table to GPT

▪ Disk conversion: GPT to MBR    (p. 178) - converts a GPT partition table to MBR

▪ Disk conversion: Basic to Dynamic    (p. 178) - converts a basic disk to dynamic

▪ Disk conversion: Dynamic to Basic    (p. 179) - converts a dynamic disk to basic

The full version of Acronis Disk Director will provide more tools and utilities for working with disks.

*Acronis Disk Director Lite must obtain exclusive access to the target disk. This means no other disk management utilities (like Windows Disk Management utility) can access it at that time. If you receive a message stating that the disk cannot be blocked, close the disk management applications that use this disk and start again. If you cannot determine which applications use the disk, close them all.*

## 9.6.1    Disk initialization

If you add any new disk to your machine, Acronis Disk Director Lite will notice the configuration change and scan the added disk to include it to the disk and volume list. If the disk is still not initialized or, possibly, has a file structure unknown to the machine system, that means that no programs can be installed on it and you will not be able to store any files there.

Acronis Disk Director Lite will detect that the disk is unusable by the system and needs to be initialized. The **Disk management** view will show the newly detected hardware as a gray block with a grayed icon, thus indicating that the disk is unusable by the system.

***If you need to initialize a disk:***

1. Select a disk to initialize.

2. Right-click on the selected volume, and then click **Initialize** in the context menu. You will be forwarded to the **Disk Initialization** window, that will provide the basic hardware details such as the disk's number, capacity and state to aid you in the choice of your possible action.

3. In the window, you will be able to set the disk partitioning scheme (MBR or GPT) and the disk type (basic or dynamic). The new disk state will be graphically represented in the **Disk Management** view of the console immediately.

4. By clicking **OK**, you'll add a pending operation of the disk initialization.

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

After the initialization, all the disk space remains unallocated and so still impossible to be used for program installation or file storage. To be able to use it, proceed normally to the **Create volume** operation.

*If you decide to change the disk settings it can be done later using the standard Acronis Disk Director Lite disk tools.*

## 9.6.2    Basic disk cloning

Sometimes it is necessary to transfer all the disk data onto a new disk. It can be a case of expanding the system volume, starting a new system layout or disk evacuation due to a hardware fault. In any case, the reason for the **Clone basic disk** operation can be summed up as the necessity to transfer all the source disk data to a target disk exactly as it is.

Acronis Disk Director Lite allows the operation to be carried out to basic MBR disks only.

To plan the **Clone basic disk** operation:

1. Select a disk you want to clone.

2. Select a disk as target for the cloning operation.

3. Select a cloning method and specify advanced options.

The new volume structure will be graphically represented in the **Disk management** view immediately.

*It is advisable that you deactivate Acronis Startup Recovery Manager (p. 227) (ASRM), if it is active, before cloning a system disk. Otherwise the cloned operating system might not boot. You can activate the ASRM again after the cloning is completed. If deactivation is not possible, choose the **As is** method to clone the disk.*

### 9.6.2.1    Selecting source and target disks

The program displays a list of partitioned disks and asks the user to select the source disk, from which data will be transferred to another disk.

The next step is selection of a disk as target for the cloning operation. The program enables the user to select a disk if its size will be sufficient to hold all the data from the source disk without any loss.

If there is some data on the disk that was chosen as the target, the user will receive a warning: "**The selected target disk is not empty. The data on its volumes will be overwritten.**", meaning that all the data currently located on the chosen target disk will be lost irrevocably.

## 9.6.2.2    Cloning method and advanced options

The **Clone basic disk** operation usually means that the information from the source disk is transferred to the target "**As is**". So, if the destination disk is the same size and even if it is larger, it is possible to transfer all the information there exactly as it is stored at the source.

But with the wide range of available hardware it is normal that the target disk would differ in size from the source. If the destination is larger, then it would be advisable to resize the source disk volumes to avoid leaving unallocated space on the target disk by selecting the **Proportionally resize volumes** option. The option to **Clone basic disk** "as is" remains, but the default method of cloning will be carried out with proportional enlargement of all the **source** disk volumes      so that no unallocated space remains on the **target** disk .

If the destination is smaller, then the **As is** option of cloning will be unavailable and proportional resizing of the **source** disk volumes will be mandatory. The program analyzes the **target** disk to establish whether its size will be sufficient to hold all the data from the **source** disk without any loss. If such transfer with proportional resizing of the **source** disk volumes is possible, but without any data loss , then the user will be allowed to proceed. If due to the size limitations safe transfer of all the **source** disk data to the **target** disk is impossible even with the proportional resizing of the volumes, then the **Clone basic disk** operation will be impossible and the user will not be able to continue.

If you are about to clone a disk comprising of a **system volume**, pay attention to the **Advanced options**.

By clicking **Finish**, you'll add the pending operation of the disk cloning.

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

**Using advanced options**

When cloning a disk comprising of a **system volume**, you need to retain an operating system bootability on the target disk volume. It means that the operating system must have the system volume information (e.g. volume letter) matched with the disk NT signature, which is kept in the MBR disk record. But two disks with the same NT signature cannot work properly under one operating system.

*If there are two disks having the same NT signature and comprising of a system volume on a machine, at the startup the operating system runs from the first disk, discovers the same signature on the second one, automatically generates a new unique NT signature and assigns it to the second disk. As a result, all the volumes on the second disk will lose their letters, all paths will be invalid on the disk, and programs won't find their files. The operating system on that disk will be unbootable.*

You have the following two alternatives to retain system bootability on the target disk volume:

1. Copy NT signature – to provide the target disk with the source disk NT signature matched with the Registry keys also copied on the target disk.

2. Leave NT signature – to keep the old target disk signature and update the operating system according to the signature.

***If you need to copy the NT signature:***

1. Select the **Copy NT signature** check box. You receive the warning: "If there is an operating system on the hard disk, uninstall either the source or the target hard disk drive from your machine prior to starting the machine again. Otherwise, the OS will start from the first of the two, and the OS on the second disk will become unbootable." The **Turn off the machine after the cloning operation** check box is selected and disabled automatically.

2. Click **Finish** to add the pending operation.

3. Click **Commit** on the toolbar and then click **Proceed** in the **Pending Operations** window.

4. Wait until the operation is finished.

5. Wait until the machine is turned off.

6. Disconnect either the source or the target hard disk drive from the machine.

7. Start up the machine.

***If you need to leave an NT signature:***

1. Click to clear the **Copy NT signature** check box, if necessary.

2. Click to clear the **Turn off the machine after the cloning operation** check box, if necessary.

3. Click **Finish** to add the pending operation.

4. Click **Commit** on the toolbar and then click **Proceed** in the **Pending Operations** window.

5. Wait until the operation is finished.

## 9.6.3    Disk conversion: MBR to GPT

You would want to convert an MBR basic disk to a GPT basic disk in the following cases:

- If you need more than 4 primary volumes on one disk.
- If you need additional disk reliability against any possible data damage.

***If you need to convert a basic MBR disk to basic GPT:***

1. Select a basic MBR disk to convert to GPT.

2. Right-click on the selected volume, and then click **Convert to GPT** in the context menu.

   You will receive a warning window, stating that you are about to convert MBR into GPT.

3. By clicking **OK**, you'll add a pending operation of MBR to GPT disk conversion.

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

*Please note: A GPT-partitioned disk reserves the space in the end of the partitioned area necessary for the backup area, which stores copies of the GPT header and the partition table. If the disk is full and the volume size cannot be automatically decreased, the conversion operation of the MBR disk to GPT will fail.*

*The operation is irreversible. If you have a primary volume, belonging to an MBR disk, and convert the disk first to GPT and then back to MBR, the volume will be logical and will not be able to be used as a system volume.*

If you plan to install an OS that does not support GPT disks, the reverse conversion of the disk to MBR is also possible through the same menu items the name of the operation will be listed as **Convert to MBR**.

### Dynamic disk conversion: MBR to GPT

Acronis Disk Director Lite does not support direct MBR to GPT conversion for dynamic disks. However you can perform the following conversions to reach the goal using the program:

1. MBR disk conversion: dynamic to basic    (p. 179) using the **Convert to basic** operation.
2. Basic disk conversion: MBR to GPT    using the **Convert to GPT** operation.
3. GPT disk conversion: basic to dynamic    (p. 178) using the **Convert to dynamic** operation.

## 9.6.4    Disk conversion: GPT to MBR

If you plan to install an OS that does not support GPT disks, conversion of the GPT disk to MBR is possible the name of the operation will be listed as **Convert to MBR**.

**If you need to convert a GPT disk to MBR:**

1. Select a GPT disk to convert to MBR.
2. Right-click on the selected volume, and then click **Convert to MBR** in the context menu.

   You will receive a warning window, stating that you are about to convert GPT into MBR.

   You will be explained the changes that will happen to the system after the chosen disk is converted from GPT to MBR. E.g. if such conversion will stop a disk from being accessed by the system, the operating system will stop loading after such conversion or some volumes on the selected GPT disk will not be accessible with MBR (e.g. volumes located more than 2 TB from the beginning of the disk) you will be warned here about such damage.

   *Please note, a volume, belonging to a GPT disk to convert, will be a logical one after the operation and is irreversible.*

3. By clicking **OK**, you'll add a pending operation of GPT to MBR disk conversion.

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

## 9.6.5    Disk conversion: basic to dynamic

You would want to convert a basic disk to dynamic in the following cases:

- If you plan to use the disk as part of a dynamic disk group.
- If you want to achieve additional disk reliability for data storage.

*If you need to convert a basic disk to dynamic:*

1. Select the basic disk to convert to dynamic.
2. Right-click on the selected volume, and then click **Convert to dynamic** in the context menu. You will receive a final warning about the basic disk being converted to dynamic.
3. If you click **OK** in this warning window, the conversion will be performed immediately and if necessary, your machine will be restarted.

   *Please note: A dynamic disk occupies the last megabyte of the physical disk to store the database, including the four-level description (Volume-Component-Partition-Disk) for each dynamic volume. If during the conversion to dynamic it turns out that the basic disk is full and the size of its volumes cannot be decreased automatically, the basic disk to dynamic conversion operation will fail.*

Should you decide to revert your dynamic disks back to basic ones, e.g. if you want to start using an OS on your machine that does not support dynamic disks, you can convert your disks using the same menu items, though the operation now will be named **Convert to basic**.

**System disk conversion**

Acronis Disk Director Lite does not require an operating system reboot after basic to dynamic conversion of the disk, if:

1. There is a single Windows 2008/Vista operating system installed on the disk.
2. The machine runs this operating system.

*Basic to dynamic conversion of the disk, comprising of system volumes, takes a certain amount of time, and any power loss, unintentional turning off of the machine or accidental pressing of the Reset button during the procedure could result in bootability loss.*

In contrast to Windows Disk Manager the program ensures bootability of an **offline operating system** on the disk after the operation.

# 9.6.6    Disk conversion: dynamic to basic

You would want to convert dynamic disks back to basic ones, e.g. if you want to start using an OS on your machine    that does not support dynamic disks.

***If you need to convert a dynamic disk to basic:***

1. Select the dynamic disk to convert to basic.
2. Right-click on the selected volume, and then click **Convert to basic** in the context menu. You will receive a final warning about the dynamic disk being converted to basic.

   You will be advised about the changes that will happen to the system if the chosen disk is converted from dynamic into basic. E.g. if such a conversion will stop the disk from being accessed by the system, the operating system will stop loading after such conversion, or if the disk you want to convert to basic contains any volumes of the types that are only supported by dynamic disks (all volume types except Simple volumes), then you will be warned here about the possible damage to the data involved in the conversion.

   *Please note, the operation is unavailable for a dynamic disk containing Spanned, Striped, or RAID-5 volumes.*

3. If you click **OK** in this warning window, the conversion will be performed immediately.

After the conversion the last 8Mb of disk space is reserved for the future conversion of the disk from basic to dynamic.

In some cases the possible unallocated space and the proposed maximum volume size might differ (e.g. when the size of one mirror establishes the size of the other mirror, or the last 8Mb of disk space are reserved for the future conversion of the disk from basic to dynamic).

**System disk conversion**

Acronis Disk Director Lite does not require an operating system reboot after dynamic to basic conversion of the disk, if:

1. There is a single Windows 2008/Vista operating system installed on the disk.
2. The machine runs this operating system.

*Dynamic to basic conversion of the disk, comprising of system volumes, takes a certain amount of time, and any power loss, unintentional turning off of the machine or accidental pressing of the Reset button during the procedure could result in bootability loss.*

In contrast to Windows Disk Manager the program ensures:

- safe conversion of a dynamic disk to basic when it contains volumes **with data** for simple and mirrored volumes
- in multiboot systems, bootability of a system that was **offline** during the operation

## 9.6.7    Changing disk status

Changing disk status is effective for Windows Vista SP1, Windows Server 2008, Windows 7 operating systems and applies to the current disk layout (p. 173).

One of the following disk statuses always appears in the graphical view of the disk next to the disk's name:

- **Online**

    The online status means that a disk is accessible in the read-write mode. This is the normal disk status. If you need a disk to be accessible in the read-only mode, select the disk and then change its status to offline by selecting **Change disk status to offline** from the **Operations** menu.

- **Offline**

    The offline status means that a disk is accessible in the read-only mode. To bring the selected offline disk back to online, select **Change disk status to online** from the **Operations** menu. If the disk has the offline status and the disk's name is **Missing**, this means that the disk cannot be located or identified by the operating system. It may be corrupted, disconnected, or powered off. For information on how to bring a disk that is offline and missing back online, please refer to the following Microsoft knowledge base article: http://technet.microsoft.com/en-us/library/cc732026.aspx.

## 9.7    Volume operations

Acronis Disk Director Lite includes the following operations that can be performed on volumes:

- Create Volume (p. 180) - Creates a new volume with the help of the Create Volume Wizard.
- Delete Volume (p. 184) - Deletes the selected volume.
- Set Active (p. 185) - Sets the selected volume Active so that the machine will be able to boot with the OS installed there.
- Change Letter (p. 185) - Changes the selected volume letter
- Change Label (p. 186) - Changes the selected volume label
- Format Volume (p. 186) - Formats a volume giving it the necessary file system

The full version of Acronis Disk Director will provide more tools and utilities for working with volumes.

*Acronis Disk Director Lite must obtain exclusive access to the target volume. This means no other disk management utilities (like Windows Disk Management utility) can access it at that time. If you receive a message stating that the volume cannot be blocked, close the disk management applications that use this volume and start again. If you can not determine which applications use the volume, close them all.*

## 9.7.1    Creating a volume

You might need a new volume to:

- Recover a previously saved backup copy in the "exactly as was" configuration;

- Store collections of similar files separately — for example, an MP3 collection or video files on a separate volume;
- Store backups (images) of other volumes/disks on a special volume;
- Install a new operating system (or swap file) on a new volume;
- Add new hardware to a machine.

In Acronis Disk Director Lite the tool for creating volumes is the **Create volume Wizard**.

## 9.7.1.1    Types of dynamic volumes

**Simple Volume**

A volume created from free space on a single physical disk. It can consist of one region on the disk or several regions, virtually united by the Logical Disk Manager (LDM). It provides no additional reliability, no speed improvement, nor extra size.

**Spanned Volume**

A volume created from free disk space virtually linked together by the LDM from several physical disks. Up to 32 disks can be included into one volume, thus overcoming the hardware size limitations, but if at least one disk fails, all data will be lost, and no part of a spanned volume may be removed without destroying the entire volume. So, a spanned volume provides no additional reliability, nor a better I/O rate.

**Striped Volume**

A volume, also sometimes called RAID 0, consisting of equal sized stripes of data, written across each disk in the volume; it means that to create a striped volume, a user will need two or more dynamic disks. The disks in a striped volume don't have to be identical, but there must be unused space available on each disk that you want to include in the volume and the size of the volume will depend on the size of the smallest space. Access to the data on a striped volume is usually faster than access to the same data on a single physical disk, because the I/O is spread across more than one disk.

Striped volumes are created for improved performance, not for their better reliability - they do not contain redundant information.

**Mirrored Volume**

A fault-tolerant volume, also sometimes called RAID 1, whose data is duplicated on two identical physical disks. All of the data on one disk is copied to another disk to provide data redundancy. Almost any volume can be mirrored, including the system and boot volumes, and if one of the disks fails, the data can still be accessed from the remaining disks. Unfortunately, the hardware limitations on size and performance are even more severe with the use of mirrored volumes.

**Mirrored-Striped Volume**

A fault-tolerant volume, also sometimes called RAID 1+0, combining the advantage of the high I/O speed of the striped layout and redundancy of the mirror type. The evident disadvantage remains inherent with the mirror architecture - a low disk-to-volume size ratio.

**RAID-5**

A fault-tolerant volume whose data is striped across an array of three or more disks. The disks do not need to be identical, but there must be equally sized blocks of unallocated space available on each disk in the volume. Parity (a calculated value that can be used to reconstruct data in case of failure) is also striped across the disk array. And it is always stored on a different disk than the

data itself. If a physical disk fails, the portion of the RAID-5 volume that was on that failed disk can be re-created from the remaining data and the parity. A RAID-5 volume provides reliability and is able to overcome the physical disk size limitations with a higher than mirrored disk-to-volume size ratio.

## 9.7.1.2    Create volume wizard

The **Create volume** wizard lets you create any type of volume      (including system and active), select a file system, label, assign a letter, and also provides other disk management functions.

Its pages will enable you to enter operation parameters, proceeding step-by-step further on and return to any previous step if necessary to change any previously selected options. To help you with your choices, each parameter is supplemented with detailed instructions.

**If you want to create a volume**:

Run the **Create volume** wizard by selecting **Create volume** on the **Wizards** bar, or right-click any unallocated space and select **Create volume** in the appearing context menu.

### Select the type of volume being created

At the first step you have to specify the type of volume you want to create. The following types of volume are available:

- Basic
- Simple/Spanned
- Striped
- Mirrored
- RAID-5

You will obtain a brief description of every type of volume    for better understanding of the advantages and limitations of each possible volume architecture.

*If the current operating system, installed on this machine, does not support the selected type of volume , you will receive the appropriate warning. In this case the **Next** button will be disabled and you will have to select another type of volume to proceed with the new volume creation.*

After you click the **Next** button, you will proceed forward to the next wizard page: Select destination disks (p. 182).

### Select destination disks

The next wizard page will prompt you to choose the disks, whose space will be used for the volume creation.

**To create a basic volume:**

- Select a destination disk and specify the unallocated space to create the basic volume on.

**To create a Simple/Spanned volume:**

- Select one or more destination disks to create the volume on.

**To create a Mirrored volume:**

- Select two destination disks to create the volume on.

**To create a Striped volume:**

- Select two or more destination disks to create the volume on.

**To create a RAID-5 volume:**

- Select three destination disks to create the volume on.

After you choose the disks, the wizard will calculate the maximum size of the resulting volume, depending on the size of the unallocated space on the disks you chose and the requirements of the volume type you have previously decided upon.

If you are creating a **dynamic** volume and select one or several **basic** disks, as its destination, you will receive a warning that the selected disk will be converted to dynamic automatically.

If need be, you will be prompted to add the necessary number of disks to your selection, according to the chosen type of the future volume.

If you click the **Back** button, you will be returned to the previous page: Select the type of volume being created (p. 182).

If you click the **Next** button, you will proceed to the next page: Set the volume size (p. 183).

## Set the volume size

On the third wizard page, you will be able to define the size of the future volume, according to the previously made selections. In order to choose the necessary size between the minimum and the maximum values, use the slider or enter the necessary values into the special windows between the minimum and the maximum values or click on the special handle, and hold and drag the borders of the disk's picture with the cursor.

The maximum value normally includes the most possible unallocated space. But in some cases the possible unallocated space and the proposed maximum volume size might differ (e.g. when the size of one mirror establishes the size of the other mirror, or the last 8Mb of the disk space is reserved for the future conversion of the disk from basic to dynamic).

For basic volumes if some unallocated space is left on the disk, you also will be able to choose the position of the new volume on the disk.

If you click the **Back** button, you will be returned to the previous page: Select destination disks (p. 182).

If you click the **Next** button, you will proceed to the next    page: Set the volume options (p. 183).

## Set the volume options

On the next wizard page you can assign the volume **Letter** (by default - the first free letter of the alphabet) and, optionally, a **Label** (by default – none). Here you will also specify the **File system** and the **Cluster size**.

The wizard will prompt you to choose one of the Windows file systems: FAT16 (disabled, if the volume size has been set at more than 2 GB), FAT32 (disabled, if the volume size has been set at more than 2 TB), NTFS or to leave the volume **Unformatted**.

In setting the cluster size you can choose between any number in the preset amount for each file system. Note, the program suggests the cluster size best suited to the volume with the chosen file system.

If you are creating a basic volume, which can be made into a system volume, this page will be different, giving you the opportunity to select the volume **Type** — **Primary** (**Active Primary**) or **Logical**.

Typically **Primary** is selected to install an operating system to a volume. Select the **Active** (default) value if you want to install an operating system on this volume to boot at machine startup. If the **Primary** button is not selected, the **Active** option will be inactive. If the volume is intended for data storage, select **Logical**.

*A Basic disk can contain up to four primary volumes. If they already exist, the disk will have to be converted into dynamic, otherwise or **Active** and **Primary** options will be disabled and you will only be able to select the **Logical** volume type. The warning message will advise you that an OS installed on this volume will not be bootable.*

*If you use characters when setting a new volume label that are unsupported by the currently installed operation system, you will get the appropriate warning and the **Next** button will be disabled. You will have to change the label to proceed with the creation of the new volume.*

If you click the **Back** button, you will be returned to the previous page: Set the volume size (p. 183).

If you click the **Finish** button, you will complete the operation planning.

To perform the planned operation click **Commit** in the toolbar, and then click **Proceed** in the **Pending Operations** window**.**

*If you set a 64K cluster size for FAT16/FAT32 or on 8KB-64KB cluster size for NTFS, Windows can mount the volume, but some programs (e.g. Setup programs) might calculate its disk space incorrectly.*

## 9.7.2    Delete volume

*This version of Acronis Disk Director Lite has reduced functionality because it is mainly a tool for preparing bare-metal systems for recovering previously saved volume images. The features of resizing the existing volumes and creating the new volumes, using free space from the existing ones, exist on the full version of the software, so with this version deleting an existing volume sometimes might be the only way to free the necessary disk space without changing the existing disk configuration.*

After a volume is deleted, its space is added to unallocated disk space. It can be used for creation of a new volume or to change another volume's type.

***If you need to delete a volume:***

1. Select a hard disk and a volume to be deleted.
2. Select **Delete volume** or a similar item in the **Operations** sidebar list, or click the **Delete the selected volume** icon on the toolbar.

   *If the volume contains any data, you will receive the warning, that all the information on this volume will be lost irrevocably.*

3. By clicking **OK** in the **Delete volume** window, you'll add the pending operation of volume deletion.

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

### 9.7.3 Set active volume

If you have several primary volumes, you must specify one to be the boot volume. For this, you can set a volume to become active. A disk can have only one active volume, so if you set a volume as active, the volume, which was active before, will be automatically unset.

***If you need to set a volume active:***

1.  Select a primary volume on a basic MBR disk to set as active.

2.  Right-click on the selected volume, and then click **Mark as active** in the context menu.

    If there is no other active volume in the system, the pending operation of setting active volume will be added.

    *Please note, that due to setting the new active volume, the former active volume letter might be changed and some of the installed programs might stop running.*

3.  If another active volume is present in the system, you will receive the warning that the previous active volume will have to be set passive first. By clicking **OK** in the **Warning** window, you'll add the pending operation of setting active volume.

    *Please note: even if you have the Operating System on the new active volume, in some cases the machine will not be able to boot from it. You will have to confirm your decision to set the new volume as active.*

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the **Disk management** view immediately.

### 9.7.4 Change volume letter

Windows operating systems assign letters (C:, D:, etc) to hard disk volumes at startup. These letters are used by applications and operating systems to locate files and folders in the volumes.

Connecting an additional disk, as well as creating or deleting a volume on existing disks, might change your system configuration. As a result, some applications might stop working normally or user files might not be automatically found and opened. To prevent this, you can manually change the letters that are automatically assigned to the volumes by the operating system.

***If you need to change a letter assigned to a volume by the operating system:***

1.  Select a volume to change a letter.
2.  Right-click on the selected volume, and then click **Change letter** in the context menu.
3.  Select a new letter in the **Change Letter** window.
4.  By clicking **OK** in the **Change Letter** window, you'll add a pending operation to volume letter assignment.

    (To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the **Disk management** view immediately.

## 9.7.5   Change volume label

The volume label is an optional attribute. It is a name assigned to a volume for easier recognition. For example, one volume could be called SYSTEM — a volume with an operating system, or PROGRAM — an application volume, DATA — a data volume, etc., but it does not imply that only the type of data stated with the label could be stored on such a volume.

In Windows, volume labels are shown in the Explorer disk and folder tree: LABEL1(C:), LABEL2(D:), LABEL3(E:), etc. LABEL1, LABEL2 and LABEL3 are volume labels. A volume label is shown in all application dialog boxes for opening and saving files.

***If you need to change a volume label:***

1.   Right-click on the selected volume, and then click **Change label**.

2.   Enter a new label in the **Change label** window text field.

3.   By clicking **OK** in the **Change label** window, you'll add the pending operation of changing the volume label .

> *If when setting a new volume label you use characters that are unsupported by the currently installed operating system, you will get the appropriate warning and the **OK** button will be disabled. You will have to use only supported characters to proceed with changing the volume label.*

(To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new label will be graphically represented in the **Disk Management** view of the console immediately.

## 9.7.6   Format volume

You might want to format a volume if you want to change its file system:

- to save additional space which is being lost due to the cluster size on the FAT16 or FAT32 file systems

- as a quick and more or less reliable way of destroying data, residing in this volume

***If you want to format a volume:***

1.   Select a volume to format.

2.   Right-click on the selected volume, and then click **Format** in the context menu.

  You will be forwarded to the **Format Volume** window, where you will be able to set the new file system options. You can choose one of the Windows file systems: FAT16 (disabled, if the Volume Size is more than 2 GB), FAT32 (disabled, if the Volume Size is more than 2 TB) or NTFS.

  In the text window you will be able to enter the volume label, if necessary: by default this window is empty.

  In setting the cluster size you can choose between any number in the preset amount    for each file system. Note, the program suggests the cluster size best suited to the volume with the chosen file system.

3.   If you click **OK** to proceed with the **Format Volume** operation, you'll add a pending operation of formatting a volume.

  (To finish the added operation you will have to commit (p. 187) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the **Disk management** view.

*If you set a 64K cluster size for FAT16/FAT32 or an 8KB-64KB cluster size for NTFS, Windows can mount the volume, but some programs (e.g. Setup programs) might calculate its disk space incorrectly.*

# 9.8    Pending operations

All operations, which were prepared by the user in manual mode or with the aid of a wizard, are considered pending until the user issues the specific command for the changes to be made permanent. Until then, Acronis Disk Director Lite will only demonstrate the new volume structure that will result from the operations that have been planned to be performed on disks and volumes. This approach enables you to control all planned operations, double-check the intended changes, and, if necessary, cancel operations before they are executed.

To prevent you from performing any unintentional change on your disk, the program will first display the list of all pending operations.

The **Disk management** view contains the toolbar with icons to launch the **Undo**, **Redo** and **Commit** actions intended for pending operations. These actions might also be launched from the **Disk management** menu of the console.

All planned operations are added to the pending operation list.

The **Undo** action lets you undo the latest operation in the list. While the list is not empty, this action is available.

The **Redo** action lets you reinstate the last pending operation that was undone.

The **Commit** action forwards you to the **Pending Operations** window, where you will be able to view the pending operation list. Clicking **Proceed** will launch their execution. You will not be able to undo any actions or operations after you choose the **Proceed** operation. You can also cancel the commitment by clicking **Cancel**. Then no changes will be done to the pending operation list.

Quitting Acronis Disk Director Lite without committing the pending operations effectively cancels them, so if you try to exit **Disk management** without committing the pending operations, you will receive the appropriate warning.

# 10 Administering a managed machine

This section describes the views that are available through the navigation tree of the console connected to a managed machine and explains how to work with each view.

## 10.1 Backup plans and tasks

The **Backup plans and tasks** view keeps you informed of data protection on a given machine. It lets you monitor and manage backup plans and tasks.

To find out what a backup plan is currently doing on the machine, check the backup plan execution state (p. 190). A backup plan execution state is a cumulative state of the plan's most recent activities. The status of a backup plan (p. 191) helps you to estimate whether the data is successfully protected.

To keep track of a task's current progress, examine its state (p. 192). Check a task status (p. 192) to ascertain the result of a task.

**Typical workflow**

- Use filters to display the desired backup plans (tasks) in the backup plans table. By default, the table displays all the plans of the managed machine sorted by name. You can also hide the unneeded columns and show the hidden ones. For details, see "Sorting, filtering and configuring table items" (p. 16).

- In the backup table, select the backup plan (task).

- Use the toolbar's buttons to take an action on the selected plan (task). For details, see "Actions on backup plans and tasks" (p. 188).

- To review detailed information on the selected plan (task), use the information panel at the bottom of the window. The panel is collapsed by default. To expand the panel, click the arrow mark ( ▲ ). The content of the panel is also duplicated in the **Plan details** (p. 197) and **Task details** (p. 198) windows respectively.

## 10.1.1 Actions on backup plans and tasks

The following is a guideline for you to perform operations with backup plans and tasks.

**Restrictions**

- Without the Administrator privileges on the machine, a user cannot run or modify plans or tasks owned by other users.

- It is not possible to modify or delete a currently running backup plan or task.

- A centralized backup plan or task can be modified or deleted only on the management server side.

| To | Do |
|---|---|
| Create a new backup plan or task | Click 📋 **New**, then select one of the following:<br><br>- Backup plan (p. 29)<br><br>- Recovery task (p. 97)<br><br>- Validation task (p. 143) |

| To | Do |
|---|---|
| View details of a plan/task | Click 🔍 **Details**.<br>In the respective **Plan Details** (p. 197) or **Task Details (p. 198)** window, review the plan or task details. |
| View plan's/task's log | Click 📋 **Log**.<br><br>You will be taken to the **Log** (p. 199) view containing the list of the log entries grouped by the plan/task-related activities. |
| Run a plan/task | **Backup plan**<br><br>1. Click ▶ **Run**.<br><br>2. In the drop-down list, select the plan's task you need run.<br><br>Running the backup plan starts the selected task of that plan immediately in spite of its schedule and conditions.<br><br>**Task**<br>Click ▶ **Run**.<br><br>The task will be executed immediately in spite of its schedule and conditions. |
| Stop a plan/task | Click 🟥 **Stop**.<br><br>**Backup plan**<br><br>Stopping the running backup plan stops all its tasks. Thus, all the task operations will be aborted.<br><br>**Task**<br><br>Stopping a task aborts its operation (recovery, validation, exporting, conversion, etc.). The task enters the **Idle** state. The task schedule, if created, remains valid. To complete the operation you will have to run the task over again.<br><br>**What will happen if I stop the recovery task?**<br><br>■ Recovering disks: the aborted operation may cause changes in the target disk. Depending on the time that has passed since the task run, the target disk may not be initialized, or the disk space may be unallocated, or some volumes may be recovered and others not. To recover the entire disk, run the task once again.<br><br>■ Recovering volumes: the target volume will be deleted and its space unallocated – the same result you will get if the recovery is unsuccessful. To recover the "lost" volume, run the task once again.<br><br>■ Recovering files or folders: the aborted operation may cause changes in the destination folder. Depending on the time that has passed since the task run, some files may be recovered, but some not. To recover all the files, run the task once again. |

| To | Do |
|---|---|
| Edit a plan/task | Click ✏ **Edit**.<br><br>Backup plan editing is performed in the same way as creation (p. 29), except for the following **limitations**:<br><br>It is not always possible to use all scheme options, when editing a backup plan if the created archive is not empty (i.e. contains backups).<br><br>1. It is not possible to change the scheme to Grandfather-Father-Son or Tower of Hanoi.<br><br>2. If the Tower of Hanoi scheme is used, it is not possible to change the number of levels.<br><br>In all other cases the scheme can be changed, and should continue to operate as if existing archives were created by a new scheme. For empty archives all changes are possible. |
| Clone a backup plan | Click ▦ **Clone**.<br><br>The clone of the original backup plan will be created with default name "*Clone of <original_plan_name>*". The cloned plan will be disabled immediately after cloning, so that it does not run concurrently with the original plan. You can edit the cloned plan settings before enabling it. |
| Enable a plan | Click 📋 **Enable**.<br><br>The previously disabled backup plan will run again as scheduled. |
| Disable a plan | Click 📋 **Disable**.<br><br>The backup plan will not run as scheduled. However, it can be started manually. After a manual run, the plan will stay disabled. The plan will run as usual if you enable it again. |
| Export a plan | Click ⬇ **Export**.<br><br>Specify the path and name of the resulting file. See Export and import of backup plans (p. 193) for more information. |
| Import a plan | Click ⬆ **Import**.<br><br>Specify the path and name of the file that contains a previously exported plan. See Export and import of backup plans (p. 193) for more information. |
| Delete a plan/task | Click ✕ **Delete**. |

# 10.1.2   States and statuses of backup plans and tasks

## 10.1.2.1   Backup plan execution states

A backup plan state is a cumulative state of the plan's tasks/activities.

| | State | How it is determined | How to handle |
|---|---|---|---|

| | | | |
|---|---|---|---|
| 1 | **Need interaction** | At least one task needs user interaction.<br><br>Otherwise, see 2. | Identify the tasks that need interaction (the program will display what action is needed) -> Stop the tasks or enable the tasks to run (change media; provide additional space on the vault; ignore the read error; create the missing Acronis Secure Zone). |
| 2 | **Running** | At least one task is running.<br><br>Otherwise, see 3. | No action is required. |
| 3 | **Waiting** | At least one task is waiting.<br><br>Otherwise, see 4. | Waiting for condition. This situation is quite normal, but delaying a backup for too long is risky. The solution may be to set the maximum delay (p. 94) after which the task will start anyway or force the condition (tell the user to log off, enable the required network connection.)<br><br>Waiting while another task locks the necessary resources. A one-time waiting case may occur when a task start is delayed or a task run lasts much longer than usual for some particular reason and prevents another task from starting. This situation is resolved automatically when the obstructing task comes to an end. Consider stopping a task if it hangs for too long to enable the next task to start.<br><br>Persistent task overlapping may result from an incorrectly scheduled plan or plans. It makes sense to edit the plan in this case. |
| 4 | **Idle** | All the tasks are idle. | No action is required. |

## 10.1.2.2   Backup plan statuses

A backup plan can have one of the following statuses: **Error; Warning; OK**.

A backup plan status is derived from the results of the last run of the plans' tasks/activities.

| | Status | How it is determined | How to handle |
|---|---|---|---|
| 1 | **Error** | At least one task has failed.<br><br>Otherwise, see 2 | Identify the failed tasks -> Check the tasks log to find out the reason of the failure, then do one or more of the following:<br><br>▪ Remove the reason of the failure -> [optionally] Start the failed task manually<br><br>▪ Edit the local plan to prevent its future failure if a local plan has failed<br><br>▪ Edit the centralized backup plan on the management server if a centralized plan has failed |
| 2 | **Warning** | At least one task has succeeded with warnings.<br><br>Otherwise, see 3. | View the log to read the warnings -> [optionally] Perform actions to prevent the future warnings or failure. |
| 3 | **OK** | All the tasks are completed successfully. | No action is required. Note that a backup plan can be OK if none of the tasks has been started yet. |

## 10.1.2.3   Task states

A task can be in one of the following states: **Idle**; **Waiting**; **Running**; **Need interaction**. The initial task state is **Idle**.

Once the task is started manually or the event specified by the schedule occurs, the task enters either the **Running** state or the **Waiting** state.

**Running**

> A task changes to the **Running** state when the event specified by the schedule occurs AND all the conditions set in the backup plan are met AND no other task that locks the necessary resources is running. In this case, nothing prevents the task from running.

**Waiting**

> A task changes to the **Waiting** state when the task is about to start, but another task using the same resources is already running. In particular, more than one backup tasks cannot run simultaneously on a machine. A backup task and a recovery task also cannot run simultaneously, if they use the same resources. Once the other task unlocks the resource, the waiting task enters the **Running** state.
>
> A task may also change to the **Waiting** state when the event specified by the schedule occurs but the condition set in the backup plan is not met. See Task start conditions (p. 94) for details.

**Need interaction**

> Any running task can put itself into the **Need interaction** state when it needs human interaction such as changing media or ignoring a read error. The next state may be **Idle** (if the user chooses to stop the task) or **Running** (on selecting Ignore/Retry or another action, such as Reboot, that can put the task to the **Running** state.)

## 10.1.2.4   Task statuses

A task can have one of the following statuses: **Error**; **Warning**; **OK**.

A task status is derived from the result of the last run of the task.

| | Status | How it is determined | How to handle |
|---|---|---|---|
| 1 | **Error** | Last result is "Failed" | Identify the failed task -> Check the task log to find out the reason of the failure, then do one or more of the following: <br>■ Remove the reason of the failure -> [optionally] Start the failed task manually <br>■ Edit the failed task to prevent its future failure |
| 2 | **Warning** | Last result is "Succeeded with warning" or the task has been stopped | View the log to read the warnings -> [optionally] Perform actions to prevent the future warnings or failure. |
| 3 | **OK** | Last result is "Succeeded" or "Not run yet" | "Not run yet" means that the task has never been started or has been started, but has not finished yet and, therefore its result is not available. You may want to find out why the task has not started so far. |

# 10.1.3 Export and import of backup plans

The export operation creates a file with complete configuration of the backup plan. You can import the file to reuse the exported backup plan on another machine.

Centralized backup plans can be exported from a management server and imported to a management server only.

You can edit plans in the Acronis Backup & Recovery 11 graphical user interface when importing them or after. Backup plans are exported to .xml files, so you can edit the export files of backup plans (p. 194) with text editors. Passwords are encrypted in the export files.

## Usage examples

- **Agent reinstallation**

  Export the backup plans before reinstalling the agent and import them after reinstalling.

- **Deploying of a backup plan to multiple machines**

  You have an environment where using of Acronis Backup & Recovery 11 Management Server is not possible; for example, because of security restrictions. Nevertheless, you want to use the same backup plan on multiple machines. Export this plan from one of the machines and deploy it as a file (p. 196) to the other machines.

## Adjusting credentials

A scheduled plan contains credentials of the user account under which the plan's tasks run. The plan will not start on a machine where a user account with identical credentials does not exist. To avoid this situation, do one of the following:

- Create an account with identical credentials on the second machine.
- Edit credentials in the export file before importing. For details, see Editing the export file (p. 194).
- Edit credentials after importing the plan.

When creating a backup plan with manual start, do not change the **Run under the current user** setting in **Plan parameters** > **Show task credentials, comments, label**. With this setting, the plan's tasks will always run under the account of the user who starts them.

## Steps to perform

### To export a backup plan

1. Select a backup plan in the **Backup plans and tasks** view.
2. Click ⬎ **Export**.
3. Specify the path and name of the export file.
4. Confirm your choice.

### To import a backup plan

1. Click ⬈ **Import** in the **Backup plans and tasks** view.
2. Specify the path and name of the export file.
3. Acronis Backup & Recovery 11 will show the **Edit backup plan** page. In most cases, you need to update the plan's credentials and the access credentials to the backup destination. Make the necessary changes and click **Save**. Otherwise, click **Cancel**, and the plan will be imported as is.

## 10.1.3.1  Editing the export file

The export file is an .xml file and can be edited with a text editor.

Here is how to make some useful changes.

**How to modify credentials**

In the export file, the **`<login>`** tags include the user name and the **`<password>`** tags include the user password.

To modify credentials, change the **`<login>`** and **`<password>`** tags in the corresponding sections:

▪ plan's credentials - the **`<plan><options><common_parameters>`** section

▪ access credentials for the backed up data - the **`<plan><targets><inclusions>`** section

▪ access credentials for the backup destination - the **`<plan><locations>`** section.

Pay special attention to modifying the **`<password>`** tag. The tag that contains an encrypted password looks like **`<password encrypted="true">...</password>`**.

*To change the encrypted password*

1. In the command line, run the **`acronis_encrypt`** utility:

   **`acronis_encrypt UserPassword#1`**

   (here **`UserPassword#1`** is the password you want to encrypt).

2. The utility outputs a string, for example "**XXXYYYZZZ888**".

3. Copy this string and paste it into the tag as follows:

   **`<password encrypted="true">XXXYYYZZZ888</password>`**

The **`acronis_encrypt`** utility is available on any machine where Acronis Backup & Recovery 11 Management Console is installed. The path to the utility is as follows:

▪ %ProgramFiles%/Common Files/Acronis/Utils - in 32-bit Windows

▪ %ProgramFiles(x86)%/Common Files/Acronis/Utils - in 64-bit Windows

▪ /usr/sbin - in Linux

**How to make a backup plan use the agent's credentials**

Before importing or deploying the export file, delete the value of the required **`<login>`** tag. Then the imported or deployed plan will use credentials of the agent service.

**Example**

To make the backup plan run under the agent's credentials, find the **`<login>`** tag in the **`<plan><options><common_parameters>`** section. The tag looks like follows:

```
<login>
    Administrator
</login>
<password encrypted="true">
    XXXYYYZZZ888
</password>
```

Delete the value of the **`<login>`** tag, so that the tag looks like follows:

```
<login>
```

```
</login>
<password encrypted="true">
    XXXYYYZZZ888
</password>
```

**How to change items to back up**

## Replacing a directly specified item with another directly specified item

Inside the **<plan><targets><inclusions>** section:

1. Delete the **<ID>** tag.
2. Edit the value of the **<Path>** tag , which contains information about data to back up; for example, replace "**C:**" with "**D:**".

## Replacing a directly specified item with a selection template

Inside the **<plan><options><specific><inclusion_rules>** section:

1. Add the **<rules_type>** tag with "**disks**" or "**files**" value, depending on the type of the template you need.
2. Add   the **<rules>** tag.
3. Inside   the **<rules>** tag, add the **<rule>** with the required template. The template must correspond to the directly specified item. For example, if the specified item has the "**disks**" value, you can use the **[SYSTEM]**, **[BOOT]** and **[Fixed Volumes]** templates; but you cannot use the **[All Files]** or **[All Profiles Folder]** templates. For more information about templates, see "Selection rules for volumes" and "Selection rules for files and folders".
4. To add another template, repeat the step 3.

## Example

The following example illustrates how to replace a directly specified item with selection templates.

The original section:

```
<specific>
   <backup_type>
      disks
   </backup_type>
   <disk_level_options />
   <file_level_options />
   <inclusion_rules />
<specific>
```

The section after applying the selection templates:

```
<specific>
    <backup_type>
        disks
    </backup_type>
    <disk_level_options />
    <file_level_options />
    <inclusion_rules>
        <rules_type>
            disks
        </rules_type>
        <rules>
            <rule>
                [BOOT]
            </rule>
            <rule>
                [SYSTEM]
            </rule>
        </rules>
    </inclusion_rules>
<specific>
```

# 10.1.4 Deploying backup plans as files

Assume that for some reason you cannot run Acronis Backup & Recovery 11 Management Server in your environment, but you need to apply one and the same backup plan to multiple machines. A good decision is to export the backup plan from one machine and deploy it to all the other machines.

**How it works**

A dedicated folder for storing deployed plans exists on every machine where an agent is installed. The agent tracks changes in the dedicated folder. As soon as a new .xml file appears in the dedicated folder, the agent imports the backup plan from that file. If you change (or delete) an .xml file in the dedicated folder, the agent automatically changes (or deletes) the appropriate backup plan.

**Editing the export file**

A backup plan imported in such way cannot be edited through the graphical user interface. You can edit the export file (p. 194) with a text editor either before or after the deployment.

If you edit the file before the deployment, the changes will take effect on all the machines where the plan will be deployed. You may want to change the direct specification of the item to backup (such as C: or C:\Users) with a template (such as [SYSTEM] or [All Profiles Folder]). For more information about templates see Selection rules for volumes and Selection rules for files and folders.

You may also want to change credentials used by the plan.

***To deploy a backup plan as file***

1. Create a backup plan on one of the machines.
2. Export it to an .xml file (p. 193).
3. [Optional] Edit the export file. See Editing the export file (p. 194) for more information.
4. Deploy this .xml file to the dedicated folder.

    The dedicated folder path

**In Windows**

The default path to the dedicated folder is
**%ALLUSERSPROFILE%\Acronis\BackupAndRecovery\import**.

The path is stored in the registry key
**HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\MMS\Configuration\Import\folderPath**.

The absence of the key means that the agent does not monitor the dedicated folder.

To change the path, edit the key. The change will be applied after a restart of the agent.

**In Linux**

The default path to the dedicated folder is **/usr/lib/Acronis/BackupAndRecovery/import**.

The path is stored in the file **/etc/Acronis/BackupAndRecovery.config**.

To change the path, edit the **/usr/lib/Acronis/BackupAndRecovery/import** value in the
following tag:

```
<key name="Settings">
…
    <value name="ImportFolderPath" type="TString">
    "/usr/lib/Acronis/BackupAndRecovery/import"
    </value>
…
</key>
```

The change will be applied after a restart of the agent.

The absence of the tag means that the agent does not monitor the dedicated folder.

## 10.1.5  Backup plan details

The **Backup plan details** window (also duplicated on the **Information** panel) aggregates all
information on the selected backup plan.

The respective message will appear at the top of the tabs, if execution of the plan requires user
interaction. The message contains a brief description of the problem and action buttons that let you
select the appropriate action or stop the plan.

**Details**

The **Backup plans and tasks** tab provides the following general information on the selected plan:

- **Name** - name of the backup plan
- **Origin** - whether the plan was created directly on the machine (local origin), or deployed to the machine from the management server (centralized origin).
- **Execution state** - execution state (p. 190) of the backup plan.
- **Status** - status (p. 191) of the backup plan.
- **Machine** - name of the machine on which the backup plan exists (only for centralized backup plans).
- **Schedule** - whether the task is scheduled, or set to start manually.
- **Last start time** - how much time has passed since the last plan or task start.
- **Deployment state** - the deployment states of the backup plan (only for centralized backup plans).

- **Last finish time** - how much time has passed since the last plan or task end.
- **Last result** - the result of the last plan or task run.
- **Type** - backup plan or task type.
- **Owner** - the name of the user who created or last modified the plan
- **Next start time** - when the plan or task will start the next time.
- **Comments** - description of the plan (if provided).

**Tasks**

The **Tasks** tab displays a list of all tasks of the selected backup plan. To view the selected task details, click **Details**.

**Progress**

The **Progress** tab lists all the selected backup plan's activities that are currently running or waiting for their turn to run.

**History**

The **History** tab lets you examine the history of all the backup plan's accomplished activities.

**What to back up**

The **Source** tab provides the following information on the data selected for backup:

- **Source type** - the type of data selected for backing up.
- **Items to back up** - items selected to back up and their size.

**Where to back up**

The **Destination** tab provides the following information:

- **Name** - name of the archive.
- **Location** - name of the vault or path to the folder, where the archive is stored.
- **Archive comments** - comments on the archive (if provided).
- **2nd**, **3rd**, **4th**, **5th location** - names of the locations to which the archive was copied or moved (if specified in the backup plan).

**Settings**

The **Settings** tab displays the following information:

- **Backup scheme** - the selected backup scheme and all its settings with schedules.
- **Validation** - if specified, events before or after which the validation is performed, and validation schedule. If the validation is not set, the **Never** value is displayed.
- **Backup options** - backup options changed against the default values.

## 10.1.6  Task/activity details

The **Task/activity details** window (also duplicated on the **Information** panel) aggregates on several tabs all information about the selected task or activity.

When a task or activity requires user interaction, a message and action buttons appear above the tabs. The message contains a brief description of the problem. The buttons allow you to retry or stop the task or the activity.

## 10.2 Log

The local event log stores the history of operations performed by Acronis Backup & Recovery 11 on the machine.

To view a plain list of log entries, select **Events** in the **Display** drop-down list; to view log entries grouped by activities, select **Activities**. The details of the selected log entry or activity are shown in the **Information** panel at the bottom of the **Log** view.

Use filters to display the desired activities and log entries in the table. You can also hide the unneeded columns and show the hidden ones. For details, see "Sorting, filtering and configuring table items" (p. 16).

Select the activity or log entry to take an action on log entries. For details, see "Actions on log entries" (p. 199) and "Log entry details" (p. 200).

### 10.2.1 Actions on log entries

All the operations described below are performed by clicking the corresponding items on the log **toolbar**. These operations can also be performed with the context menu (by right-clicking the log entry or the activity).

The following is a guideline for you to perform actions on log entries.

| To | Do |
|---|---|
| Select a single activity | Select **Activities** in the **Display** drop-down list and click an activity. <br><br> The **Information** pane will show log entries for the selected activity. |
| Select a single log entry | Click on it. |
| Select multiple log entries | ■ *non-contiguous*: hold down **CTRL** and click the log entries one by one <br><br> ■ *contiguous*: select a single log entry, then hold down **SHIFT** and click another log entry. All the log entries between the first and last selections will be selected too. |
| View a log entry's details | 1. Select a log entry. <br><br> 2. Do one of the following: <br><br>   ■ Double click the selection. <br><br>   ■ Click 🔍 **Details**. <br><br> The log entry's details will be displayed. See Log entry details for details of the log entry's operations. |
| Save the selected log entries to a file | 1. Display **Activities** and select activities or display **Events** and select log entries. <br><br> 2. Click 🖫 **Save selected to file**. <br><br> 3. In the opened window, specify a path and a name for the file. <br><br> All log entries of the selected activities or selected log entries will be saved to the specified file. |
| Save all the log entries to a file | 1. Make sure, that the filters are not set. <br><br> 2. Click 💾 **Save all to file**. <br><br> 3. In the opened window, specify a path and a name for the file. All log entries will be saved to the specified file. |

| Save all the filtered log entries to a file | 1. Set filters to get a list of the log entries that satisfy the filtering criteria. |
|---|---|
| | 2. Click 💾 **Save all to file**. |
| | 3. In the opened window, specify a path and a name for the file. |
| | All log entries in the list will be saved to the specified file. |
| Delete all the log entries | Click 🧹 **Clear log**. |
| | All the log entries will be deleted from the log, and a new log entry will be created. It will contain information about who deleted the log entries and when. |

## 10.2.2 Log entry details

Displays detailed information on the log entry you have selected and lets you copy the details to the clipboard.

To view details of the next or the previous log entry, click the down arrow button or correspondingly the up arrow button.

To copy the details, click the **Copy to clipboard** button.

**Log entry data fields**

A log entry contains the following data fields:

- **Type** - Type of event (Error; Warning; Information).
- **Date and time** - Date and time when the event took place.
- **Backup plan** - The backup plan the event relates to (if any).
- **Task** - The task the event relates to (if any).
- **Code** - It can be blank or the program error code if the event type is error. Error code is an integer number that may be used by Acronis support service to solve the problem.
- **Module** - It can be blank or the number of the program module where an error was occurred. It is an integer number that may be used by Acronis support service to solve the problem.
- **Owner** - User name of the backup plan owner (p. 21).
- **Message** - The event text description.

Date and time presentation varies depending on your locale settings.

## 10.3 Alerts

An alert is a message that warns about actual or potential problems. The **Alerts** view lets you rapidly identify and solve the problems by monitoring the current alerts and view the alerts history.

**Active and inactive alerts**

An alert can be either in an active, or inactive state. The active state indicates that the issue that caused the alert still exists. An active alert becomes inactive when the problem that caused the alert is resolved either manually or on its own.

*Note: There is one alert type that is always active: "Backup not created". This is because even if the cause of this alert was resolved and the following backups successfully created, the fact that the backup was not created remains.*

## Fixing issues that caused alerts

To find and fix the issue that caused the alert, click **Fix the issue**. You will be taken to the corresponding view,where you can examine the issue and take the necessary steps to resolve it.

Optionally, you can click **View details** to get more information about the alert you select.

## Accepting alerts

By default, the **Current alerts** table lists both active and inactive alerts until they are not accepted. To accept an alert, select it and then click **Accept**. By accepting an alert you acknowledge the alert and agree to takeresponsibility for it. The accepted alerts are then moved to the **Accepted alerts** table, with the alert state unchanged.

The **Accepted alerts** table stores the history of the accepted alerts. Here, you can find out who accepted the alert and when it happen. The accepted alerts of both states can be removed from the table either manually, by using **Delete** and **Delete all** buttons, or automatically (see "Configuring alerts" later in this section).

To export entire table contents to a *.txt or *.csv file, click **Save all to file**.

## Configuring alerts

Use the following options at the top of the **Alerts** view to configure alerts:

- **Show/hide alerts** (p. 18) - specify the alert types to display in the **Alerts** view.
- **Notifications** (p. 203) - set up e-mail notifications about alerts.
- **Settings** (p. 202) - specify whether to move inactive alerts to the **Accepted alerts** table automatically; set how long to keep the accepted alerts in the **Accepted alerts** table.

# 10.4 Collecting system information

The system information collection tool gathers information about the machine to which the management console is connected, and saves it to a file. You may want to provide this file when contacting Acronis technical support.

This option is available under bootable media and for machines where Agent for Windows, Agent for Linux or Acronis Backup & Recovery 11 Management Server is installed.

### *To collect system information*

1. In the management console, select from the top menu **Help** > **Collect system information from** 'machine name'.
2. Specify where to save the file with system information.

# 10.5 Adjusting machine options

The machine options define the general behavior of all Acronis Backup & Recovery 11 agents operating on the managed machine, and so the options are considered machine-specific.

To access the machine options, connect the console to the managed machine and then select **Options > Machine options** from the top menu.

## 10.5.1  Customer Experience Program

This option defines whether the machine will participate in the Acronis Customer Experience Program (ACEP).

If you choose **Yes, I want to participate in the ACEP**, information about the hardware configuration, the most and least used features and about any problems will be automatically collected from the machine and sent to Acronis on a regular basis. The end results are intended to provide software improvements and enhanced functionality to better meet the needs of Acronis customers.

Acronis does not collect any personal data. To learn more about the ACEP, read the terms of participation on the Acronis Web site or in the product GUI.

Initially the option is configured during the Acronis Backup & Recovery 11 agent installation. This setting can be changed at any time using the product GUI (**Options** > **Machine options** > **Customer Experience Program**). The option can also be configured using the Group Policy infrastructure. A setting defined by a Group Policy cannot be changed using the product GUI unless the Group Policy is disabled on the machine.

# 10.5.2  Alerts

## 10.5.2.1  Alert management

**Remove from "Accepted alerts" items older than**

This option defines whether to delete the accepted alerts from the **Accepted alerts** table.

The preset is: **Disabled**.

When enabled, you can specify the keeping period for the accepted alerts. The accepted alerts older than this period will be deleted from the table automatically.

**Automatically move inactive alerts to "Accepted alerts"**

This option defines whether to accept all the alerts that become inactive and move them to the **Accepted alerts** table automatically.

The preset is: **Disabled**.

When enabled, you can specify the alert types to apply this option to.

## 10.5.2.2  Time-based alerts

**Last backup**

This option is effective when the console is connected to a managed machine (p. 237) or to the management server (p. 237).

The option defines whether to alert if no backup was performed on a given machine for a period of time. You can configure the time period that is considered critical for your business.

The preset is: alert if the last successful backup on a machine was completed more than **5 days** ago.

The alert is displayed in the **Alerts** view of the **Navigation** pane. When the console is connected to the management server, this setting will also control the color scheme of the **Last backup** column's value for each machine.

**Last connection**

This option is effective when the console is connected to the management server or to a registered machine (p. 238).

The option defines whether to alert if no connection was established between a registered machine and the management server for a period of time so indicating that the machine might not be centrally managed (for instance in the case of network connection failure to that machine). You can configure the length of time that is considered critical.

The preset is: alert if the machine's last connection to the management server was more than **5 days** ago.

The alert is displayed in the **Alerts** view of the **Navigation** pane. When the console is connected to the management server, this setting will also control the color scheme of the **Last connect** column's value for each machine.

## 10.5.3   E-mail notifications

The option enables you to configure e-mail notifications.

The preset is: **Disabled**.

***To configure e-mail notification***

1.  In the **SMTP server** field**,** enter the name of the SMTP server.
2.  In the **Port** field– set the port of the SMTP server. By default, the port is set to 25.
3.  In the **User name** filed, enter the user name.
4.  In the **Password** field, enter the password.
5.  Click **Additional e-mail parameters...** to configure additional e-mail parameters as follows, then click **OK**:
    -   **From** - type the e-mail address of the user from whom the message will be sent. If you leave this field empty, messages will be constructed as if they are from the destination address.
    -   **Use encryption** – you can opt for encrypted connection to the mail server. SSL and TLS encryption types are available for selection.
    -   Some Internet service providers require authentication on the incoming mail server before being allowed to send something. If this is your case, select the **Log on to incoming mail server** check box to enable a POP server and to set up its settings:
        -   **Incoming mail server (POP)** – enter the name of the POP server.
        -   **Port** – set the port of the POP server. By default, the port is set to **110**.
        -   **User name** – enter the user name
        -   **Password** – enter the password.
6.  Click **OK**.

## 10.5.3.1   Alert notifications

Acronis Backup & Recovery 11 provides the ability of notifying users about alerts by e-mail.

This option enables you to specify when and how often to receive notifications about the certain types of alerts.

The preset is: **Disabled**.

*Note: Before configuring alert notifications, specify the SMTP server settings in the E-mail notifications (p. 203).*

### To configure alert notifications

1.  Select the **Send e-mail notifications** check box.
2.  In the **E-mail addresses** field, type the e-mail address to which notifications will be sent. You can enter several addresses separated by semicolons.
3.  In the **Subject** field, type the notification subject or leave the default value.
4.  Select the necessary notification method:

    -   Per-alert – the notification will be sent as soon as a new alert occurs:

        Select the **As soon as an alert appears** check box.

        Click **Select the types of alerts...** to specify the types of alerts to receive notifications about.

    -   On schedule – the notification including all alerts that have occurred over a certain period of time. To receive notifications on schedule:

        Select the **On schedule** check box.

        Click **Select the types of alerts...** to specify the types of alerts to receive notifications about.

        Click **Notification schedule** to set up the notification frequency and time.
5.  Click **OK**.
6.  Click **Send test e-mail message** to check if the settings are correct.

## 10.5.4   Event tracing

It is possible to duplicate log events generated by the agent(s), operating on the managed machine, in the Application Event Log of Windows; or send the events to the specified SNMP managers. If you do not modify the event tracing options anywhere except for here, your settings will be effective for each local backup plan and each task created on the machine.

You can override the settings set here, exclusively for the events occurred during backup or during recovery (see Default backup and recovery options.) In this case, the settings set here will be effective for operations other than backup and recovery, such as archive validation or cleanup.

You can further override the settings set in the default backup and recovery options, when creating a backup plan or a recovery task. The settings you obtain in this case will be plan-specific or task-specific.

### 10.5.4.1   SNMP notifications

This option is effective for both Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option defines whether the agent(s) operating on the managed machine have to send the log events to the specified Simple Network Management Protocol (SNMP) managers. You can choose the types of events to be sent.

You can override the settings set here, exclusively for the events that occur during backup or during recovery, in the Default backup and recovery options. In this case, the settings set here will be effective for operations other than backup and recovery, such as archive validation or cleanup.

You can further override the settings set in the default backup and recovery options, when creating a backup plan or a recovery task. The settings you obtain in this case will be plan-specific or task-specific.

For detailed information about using SNMP with Acronis Backup & Recovery 11, please see "Support for SNMP (p. 28)".

The preset is: **Disabled**.

### *To set up sending SNMP messages*

1. Select the **Send messages to SNMP server** check box.
2. Specify the appropriate options as follows:
   - **Types of events to send** – choose the types of events: **All events**, **Errors and warnings**, or **Errors only**.
   - **Server name/IP** – type the name or IP address of the host running the SNMP management application, the messages will be sent to.
   - **Community** – type the name of the SNMP community to which both the host running SNMP management application and the sending machine belong. The typical community is "public".

   Click **Send test message** to check if the settings are correct.

To disable sending SNMP messages, clear the **Send messages to SNMP server** check box.

The messages are sent over UDP.

The next section contains additional information about Setting up SNMP services on the receiving machine (p. 205).

## 10.5.4.2    Setting up SNMP services on the receiving machine

### Windows

To install the SNMP service on a machine running Windows:

1. **Start > Control Panel > Add or Remove Programs > Add/Remove Windows Components.**
2. Select **Management and Monitoring Tools.**
3. Click **Details.**
4. Select the **Simple Network Management Protocol** check box.
5. Click **OK.**

You might be asked for lmmib2.dll that can be found on the installation disc of your operating system.

### Linux

To receive SNMP messages on a machine running Linux, the net-snmp (for RHEL and SUSE) or the snmpd (for Debian) package has to be installed.

SNMP can be configured using the **snmpconf** command. The default configuration files are located in the /etc/snmp directory:

- /etc/snmp/snmpd.conf - configuration file for the Net-SNMP SNMP agent
- /etc/snmp/snmptrapd.conf - configuration file for the Net-SNMP trap daemon.

### 10.5.4.3    Windows event log

This option is effective only in Windows operating systems.

This option is not available when operating under the bootable media.

This option defines whether the agent(s) operating on the managed machine have to log events in the Application Event Log of Windows (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

You can override the settings set here, exclusively for the events that occur during backup or during recovery, in the Default backup and recovery options. In this case, the settings set here will be effective for operations other than backup and recovery, such as archive validation or cleanup.

You can further override the settings set in the default backup and recovery options, when creating a backup plan or a recovery task. The settings you obtain in this case will be plan-specific or task-specific.

The preset is: **Disabled**.

To enable this option, select the **Log events** check box.

Use the **Types of events to log** check box to filter the events to be logged in the Application Event Log of Windows:

- **All events** - all events (information, warnings and errors)
- **Errors and warnings**
- **Errors only**.

To disable this option, clear the **Log events** check box.

## 10.5.5   Log cleanup rules

This option specifies how to clean up the Acronis Backup & Recovery 11 agent log.

This option defines the maximum size of the agent log folder (in Windows XP/2003 Server, %ALLUSERSPROFILE%\Application Data\Acronis\BackupAndRecovery\MMS\LogEvents).

The preset is: **Maximum log size: 50 MB**. **On cleanup, keep 95% of the maximum log size**.

When the option is enabled, the program compares the actual log size with the maximum size after every 100 log entries. Once the maximum log size is exceeded, the program deletes the oldest log entries. You can select the amount of log entries to retain. The default 95% setting will keep most of the log. With the minimum 1% setting, the log will be nearly cleared.

This parameter can also be set by using Acronis Administrative Template.

## 10.5.6   Online backup proxy

This option is effective only for backup to and recovery from Acronis Online Backup Storage over the Internet.

This option defines whether the Acronis agent will connect to the Internet through a proxy server.

*Note: Acronis Backup & Recovery Online supports only HTTP and HTTPS proxy servers.*

### To set up proxy server settings

1. Select the **Use a proxy server** check box.
2. In **Address**, specify the network name or IP address of the proxy server—for example: **proxy.example.com** or **192.168.0.1**
3. In **Port**, specify the port number of the proxy server—for example: **80**
4. If the proxy server requires authentication, specify the credentials in **User name** and **Password**.
5. To test the proxy server settings, click **Test connection**.

If you do not know the proxy server settings, contact your network administrator or Internet service provider for assistance.

Alternatively, you can try to take these settings from your Web browser's configuration. This is how to find them in three popular browsers.

- **Microsoft Internet Explorer.** On the **Tools** menu, click **Internet Options**. On the **Connections** tab, click **LAN settings**.
- **Mozilla Firefox.** On the **Tools** menu, click **Options** and then click **Advanced**. On the **Network** tab, under **Connection**, click **Settings**.
- **Google Chrome.** In **Options**, click **Under the Hood**. Under **Network**, click **Change proxy settings**.

# 11 Online backup

This section provides details about using the Acronis Backup & Recovery Online service. This service enables you to do online backups to Acronis Online Backup Storage.

*Acronis Backup & Recovery Online might be unavailable in your region. To find more information, click here: http://www.acronis.com/my/backup-recovery-online/*

To configure backup to the online storage or recovery from the storage, follow the regular steps described in the corresponding sections:

Creating a backup plan (p. 29)

Creating a centralized backup plan

Recovering data (p. 97)

The main difference is that you select the online storage as the backup destination.

Host-based backups of virtual machines are possible with Acronis Backup & Recovery 11 Virtual Edition. You can back up all virtual machines managed by Agent for ESX(i) or Agent for Hyper-V with a single subscription for virtual machines.

## 11.1  Introduction to Acronis Backup & Recovery Online

This section contains a brief overview of Acronis Backup & Recovery Online and answers questions that may arise during evaluation and usage of this product.

### 11.1.1  What is Acronis Backup & Recovery Online?

Acronis Backup & Recovery Online is a service that enables you to back up data to Acronis Online Backup Storage. To use this service, you need to buy a subscription that determines the amount of storage space reserved for your backups (storage quota) and how long the online service will be available to you.

A separate subscription is required for each physical machine you are going to back up. The server or workstation subscription that you buy will depend on the Windows operating system that the machine is running. Other operating systems are not supported by Acronis Backup & Recovery Online.

**Example:** A 250 GB/ 1 year workstation subscription means that you can back up data from a machine running a non-server Windows operating system for a period of 1 year. The backups can occupy no more than 250 GB.

### 11.1.2  What data can I back up and recover?

You can back up any files, volumes, or the entire physical machine as often as you wish. Unlike most online backup solutions, Acronis Backup & Recovery Online enables bare metal recovery directly from the online storage. Files can be recovered from disk-level backups as well as from file-level backups.

### 11.1.3 How long will my backups be kept in the online storage?

Your backups remain in the online storage until you delete them or until the subscription expires. Recovering data from the online storage is possible for 30 days following the subscription expiration date.

For effective use of the storage space, you have the option to set up the "**Delete backups older than**" retention rule.

**Example**

You might want to use the following backup strategy for a file server.

Back up the critical files twice a day on a schedule. Set the retention rule "**Delete backups older than**" 7 days. This means that after every backup the software will check for backups older than 7 days and delete them automatically.

Run backup of the server's system volume manually as required. For example, after the operating system updates. Manually delete the backups that you do not need.

### 11.1.4 How to secure my data?

Backups can be encrypted using the Advanced Encryption Standard (AES) cryptographic algorithm and the password you set. This guarantees that your data is not accessed by anyone else.

### 11.1.5 Backup and recovery FAQ

This section answers questions related to backup and recovery processes.

#### 11.1.5.1 What backup methods are available?

Full and incremental backup methods are available through the following backup schemes.

**Run now** (immediate start) or **Manual start** (postponed start). On the first task run a full backup is created. If you run the backup task again, an incremental backup will be created.

**Simple** (start on schedule). The first backup is full, the later backups are incremental. With this backup scheme, you can set up a retention rule to automatically delete old backups.

An additional backup scheme that is available only for online storage is **Initial seeding**. This is the **Run now** scheme using local destination and full backup method. To use this scheme, you need a license for the Initial Seeding (p. 210) service.

#### 11.1.5.2 Is the online storage available under Acronis bootable media?

Recovery from Acronis Online Backup Storage is available but backup to the storage is not.

#### 11.1.5.3 Can I use Acronis Universal Restore when recovering a system from the online storage?

Yes. Acronis Universal Restore is always available when recovering a system from the online storage. Using Acronis Universal Restore when recovering from other types of storage will require a separate license.

### 11.1.5.4 What if a network connection is lost during online backup or recovery?

The software will try to reach the online storage every 30 seconds. After five unsuccessful attempts the backup or recovery task will fail.

You can change the number of attempts and the interval between the attempts in the **Error handling > Re-attempt, if an error occurs** option. Every backup plan or recovery task includes this option.

### 11.1.5.5 What happens if I run out of space?

When a machine's backups are about to exceed the storage space allowed by its subscription, you receive an e-mail notification with an alert. In addition, you can see this alert on the account management web page near the machine. This means you have to free some space for future backups. Or, you can consider increasing the storage quota (p. 219). You may also want to set or edit the retention rule (p. 209) so that an overflow does not occur in future. Once the occupied space reaches the limit, the backups will cease to run.

### 11.1.5.6 What is the cleanup task for?

Any backup plan where the retention rule is set contains a cleanup task in addition to a backup task. The cleanup task checks the archive created by the backup plan for backups that have outlived their lifetime. If such backups are found, the task makes the online storage delete them. Since the deletion is performed on the online storage side, it does not take your machine's CPU resource.

The cleanup task runs after every online backup, even if the backup has failed. The last successful backup is always kept though. For more information about the retention rule please refer to "How long will my backups be kept in the online storage? (p. 209)"

Normally, there is no need to start and stop the cleanup task manually. But it is possible to do so in the **Backup plans and tasks** view.

### 11.1.5.7 How to make a recovered machine recognize its subscription?

When you recover a physical machine from a backup, a new machine identifier is created. Therefore, the machine is not able to back up to the subscription it used before recovery.

To continue backing up the machine to the same subscription, reassign (p. 222) the subscription to the machine. If you do this, the next machine's backup can be incremental. If you assign a new subscription to the machine, the software will have to do a new full backup.

## 11.1.6 Initial Seeding FAQ

This section explains what Initial Seeding is, why you would want to use it and provides some usage details.

### 11.1.6.1 What is Initial Seeding?

Initial Seeding is an extra service that lets you save an initial full backup locally and then send it to Acronis on a hard disk drive.

Acronis uploads the backup to the online storage. After that, you can add incremental backups to this full backup, either manually or on a schedule.

The hard disk drive is sent back to you but it is not possible to recover from it. However, recovery from a locally attached device is possible with the Large scale recovery (p. 215) option.

### 11.1.6.2    Why would I want to use Initial Seeding?

This service helps you save time and network traffic during the initial full backup. It is useful when backing up very large volumes of data or entire machines to the online storage.

### 11.1.6.3    Is Initial Seeding a paid service?

Yes, you need to buy one Initial Seeding license per machine.

### 11.1.6.4    What types of hard drive can I use for Initial Seeding?

Acronis accepts hard disk drives of the following interface types: IDE, ATA, SATA, USB connected drives. SCSI drives are not accepted.

You can back up directly to the device or back up to a local or network folder and then copy the backup to the device. Make sure that the device has only one volume and that the file system on that volume is NTFS or FAT32.

### 11.1.6.5    Can I send more than one backup under a single Initial Seeding license?

No. An Initial Seeding license allows you to create only one backup on the machine.

However, if you have made a mistake or have decided to create another backup for any reason, you can cancel the initial seeding order. As a result, the license will become available again.

### 11.1.6.6    Can I send backups taken from a number of machines on a single hard drive?

Yes. However, the number of required licenses is still one per machine.

### 11.1.6.7    How to buy an Initial Seeding license?

You can buy an Initial Seeding license from an Acronis partner or in Acronis online store. Follow the link www.acronis.com/my/backup-recovery-online/#buy to locate a partner or to buy online.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Click **Enter new registration code** on the same web page and register the license. The license becomes available on the **Initial Seeding / Recovery** tab.

A license purchased in Acronis online store becomes available immediately after the payment is processed.

### 11.1.6.8    How to perform initial seeding?

1.  Make sure you have activated an Acronis Backup & Recovery Online subscription on the machine where you will do initial seeding. If you are currently using a trial subscription, make sure that you also have a paid subscription available and assigned to this machine. Do not use the Initial Seeding service if you do not have a paid subscription.

2. Decide on the media (p. 211) that you will send.

3. Attach the media to the machine you are going to back up. Alternatively, you can back up to a local or network folder and then copy/move the backup to the media.

4. Start Acronis Backup & Recovery 11, click **Create backup plan** and create a backup plan on this machine:

   - Under **What to back up**, select disks, volumes or files/folders you want to back up.

   - Under **Where to back up**, specify **Online Backup Storage**.

   - In **Backup scheme**, select **Initial seeding**. Specify the said media as the backup destination.

   - [Optional, but strongly recommended] Enable backup encryption in **Backup options** > **Archive protection**.

   The backup starts immediately after you click the final **OK**.

5. [Optional] If you want to add backups from another machine, attach the media to that machine and perform the same steps. You need a separate Initial Seeding license for each machine that you want to back up.

6. Package (p. 212) the media along with a prepaid return shipping label and send it to Acronis by physical mail. The address is available on your account management Web page > **Initial Seeding / Recovery** tab > **Orders in process** > **Initial Seeding orders** > **Datacenter address**.

7. On the same Web page, mark the order as "shipped" and track (p. 214) the order status.

8. Once you observe that the backup has been uploaded on the online storage, you can edit the backup plan to do incremental backups:

   - In **Backup scheme**, select **Run now**, **Manual start** for manual backups or **Simple** for scheduled backups.

   - For **Simple**, specify the schedule and (optionally) the retention rule.

   - Click **Save**.

   When started manually or on schedule, your backup plan will add incremental backups to the initial backup stored in the online storage.

## 11.1.6.9   How to package a hard drive for shipment?

It is very important that your hard drive be packaged carefully. Careful packaging will protect your drive from any damage during shipment.

**Hard drive types**

Acronis accepts hard disk drives of the following interface types: IDE, ATA, SATA, USB connected drives.

*SCSI drives are not accepted.*

**Packaging**

If possible, use the original packaging. Otherwise, packaging materials can be obtained at any shipping outlet or stationary store. You should also include all necessary cables or adapters to the drive. Acronis will not be able to process your initial seeding request if there are no cables included.

The following are instructions about how to package your hard disk drive.

## Step 1

Delicately remove your hard disk drive from the machine.



## Step 2

Place the hard drive into an anti-static bag to protect the drive from electrostatic discharge. If you do not have an anti-static bag, simply wrap the hard drive into aluminum foil.



## Step 3

Use a sturdy box that is at least twice the size of the drive. Pack the drive with a bubble wrap around all 6 sides so it can fit tight into the box and cannot be moved within.

*DO NOT* use Styrofoam **peanuts** for packing as they do not provide enough protection.
*DO NOT* send your media in **jiffy** bags

*Step 4*

Choose the transport company that you will use for shipping. On that company's web site, prepare and print two prepaid shipping labels:

1. **Shipping label for sending** your hard drive. This label is placed on the top of the box. You should send your package to one of the Acronis data centers. The data center address can be obtained on the **Initial seeding / Recovery** tab of your account management page by clicking **Datacenter address**.

   We recommend that you use overnight shipping, if you want to start doing incremental backups as soon as possible. The data is generally available the next business day after the data center receives it.

2. **Shipping label for returning** your hard drive. Put this label in the box. When returning your hard drive, we will reuse the same packaging unless it is damaged. If you do not enclose the label, your drive will be **securely discarded**.

   You might want to use the most cost-efficient delivery method for having your hard drive returned.



*Step 5*

Securely seal the box with a sturdy tape. Then, stick the **shipping label for sending** your hard drive to the top of the box, so the label does not wrap around the edge of the package.



## 11.1.6.10  How do I track an Initial Seeding order status?

On the Acronis Web site, the **Initial Seeding / Recovery** tab shows you the status of all your orders. In addition, you will receive e-mail notifications about the most important events.

- **Available** – The license is available for using on any machine.

- **An order was created** – The backup is about to start and the license cannot be used for the same or any other machine. From this point on, you can cancel the order if something goes wrong. This will return the license to the pool of available licenses.

- **A full backup has started** – This status is set when the first backup starts. The order start time occurs at this moment.

- **A full backup has been successfully completed** – The backup has been completed and the order is ready to ship. You can now ship the media:

  **Step 1**. Package the media following the drive packaging and shipment instructions (p. 212) to avoid damage during shipment. If you want the media to be returned to you after the data is uploaded, prepare a prepaid return shipping label and place it inside the package together with the drive.

  **Step 2**. Send the drive via your preferred carrier to the Acronis datacenter.

  **Step 3**. Let us know when you have shipped the order by marking your order as "shipped".

  You will receive a notification message when Acronis receives the order and when the order is completed. If necessary, Acronis may contact you during order processing.

- [Occasional] **Backup creation error** – An error occurred when backing up. Please check the backup plan parameters and try again.

- **The media has been shipped** – This status is set after you mark the order as "shipped".

- **The media has been received by Acronis** – Acronis has started processing your order. From this point on, you cannot cancel the order. Creating a new initial seeding backup will require a new Initial Seeding license.

- **The data upload has started** – The process of uploading data to Acronis Online Backup Storage has started.

- **The data upload has been completed** – The initial full backup has been successfully uploaded to the online storage. You can do incremental online backups now.

- **The order has been completed. The media has been returned** (or: **Returning the media was not requested**) – Your media has been shipped back (the carrier and the tracking number are specified). If a prepaid shipping label was not provided with the media, the media will be discarded.

- [Occasional] **The order is on hold** – Your order was placed on hold due to technical difficulties processing the order. Acronis is working on resolving these issues.

- [Occasional] **The order has been cancelled** – The order had been cancelled before the media was shipped, so returning the media is not required.

- [Occasional] **The order has been cancelled. The media has been returned** (or: **Returning the media was not requested**) – The order was cancelled while the media was in the datacenter. The media has been shipped back (the carrier and the tracking number are specified). If a prepaid shipping label was not provided with the media, the media will be discarded.

## 11.1.7    Large Scale Recovery FAQ

This section explains what Large Scale Recovery is, why you would want to use it and provides some usage details.

### 11.1.7.1    What is Large Scale Recovery?

Large Scale Recovery is an extra service that enables you to obtain a copy of the backups you have in the online storage. You can then recover data from this copy.

Once you order Large Scale Recovery for a particular machine, Acronis sends you a USB hard disk drive with all of the backups made from this machine. You can recover data directly from the disk or copy the backups to a local or network folder.

### 11.1.7.2 Why would I use Large Scale Recovery?

In the event of a disaster or the need to recover large volumes of data or the entire machines quickly, this service helps you save time and network traffic. Recovering hundreds of gigabytes over the Internet may take days. This process will deliver a faster recovery.

### 11.1.7.3 Do I need to perform initial seeding to be able to use Large Scale Recovery?

No, these services are independent.

### 11.1.7.4 Is Large Scale Recovery a paid service?

Yes, you need to buy one Large Scale Recovery license per machine. The license enables you to get a disk with all of the currently available backups of this machine. To obtain backups that will be created in the future, you will need a new Large Scale Recovery license.

### 11.1.7.5 Can I perform large scale recovery on a different machine?

Yes. You can recover the data an unlimited number of times on any machine you wish. Acronis Universal Restore is included to help you recover an operating system to dissimilar hardware.

### 11.1.7.6 Can I obtain backups taken from a number of machines on a single hard drive?

No. A separate hard drive is required for each machine.

### 11.1.7.7 How to buy a Large Scale Recovery license?

You can buy a Large Scale Recovery license from an Acronis partner or in Acronis online store. Follow the link www.acronis.com/my/backup-recovery-online/#buy to locate a partner or to buy online.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Click **Enter new registration code** on the same web page and register the license. The license becomes available on the **Initial Seeding / Recovery** tab.

A license purchased in Acronis online store becomes available immediately after the payment is processed.

### 11.1.7.8 How do I track a Large Scale Recovery order status?

On the Acronis Web site, the **Initial Seeding / Recovery** tab shows you the status of all your orders. In addition, you will receive e-mail notifications about most important events.

- **Available** – The license can be used for any machine.
- **An order was created** – This status is set upon completion of the Large Scale Recovery order form. This means that the license cannot be used for any other machine. From this point on, you

can cancel the order if something goes wrong. This will return the license to the pool of available licenses.

- **The order is being processed** - Order processing in the datacenter started.
- **Writing data** – Your backups are being written onto the media. From this point on, you cannot cancel the order.
- **Writing data has been completed** – Your backups have been successfully written to the media.
- **Ready to ship the media** – Your order has been processed and the media will be shipped shortly.
- **The order has been completed. The media has been shipped** – The media has been shipped to you (the carrier and the tracking number are specified).
- [Occasional] **The order is on hold** – Your order was placed on hold due to technical difficulties processing the order. Acronis is working on resolving these issues.
- [Occasional] **The order has been cancelled** – The order has been cancelled.
- [Occasional] **Address is undeliverable** – Acronis cannot send the disk. On the same Web page, click **Change my delivery address** and specify the correct address for the order.
- [Occasional] **Address has been updated** – This status is set after you have updated the delivery address on Acronis web site.

### 11.1.7.9   How to perform large scale recovery?

The recovery procedure is the same as when recovering from the online storage. Just specify the path to the location where your backups are. For detailed information about recovery please refer to the context-sensitive help.

## 11.1.8   Subscription lifecycle FAQ

This section explains a subscription lifecycle and subscription operations that you can perform on your account management Web page.

### 11.1.8.1   How to access my account management Web page?

To access this Web page from the Acronis Web site:

1. Select **User Login**.
2. Log in to your account (create one if you are not registered yet).
3. Navigate to **Online backup** > **for Business**.

To access this Web page from Acronis Backup & Recovery 11:

1. On the **Actions** menu, click **Back up now** or **Create backup plan**.
2. Click **Location**, and then click **Buy or manage your subscriptions**.
3. Log in to your account (create one if you are not registered yet).

### 11.1.8.2   Where do I find the subscriptions that I purchased?

If you purchased your subscriptions from an Acronis partner, you should have received an e-mail confirming the registration codes for each subscription. Create an account on the Acronis web site, if you do not have one already, and log in to it. Navigate to **Online backup** > **For Business**. This is your *account management Web page*. Click **Enter new registration code** and enter the registration codes. The subscriptions will appear in the list of available subscriptions under the **Manage Subscriptions** tab.

If you purchased your subscriptions online using the Acronis Web site, they are available immediately on your account management Web page. The newly obtained subscriptions are listed at the **Manage Subscriptions** tab.

## 11.1.8.3    When does my subscription begin?

Your subscription begins when you choose to have it begin, not at the time of purchase.

The time count will begin as soon as the subscription is activated. The first activation occurs when you assign a subscription to a certain machine. To do so, you need to have Acronis software installed.

## 11.1.8.4    What happens when my subscription expires?

A month before the subscription expiration date you receive an e-mail notification with an alert. In addition, you can see this alert on the account management web page near the machine. This means you need to renew (p. 218) the subscription to continue backing up the machine.

If you do not renew the subscription, you will be able to back up data to the online storage for five days following the expiration date. You will be able to recover data from the online storage for 30 days following the expiration date.

## 11.1.8.5    How do I renew a subscription?

Buy another subscription and specify it as the next subscription of the same machine. The new subscription will be activated as soon as the current subscription expires.

An expired subscription can be renewed within five days after expiration. In such cases, the new subscription will be activated immediately.

### Renewing a single subscription

***To renew a subscription***

1.    Go to the account management Web page.
2.    Make sure that you have an available subscription with the *same* storage quota.
3.    Select the machine that you want to renew the subscription for, and then click **Renew**.

The subscription appears in the **Next subscription** column for the selected machine.

### Renewing a number of activated subscriptions at once

This operation is possible if the appropriate number of new subscriptions are identical to the currently used subscriptions.

Make sure the new subscriptions are available on your account management Web page. Then click **Renew all**. The confirmation window will summarize which subscriptions will be renewed. If identical subscriptions are not found for some of the machines, you have the option to cancel automatic renewal and renew each subscription individually.

## What does "Auto-renew" mean?

Auto-renewal means that when the current subscription expires, the next subscription will be automatically selected from the available subscriptions. The next subscription must be identical to the current subscription.

If an identical subscription is not found, auto-renewal will not take place and your backups may fail. No subscriptions will be bought automatically. Only those subscriptions available at the time of auto-renewal can be used. You can select auto-renewal for each individual subscription or set up bulk auto-renewal of all of the activated subscriptions you have.

### 11.1.8.6    How do I increase a storage quota for a machine?

Replace the subscription assigned to the machine with another subscription that has a larger storage quota. The new subscription has the same type as the old one. For example, you can replace a workstation subscription only with another workstation subscription with a larger quota.

The storage quota increase is free of charge. After the completion, the operation cannot be reverted.

The new subscription will have a lesser subscription period. It is calculated in the following way:

$$Pn = Po * (Qo / Qn)$$

where

- Pn – new subscription remaining period
- Po – old subscription remaining period
- Qo – old subscription storage quota
- Qn – new subscription storage quota.

**Example**. You have a 250 GB subscription and you decide to increase its storage quota 2 months before the subscription expiration. The new subscription storage quota is 500 GB. So, the new subscription remaining period will be

$$Pn = 2 \text{ months} * (250 \text{ GB} / 500 \text{ GB}) = 1 \text{ month}$$

***To increase a storage quota***

1. Go to the account management Web page.
2. Select the machine for which you want to increase a storage quota, and then click **Increase**.
3. If the existing subscription has the maximum available storage quota for this subscription type, the software notifies you. Otherwise, select the new storage quota.
4. Click **Increase**, and then click **OK** to confirm the operation.

### 11.1.8.7    What is the "Group" column for?

So you can apply actions, such as **Renew all** or **Auto-renew all**, to your selection of the subscriptions. Specify the desired group name, for example, SalesDept, near each of the subscriptions you want to group. Click the **Group** column header to sort the subscriptions and then apply the desired actions to the group.

### 11.1.8.8 Can I revoke a subscription from a machine?

You cannot return an activated subscription to the list of available subscriptions, but you can reassign (p. 222) it to a different machine in Acronis Backup & Recovery 11 GUI.

### 11.1.8.9 Can I cancel my subscription?

Just wait until the subscription expires. Refunds are not available for the online backup subscriptions.

# 11.2 Where do I start?

On the Acronis Web site, log in to your account (create one if you are not registered yet) and navigate to **Online backup** > **For Business**. This is your *account management Web page*. Here you can get a trial subscription, locate an Acronis partner or buy subscriptions online. The newly obtained subscriptions are listed as available subscriptions in the **Manage Subscriptions** tab.

If you purchased your subscriptions from an Acronis partner, register them manually using the **Enter new registration code** link. The registration codes come with the purchase confirmation e-mail.

Next, install Acronis software (if not yet installed) and assign (p. 221) each subscription to a machine. The subscriptions become activated. After that, you can start backing up to Acronis Online Backup Storage.

# 11.3 Choosing a subscription

Normally, you choose a subscription based on the operating system your machine is running.

**Server operating systems supported by Acronis Backup & Recovery Online:**

- Windows 2000 SP4 – all editions except for the Datacenter and Professional editions.
- Windows Server 2003/2003 R2 – the Standard, Enterprise editions (x86, x64).
- Windows Small Business Server 2003/2003 R2.
- Windows Server 2008 – the Standard, Enterprise editions (x86, x64).
- Windows Small Business Server 2008.
- Windows Server 2008 R2 – the Standard, Enterprise, Datacenter, Foundation editions.
- Windows MultiPoint Server 2010.
- Windows Small Business Server 2011.

**Workstation operating systems supported by Acronis Backup & Recovery Online:**

- Windows 2000 Professional SP4.
- Windows XP Professional SP2+ (x86, x64).
- Windows Vista - all editions except for Vista Home Basic and Vista Home Premium (x86, x64).
- Windows 7 - all editions except for the Starter and Home editions (x86, x64).

If your backups are likely to exceed the storage quota for this type of subscription, you may want to use a subscription with larger storage quota. For example, you can use a server subscription on a workstation. The inverse usage is not possible. You cannot back up a server using a workstation subscription.

**Trial subscriptions**

You can get one free workstation or server subscription per account. The storage quota of the trial subscription is equal to that of the standard subscription. The subscription period is limited to 2 months.

Obtaining a trial subscription is possible until you buy a paid subscription. You can use a trial subscription along with paid ones. The same expiration rules apply to trial and paid subscriptions.

To continue using the service after the trial subscription expires, buy the same type subscription and renew the trial subscription specifying the paid subscription. Your backed up data will be kept online. Regular backups of your machines will continue uninterrupted. The service will not need to perform a new full backup.

To get a trial subscription, do either of the following:

- Go to the account management Web page, click **Try Now Free for 60 Days**, and select the necessary subscription type.
- Install Acronis Backup & Recovery 11 on the machine you want to back up, start the product, click **Back up now** or **Create backup plan**, click **Location**, and click **Get trial subscription**. Log in to your account (create one if you are not registered yet). A trial subscription will be automatically created and assigned to the machine.

# 11.4 Activating online backup subscriptions

To be able to back up a machine to the online storage, you need to purchase and activate a subscription to the Acronis Backup & Recovery Online service. You can purchase subscriptions on the Acronis Web site or from an Acronis reseller.

Before activating a subscription, please take into account the following considerations:

- As soon as a subscription is activated, its subscription period starts. To avoid losing subscription time, activate the subscription only when you are ready to back up the machine.
- If a machine already has a subscription, the new subscription will replace the old one. You can reassign the old subscription to a different machine—see "Reassigning an activated subscription" later in this section.

## 11.4.1 Activating subscriptions

To begin with, make sure that the machines whose subscriptions you want to activate are registered on the management server and available (turned on).

***To activate subscriptions***

1. Connect the console to the management server.
2. In the **Actions** pane, click **Activate online backup subscriptions**.
3. Specify the credentials to log in to the online storage.
4. Select the machine and then click **Select subscription**.
5. From **Available subscriptions**, select the subscription that you want to activate for the machine.
6. Click **Activate now**.
7. Perform the previous three steps for each machine for which you want to activate a subscription.

Alternatively, you can activate a subscription when the console is connected to a machine instead of the management server.

## 11.4.2 Activating subscriptions

***To activate a subscription***

1. Start Acronis Backup & Recovery 11.
2. In the **Actions** pane, click **Activate online backup subscription**.
3. Specify the credentials to log in to the online storage.
4. From **Available subscriptions**, select the subscription that you want to activate for the machine.
5. Click **Activate now**.

## 11.4.3 Reassigning an activated subscription

Sometimes you may want to use an already activated subscription instead of an available subscription. In these cases, for example:

- You no longer need to back up one of your machines and you want to reuse that machine's subscription for another machine.
- You reinstalled Acronis Backup & Recovery 11 on a machine and want to resume its online backups.
- You recovered a machine to bare metal (or to a state when it did not yet have an activated subscription) and want to resume its online backups.

Reassigning a subscription does not restart its subscription period.

***To assign an activated subscription to a machine***

1. On the machine to which you want to assign an activated subscription, go to the subscription activation window.
2. From **Activated subscriptions**, select the activated subscription that you want to reassign to the machine.
3. Click **Activate now**.

**Example**

The diagram below shows what happens if you reassign a subscription to a different machine. Let's assume Machine 1 has four backups in Subscription 1. Machine 2 has two backups in Subscription 2. At that point, you reassign Subscription 1 to Machine 2. Machine 2 does its third backup to Subscription 1.

Depending on your settings, this backup will be either full or incremental. But its size is not likely to be less than a full backup size. Therefore, it is not practical to reassign a subscription to a machine whose first backup was done as an initial seeding. You will need to either redo the initial seeding (which requires a new license) or to transfer the sizeable backup over the Internet.



All earlier created backups remain intact. You can delete them manually if necessary. Keep in mind though, backups can be deleted from a subscription only by the machine to which the subscription is assigned. In our example, you have the following options.

**Before reassigning**

Delete backups from Subscription 1 using Machine 1 (if it is available and turned on). Delete backups from Subscription 2 using Machine 2.

**After reassigning**

Delete backups from Subscription 1 using Machine 2. You cannot delete backups from Subscription 2, unless you assign this subscription to any machine.

# 11.5 Configuring proxy settings

If the machine connects to the Internet through a proxy server, configure Acronis Backup & Recovery 11 to use the proxy server.

1. Start Acronis Backup & Recovery 11.
2. On the toolbar, on the **Options** menu, click **Machine options**.
3. Click **Online backup proxy**.
4. Enter the proxy server settings. For detailed information about the settings please refer to the context-sensitive help.

# 11.6 Limitations of the online storage

Unlike other types of storage available in Acronis Backup & Recovery 11, the online storage has the following limitations.

## Operations

The following operations are not possible.

**Backup operations:**

- Backing up from bootable media
- Backing up under Linux
- Creating differential backups
- Using the **Grandfather-Father-Son (GFS)**, **Tower of Hanoi**, and **Custom** backup schemes
- Simplified naming of backup files
- Converting a backup to a virtual machine

**Recovery operation:**

- Recovering a backup as a virtual machine

**Operations with backups:**

- Exporting a backup
- Mounting a backup
- Copying or moving backups to the online storage

**Operation with archives** (an archive is a set of backups):

- Exporting an archive

These limitations also apply to backing up data using Initial Seeding and to recovering data using Large Scale Recovery.

## Backup and recovery options

Some backup and recovery options are not supported by online backups. For example, **Backup splitting** (p. 81).

By using the **Backup performance > Network connection speed** option, you can vary the transferring speed as kilobytes per second, but not as a percentage.

# 11.7 Terminology reference

The following is the list of terms related to the Acronis Backup & Recovery Online service.

## Activate a subscription

Allow the machine to use the online storage according to the subscription. Subscription period starts counting down when the subscription is activated.

## Activated subscription

A subscription that is currently being used by a machine.

**Assign a subscription to a machine**

Reserve a subscription for a particular machine. Subscription period does not start counting down until the subscription is activated.

**Assigned subscription**

A subscription that has been assigned to a machine.

**Available subscription**

A subscription that is not assigned to any machine.

**Extra service**

A service that you can use in addition to online backup subscriptions.

**Increase storage quota**

Replace a subscription with another one that has a greater storage quota. The remaining subscription period is reduced in proportion to the capacity increase.

**Initial Seeding**

An extra service that enables you to save an initial full backup locally and then send it to Acronis on a hard disk drive. Acronis uploads the backup to the online storage. After that, you can add incremental backups to this full backup, either manually or on a schedule.

**Large Scale Recovery**

An extra service that enables you to obtain a copy of the backups you have in the online storage. You can then recover data from this copy.

**License**

*Not to be confused with Acronis Backup & Recovery 11 product license.*

Permission for a machine to use an extra service of Acronis Backup & Recovery Online.

You can buy Initial Seeding licenses and/or Large Scale Recovery licenses.

**Reassign a subscription**

Assign a subscription that is already activated, to a different machine.

**Registration code**

A character string for registering a subscription or license that was bought from an Acronis partner.

When you purchase such subscriptions or licenses, you receive a confirmation e-mail containing the registration codes for each of them. You then enter the registration codes on the account management Web page, and these subscriptions and licenses become available for use.

**Renew a subscription**

Assign a subscription that has the same storage quota as the current, activated subscription.

This subscription will become activated as soon as the current subscription expires.

**Storage quota**

The amount of storage space that a machine can use according to the subscription.

**Subscription**

Permission for a machine to use a specific amount of space in the online storage for a specific period of time.

**Subscription period**

The period during which the subscription remains activated. You can back up and recover the machine during this period. Recovery is possible for extra 30 days after this period ends.

**Unassign a subscription**

Make an assigned subscription available again.

You can unassign a subscription as long as it is not activated.

# 12 Glossary

## A

### Acronis Active Restore

The Acronis proprietary technology that brings a system online immediately after the system recovery is started. The system boots from the backup (p. 233) and the machine becomes operational and ready to provide necessary services. The data required to serve incoming requests is recovered with the highest priority; everything else is recovered in the background. Limitations:

- the backup must be located on the local drive (any device available through the BIOS except for network boot)
- does not work with Linux images.

### Acronis Plug-in for WinPE

A modification of Acronis Backup & Recovery 11 Agent for Windows that can run in the preinstallation environment. The plug-in can be added to a WinPE (p. 240) image using Bootable Media Builder. The resulting bootable media (p. 229) can be used to boot any PC-compatible machine and perform, with certain limitations, most of the direct management (p. 232) operations without the help of an operating system. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 231).

### Acronis Secure Zone

A secure volume for storing backup archives (p. 228) within a managed machine (p. 237). Advantages:

- enables recovery of a disk to the same disk where the disk's backup resides
- offers a cost-effective and handy method for protecting data from software malfunction, virus attack, operator error
- eliminates the need for a separate media or network connection to back up or recover the data. This is especially useful for mobile users
- can serve as the primary location from which backups are replicated further.

Limitation: Acronis Secure Zone cannot be organized on a dynamic disk (p. 233).

Acronis Secure Zone is considered as a personal vault (p. 237).

### Acronis Startup Recovery Manager (ASRM)

A modification of the bootable agent (p. 229), residing on the system disk and configured to start at boot time when F11 is pressed. Acronis Startup Recovery Manager eliminates the need for rescue media or network connection to start the bootable rescue utility.

Acronis Startup Recovery Manager is especially useful for mobile users. If a failure occurs, the user reboots the machine, hits F11 on prompt "Press F11 for Acronis Startup Recovery Manager…" and performs data recovery in the same way as with ordinary bootable media.

Limitation: requires re-activation of loaders other than Windows loaders and GRUB.

## Acronis Universal Restore

The Acronis proprietary technology that helps boot up Windows or Linux on dissimilar hardware or a virtual machine. Universal Restore handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

Universal Restore is not available:

- when the image being recovered is located in Acronis Secure Zone (p. 227) or
- when using Acronis Active Restore (p. 227),

because these features are primarily meant for instant data recovery on the same machine.

## Activity

An action performed by Acronis Backup & Recovery 11 for achievement of some user goal. Examples: backing up, recovery, exporting a backup, cataloging a vault. An activity may be initiated by a user or by the software itself. Execution of a task (p. 239) always causes one or more activities.

## Agent (Acronis Backup & Recovery 11 Agent)

An application that performs data backup and recovery and enables other management operations on the machine (p. 236), such as task management and operations with hard disks.

The type of data that can be backed up depends on the agent type. Acronis Backup & Recovery 11 includes the agents for backing up disks and files and the agents for backing up virtual machines residing on virtualization servers.

## Archive

See Backup archive (p. 228).

# B

## Backup

A backup is the result of a single backup operation (p. 228). Physically, it is a file or a tape record that contains a copy of the backed up data as of a specific date and time. Backup files created by Acronis Backup & Recovery 11 have a TIB extension. The TIB files which are the result of a backup export (p. 235) or consolidation (p. 231) are also called backups.

## Backup archive (Archive)

A set of backups (p. 228) created and managed by a backup plan (p. 229). An archive can contain multiple full backups (p. 235) as well as incremental (p. 236) and differential backups (p. 232). Backups belonging to the same archive are always stored in the same location. If the backup plan includes replication (p. 238) or moving of backups to multiple locations, the backups in each location form a separate archive.

## Backup operation

An operation that creates a copy of the data that exists on a machine's (p. 236) hard disk for the purpose of recovering or reverting the data to a specified date and time.

## Backup options

Configuration parameters of a backup operation (p. 228), such as pre/post backup commands, maximum network bandwidth allotted for the backup stream or data compression level. Backup options are a part of a backup plan (p. 229).

## Backup plan (Plan)

A set of rules that specify how the given data will be protected on a given machine. A backup plan specifies:

- what data to back up
- the backup archive (p. 228) name and location
- the backup scheme (p. 229). This includes the backup schedule and [optionally] the retention rules (p. 238)
- [optionally] additional operations to perform with the backups (replication (p. 238), validation (p. 239), conversion to a virtual machine)
- the backup options (p. 228).

For example, a backup plan can contain the following information:

- back up volume C: **(this is the data the plan will protect)**
- name the archive MySystemVolume and place it in \\server\backups\ **(this is the backup archive name and location)**
- perform a full backup monthly on the last day of the month at 10:00AM and an incremental backup on Sundays at 10:00PM. Delete backups that are older than 3 months **(this is a backup scheme)**
- validate the last backup immediately after its creation **(this is a validation rule)**
- protect the archive with a password **(this is an option)**.

Physically, a backup plan is a bundle of tasks (p. 239) executed on a managed machine (p. 237).

A backup plan can be created directly on the machine, imported from another machine (local plan) or propagated to the machine from the management server (centralized plan (p. 230)).

## Backup scheme

A part of the backup plan (p. 229) that includes the backup schedule and [optionally] the retention rules and the cleanup (p. 231) schedule. For example, perform a full backup (p. 235) monthly on the last day of the month at 10:00AM and an incremental backup (p. 236) on Sundays at 10:00PM. Delete backups that are older than 3 months. Check for such backups every time the backup operation is completed.

Acronis Backup & Recovery 11 provides the ability to use well-known optimized backup schemes such as GFS and Tower of Hanoi, to create a custom backup scheme or to back up data once.

## Bootable agent

A bootable rescue utility that includes most of the functionality of the Acronis Backup & Recovery 11 Agent (p. 228). Bootable agent is based on Linux kernel. A machine (p. 236) can be booted into a bootable agent using either bootable media (p. 229) or Acronis PXE Server. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 231).

# Bootable media

A physical media (CD, DVD, USB flash drive or other media supported by a machine (p. 236) as a boot device) that contains the bootable agent (p. 229) or Windows Preinstallation Environment (WinPE) (p. 240) with the Acronis Plug-in for WinPE (p. 227). A machine can also be booted into the above environments using the network boot from Acronis PXE Server or Windows Deployment Service (WDS). These servers with uploaded bootable components can also be thought of as a kind of bootable media.

Bootable media is most often used to:

- recover an operating system that cannot start
- access and back up the data that has survived in a corrupted system
- deploy an operating system on bare metal
- create basic or dynamic volumes (p. 234) on bare metal
- back up sector-by-sector a disk that has an unsupported file system
- back up offline any data that cannot be backed up online because of restricted access, being permanently locked by the running applications or for any other reason.

# Built-in group

A group of machines permanently located on a management server (p. 237).

Built-in groups cannot be deleted, moved to other groups or manually modified. Custom groups cannot be created within built-in groups. There is no way to remove a machine from the built-in group except by removing the machine from the management server.

# C

# Cataloging

Cataloging a backup (p. 228) adds the contents of the backup to the data catalog (p. 232). Backups are cataloged automatically by the agent (p. 228) as soon as they are created. A user has the option to turn off automatic cataloging and start it manually when appropriate. Backups that are stored on a storage node (p. 239) will be cataloged by the node in this case.

# Centralized backup plan

A backup plan (p. 229) that is deployed to a managed machine (p. 237) from the management server (p. 237). Such plan can be modified only by editing the original backup plan on the management server.

# Centralized management

Management of the Acronis Backup & Recovery 11 infrastructure through a central management unit known as Acronis Backup & Recovery 11 Management Server (p. 237). The centralized management operations include:

- creating centralized backup plans (p. 230) for the registered machines (p. 238) and groups of machines
- creating and managing static (p. 238) and dynamic groups (p. 234) of machines (p. 236)
- managing the tasks (p. 239) existing on the machines

- creating and managing centralized vaults (p. 231) for storing archives
- managing storage nodes (p. 239)
- monitoring activities of the Acronis Backup & Recovery 11 components, creating reports, viewing the centralized log and more.

## Centralized task

A task (p. 239) propagated to a machine from the management server (p. 237). Such task can be modified only by editing the original task or centralized backup plan (p. 230) on the management server.

## Centralized vault

A networked location allotted by the management server (p. 237) administrator to serve as storage for the backup archives (p. 228). A centralized vault can be managed by a storage node (p. 239) or be unmanaged. The total number and size of archives stored in a centralized vault are limited by the storage size only.

As soon as the management server administrator creates a centralized vault, the vault name and path to the vault are distributed to all machines registered (p. 238) on the server. The shortcut to the vault appears on the machines in the **Vaults** list. Any backup plan (p. 229) existing on the machines, including local plans, can use the centralized vault.

On a machine that is not registered on the management server, a user having the privilege to back up to the centralized vault can do so by specifying the full path to the vault. If the vault is managed, the user's archives will be managed by the storage node as well as other archives stored in the vault.

## Cleanup

Deleting backups (p. 228) from a backup archive (p. 228) or moving them to a different location in order to get rid of outdated backups or prevent the archive from exceeding the desired size.

Cleanup consists ofapplying retention rules (p. 238) to an archive. The retention rules are set by the backup plan (p. 229) that produces the archive. Cleanup may or may not result in deleting or moving backups depending on whether the retention rules are violated or not.

## Console (Acronis Backup & Recovery 11 Management Console)

A tool for remote or local access to Acronis agents (p. 228) and Acronis Backup & Recovery 11 Management Server (p. 237).

Having connected the console to the management server, the administrator sets up centralized backup plans (p. 230) and accesses other management server functionality, that is, performs centralized management (p. 230). Using the direct console-agent connection, the administrator performs direct management (p. 232).

## Consolidation

Combining two or more subsequent backups (p. 228) belonging to the same archive (p. 228) into a single backup.

Consolidation might be needed when deleting backups, either manually or during cleanup (p. 231). For example, the retention rules require to delete a full backup (p. 235) that has expired but retain the next incremental (p. 236) one. The backups will be combined into a single full backup which will be dated with the incremental backup's date. Since consolidation may take a lot of time and system resources, retention rules provide an option to not delete backups with dependencies. In our example, the full backup will be retained until the incremental one also becomes obsolete. Then both backups will be deleted.

# D

## Data catalog

Allows a user to easily find the required version of data and select it for recovery. On a managed machine (p. 237), users can view and search data in any vault (p. 240) accessible from this machine. The centralized catalog available on the management server (p. 237) contains all data stored on its storage nodes (p. 239).

Physically, data catalog is stored in catalog files. Every vault uses its own set of catalog files which normally are located directly in the vault. If this is not possible, such as for tape storages, the catalog files are stored in the managed machine's or storage node's local folder. Also, a storage node locally stores catalog files of its remote vaults, for the purpose of fast access.

## Deduplicating vault

A managed vault (p. 237) in which deduplication (p. 232) is enabled.

## Deduplication

A method of storing different duplicates of the same information only once.

Acronis Backup & Recovery 11 can apply the deduplication technology to backup archives (p. 228) stored on storage nodes (p. 239). This minimizes storage space taken by the archives, backup traffic and network usage during backup.

## Differential backup

A differential backup stores changes to the data against the latest full backup (p. 235). You need access to the corresponding full backup to recover the data from a differential backup.

## Direct management

An operation that is performed on a managed machine (p. 237) using the direct console (p. 231)-agent (p. 228) connection (as opposed to centralized management (p. 230) when the operations are configured on the management server (p. 237) and propagated by the server to the managed machines).

The direct management operations include:

- creating and managing local backup plans (p. 236)
- creating and managing local tasks (p. 236) such as recovery tasks
- creating and managing personal vaults (p. 237) and archives stored there

- viewing the state, progress and properties of the centralized tasks (p. 231) existing on the machine
- viewing and managing the log of the agent's operations
- disk management operations, such as clone a disk, create volume, convert volume.

A kind of direct management is performed when using bootable media (p. 229).

## Disaster recovery plan (DRP)

An e-mail message that contains a list of backed up data items and detailed instructions on how to recover these items from a backup.

If the corresponding backup option (p. 228) is enabled, a DRP is sent to the specified e-mail addresses after the first successful backup performed by the backup plan, and also after any change to the list of data items or the DRP parameters.

## Disk backup (Image)

A backup (p. 228) that contains a sector-based copy of a disk or a volume in a packaged form. Normally, only sectors that contain data are copied. Acronis Backup & Recovery 11 provides an option to take a raw image, that is, copy all the disk sectors, which enables imaging of unsupported file systems.

## Disk group

A number of dynamic disks (p. 233) that store the common configuration data in their LDM databases and therefore can be managed as a whole. Normally, all dynamic disks created within the same machine (p. 236) are members of the same disk group.

As soon as the first dynamic disk is created by the LDM or another disk management tool, the disk group name can be found in the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\dmio\Boot Info\Primary Disk Group\Name.

The next created or imported disks are added to the same disk group. The group exists until at least one of its members exists. Once the last dynamic disk is disconnected or converted to basic, the group is discontinued, though its name is kept in the above registry key. In case a dynamic disk is created or connected again, a disk group with an incremental name is created.

When moved to another machine, a disk group is considered as 'foreign' and cannot be used until imported into the existing disk group. The import updates the configuration data on both the local and the foreign disks so that they form a single entity. A foreign group is imported as is (will have the original name) if no disk group exists on the machine.
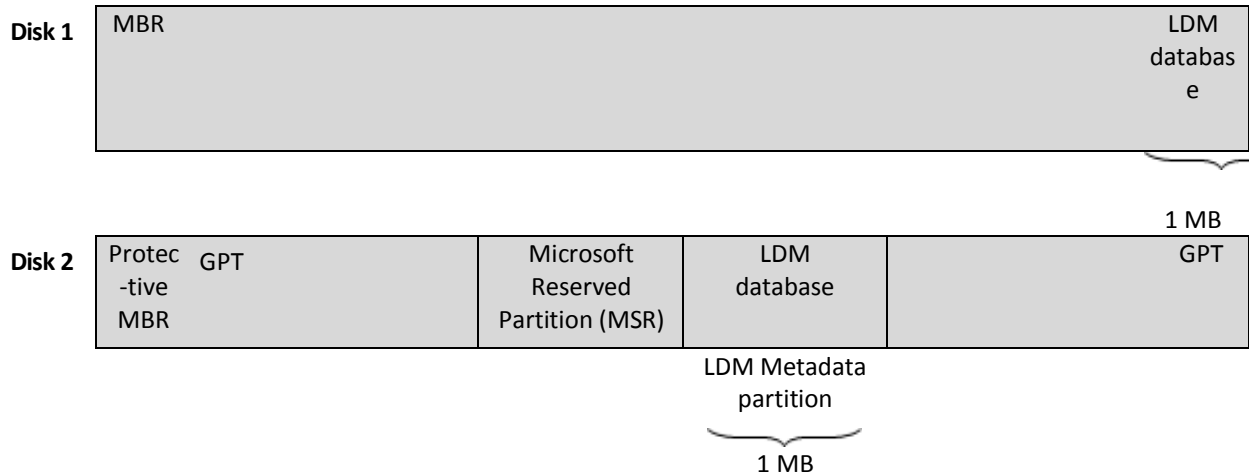
For more information about disk groups please refer to the following Microsoft knowledge base article:

222189 Description of Disk Groups in Windows Disk Management
http://support.microsoft.com/kb/222189/EN-US/

# Dynamic disk

A hard disk managed by Logical Disk Manager (LDM) that is available in Windows starting with Windows 2000. LDM helps flexibly allocate volumes on a storage device for better fault tolerance, better performance or larger volume size.

A dynamic disk can use either the master boot record (MBR) or GUID partition table (GPT) partition style. In addition to MBR or GPT, each dynamic disk has a hidden database where the LDM stores the dynamic volumes' configuration. Each dynamic disk holds the complete information about all dynamic volumes existing in the disk group which makes for better storage reliability. The database occupies the last 1MB of an MBR disk. On a GPT disk, Windows creates the dedicated LDM Metadata partition, taking space from the Microsoft Reserved Partition (MSR.)

| Disk 1 | MBR | | | | LDM database |
|---|---|---|---|---|---|

1 MB

| Disk 2 | Protec-tive MBR | GPT | Microsoft Reserved Partition (MSR) | LDM database | | GPT |
|---|---|---|---|---|---|---|

LDM Metadata partition

1 MB

**Dynamic disks organized on MBR (Disk 1) and GPT (Disk 2) disks.**

For more information about dynamic disks please refer to the following Microsoft knowledge base articles:

Disk Management (Windows XP Professional Resource Kit)
http://technet.microsoft.com/en-us/library/bb457110.aspx

816307 Best practices for using dynamic disks on Windows Server 2003-based computers
http://support.microsoft.com/kb/816307

# Dynamic group

A group of machines (p. 236) which is populated automatically by the management server (p. 237) according to membership criteria specified by the administrator. Acronis Backup & Recovery 11 offers the following membership criteria:

- Operating system
- Active Directory organizational unit
- IP address range
- Listed in txt/csv file.

A machine remains in a dynamic group as long as the machine meets the group's criteria. However, the administrator can specify exclusions and not include certain machines in the dynamic group even if they meet the criteria.

## Dynamic volume

Any volume located on dynamic disks (p. 233), or more precisely, on a disk group (p. 233). Dynamic volumes can span multiple disks. Dynamic volumes are usually configured depending on the desired goal:

- to increase the volume size (a spanned volume)
- to reduce the access time (a striped volume)
- to achieve fault tolerance by introducing redundancy (mirrored and RAID-5 volumes.)

# E

## Encrypted archive

A backup archive (p. 228) encrypted according to the Advanced Encryption Standard (AES). When the encryption option and a password for the archive are set in the backup options (p. 228), each backup belonging to the archive is encrypted by the agent (p. 228) before saving the backup to its destination.

## Encrypted vault

A managed vault (p. 237) to which anything written is encrypted and anything read is decrypted transparently by the storage node (p. 239), using a vault-specific encryption key stored on the node. In case the storage medium is stolen or accessed by an unauthorized person, the malefactor will not be able to decrypt the vault contents without access to the storage node. Encrypted archives (p. 235) will be encrypted over the encryption performed by the agent (p. 228).

## Export

An operation that creates a copy of an archive (p. 228) or a self-sufficient part copy of an archive in the location you specify. The export operation can be applied to a single archive, a single backup (p. 228) or to your choice of backups belonging to the same archive. An entire vault (p. 240) can be exported by using the command line interface.

# F

## Full backup

A self-sufficient backup (p. 228) containing all data chosen for backup. You do not need access to any other backup to recover the data from a full backup.

# G

## GFS   (Grandfather-Father-Son)

A popular backup scheme (p. 229) aimed to maintain the optimal balance between a backup archive (p. 228) size and the number of recovery points (p. 238) available from the archive. GFS enables recovering with daily resolution for the last several days, weekly resolution for the last several weeks and monthly resolution for any time in the past.

For more information please refer to GFS backup scheme.

# I

## Image

The same as Disk backup (p. 233).

## Incremental backup

A backup (p. 228) that stores changes to the data against the latest backup. You need access to other backups from the same archive (p. 228) to restore data from an incremental backup.

## Indexing

An activity (p. 228) performed by a storage node (p. 239) after a backup (p. 228) has been saved to a deduplicating vault (p. 232).

During indexing, the storage node performs the following operations:

- Moves data blocks from the backup to a special file within the vault. This file is called the deduplication data store.
- In the backup, replaces the moved blocks with their fingerprints ("hashes")
- Saves the hashes and the links that are necessary to "assemble" the deduplicated data, to the deduplication database.

Indexing can be thought of as "deduplication at target", as opposed to "deduplication at source" which is performed by the agent (p. 228) during the backup operation (p. 228). A user can suspend and resume indexing.

# L

## Local backup plan

A backup plan (p. 229) created on a managed machine (p. 237) using direct management (p. 232).

## Local task

A task (p. 239) created on a managed machine (p. 237) using direct management (p. 232).

## Logical volume

This term has two meanings, depending on the context.

- A volume, information about which is stored in the extended partition table. (In contrast to a primary volume, information about which is stored in the Master Boot Record.)
- A volume created using Logical Volume Manager (LVM) for Linux kernel. LVM gives an administrator the flexibility to redistribute large storage space on demand, add new and take out old physical disks without interrupting user service. Acronis Backup & Recovery 11 Agent (p. 228) for Linux can access, back up and recover logical volumes when running in Linux with 2.6.x kernel or a Linux-based bootable media (p. 229).

# M

## Machine

A physical or virtual computer uniquely identified by an operating system installation. Machines with multiple operating systems (multi-boot systems) are considered as multiple machines.

## Managed machine

A machine (p. 236), either physical or virtual, where at least one Acronis Backup & Recovery 11 Agent (p. 228) is installed.

## Managed vault

A centralized vault (p. 231) managed by a storage node (p. 239). Archives (p. 228) in a managed vault can be accessed as follows:

        bsp://node_address/vault_name/archive_name/

Physically, managed vaults can reside on a network share, SAN, NAS, on a hard drive local to the storage node or on a tape library locally attached to the storage node. The storage node performs cleanup (p. 231) and validation (p. 239) for each archive stored in the managed vault. An administrator can specify additional operations that the storage node will perform (deduplication (p. 232), encryption).

## Management server (Acronis Backup & Recovery 11 Management Server)

A central server that drives data protection within the enterprise network. Acronis Backup & Recovery 11 Management Server provides the administrator with:

- a single entry point to the Acronis Backup & Recovery 11 infrastructure
- an easy way to protect data on numerous machines (p. 236) using centralized backup plans (p. 230) and grouping
- enterprise-wide monitoring and reporting functionality
- the ability to create centralized vaults (p. 231) for storing enterprise backup archives (p. 228)
- the ability to manage storage nodes (p. 239)
- the centralized catalog (p. 232) of all data stored on the storage nodes.

If there are multiple management servers on the network, they operate independently, manage different machines and use different centralized vaults for storing archives.

## Media builder

A dedicated tool for creating bootable media (p. 229).

# P

## Personal vault

A local or networked vault (p. 240) created using direct management (p. 232). Once a personal vault is created, a shortcut to it appears on the managed machine in the **Vaults** list. Multiple machines can use the same physical location; for example, a network share; as a personal vault.

## Plan

See Backup plan (p. 229).

# R

## Recovery point

Date and time to which the backed up data can be reverted to.

## Registered machine

A machine (p. 236) managed by a management server (p. 237). A machine can be registered on only one management server at a time. A machine becomes registered as a result of the registration (p. 238) procedure.

## Registration

A procedure that adds a managed machine (p. 237) to a management server (p. 237).

Registration sets up a trust relationship between the agent (p. 228) residing on the machine and the server. During registration, the console retrieves the management server's client certificate and passes it to the agent which uses it later to authenticate clients attempting to connect. This helps prevent any attempts by network attackers from establishing a fake connection on behalf of a trusted principal (the management server).

## Replenishable pool

A tape pool that is allowed to take tapes from the **Free tapes** pool when required.

## Replication

Copying a backup (p. 228) to another location. By default, the backup is copied immediately after creation. A user has the option to postpone copying the backup by setting up replication inactivity time.

This feature replaces and enhances the dual destination backup feature, which was available in Acronis Backup & Recovery 10.

## Retention rules

A part of backup plan (p. 229) that specifies when and how to delete or move the backups (p. 228) created by the plan.

# S

## Static group

A group of machines which a management server (p. 237) administrator populates by manually adding machines to the group. A machine remains in a static group until the administrator removes it from the group or from the management server.

## Storage node (Acronis Backup & Recovery 11 Storage Node)

A server aimed to optimize usage of various resources required for protection of enterprise data. This goal is achieved by organizing managed vaults (p. 237). Storage Node enables the administrator to:

- use a single centralized catalog (p. 232) of data stored in the managed vaults
- relieve managed machines (p. 237) of unnecessary CPU load by performing cleanup (p. 231), validation (p. 239) and other operations with backup archives (p. 228) which otherwise would be performed by agents (p. 228)
- drastically reduce backup traffic and storage space taken by the archives (p. 228) by using deduplication (p. 232)
- prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted vaults (p. 235).

# T

## Task

A set of actions to be performed by Acronis Backup & Recovery 11 at a certain time or event. The actions are described in a non human-readable service file. The time or event (schedule) is stored in the protected registry keys (in Windows) or on the file system (in Linux).

## Tower of Hanoi

A popular backup scheme (p. 229) aimed to maintain the optimal balance between a backup archive (p. 228) size and the number of recovery points (p. 238) available from the archive. Unlike the GFS (p. 235) scheme that has only three levels of recovery resolution (daily, weekly, monthly resolution), the Tower of Hanoi scheme continuously reduces the time interval between recovery points as the backup age increases. This allows for very efficient usage of the backup storage.

For more information please refer to "Tower of Hanoi backup scheme (p. 41)".

# U

## Unmanaged vault

Any vault (p. 240) that is not a managed vault (p. 237).

# V

## Validation

An operation that checks the possibility of data recovery from a backup (p. 228).

Validation of a file backup imitates recovery of all files from the backup to a dummy destination. Validation of a disk backup calculates a checksum for every data block saved in the backup. Both procedures are resource-intensive.

While the successful validation means a high probability of successful recovery, it does not check all factors that influence the recovery process. If you back up the operating system, only a test recovery under the bootable media to a spare hard drive can guarantee successful recovery in the future.

## Vault

A place for storing backup archives (p. 228). A vault can be organized on a local or networked drive or detachable media, such as an external USB drive. There are no settings for limiting a vault size or the number of backups in a vault. You can limit the size of each archive using cleanup (p. 231), but the total size of archives stored in the vault is limited by the storage size only.

## Virtual machine

On Acronis Backup & Recovery 11 Management Server (p. 237), a machine (p. 236) is considered virtual if it can be backed up from the virtualization host without installing an agent (p. 228) on the machine. Such machine appears in the **Virtual machines** section. If an agent is installed into the guest system, the machine appears in the **Machines with agents** section.

# W

## WinPE (Windows Preinstallation Environment)

A minimal Windows system based on any of the following kernels:

- Windows XP Professional with Service Pack 2 (PE 1.5)
- Windows Server 2003 with Service Pack 1 (PE 1.6)
- Windows Vista (PE 2.0)
- Windows Vista SP1 and Windows Server 2008 (PE 2.1)
- Windows 7 (PE 3.0).

WinPE is commonly used by OEMs and corporations for deployment, test, diagnostic and system repair purposes. A machine can be booted into WinPE via PXE, CD-ROM, USB flash drive or hard disk. Acronis Plug-in for WinPE (p. 227) enables running the Acronis Backup & Recovery 11 Agent (p. 228) in the preinstallation environment.