

Acronis



Acronis Backup

Version 11.5 Update 6

APPLIES TO THE FOLLOWING PRODUCTS

For Linux Server

USER GUIDE

Copyright Statement

Copyright © Acronis International GmbH, 2002-2015. All rights reserved.

“Acronis” and “Acronis Secure Zone” are registered trademarks of Acronis International GmbH.

"Acronis Compute with Confidence", "Acronis Startup Recovery Manager", "Acronis Active Restore", "Acronis Instant Restore" and the Acronis logo are trademarks of Acronis International GmbH.

Linux is a registered trademark of Linus Torvalds.

VMware and VMware Ready are trademarks and/or registered trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Windows and MS-DOS are registered trademarks of Microsoft Corporation.

All other trademarks and copyrights referred to are the property of their respective owners.

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of this work or derivative work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Third party code may be provided with the Software and/or Service. The license terms for such third-parties are detailed in the license.txt file located in the root installation directory. You can always find the latest up-to-date list of the third party code and the associated license terms used with the Software and/or Service at <http://kb.acronis.com/content/7696>

Acronis patented technologies

Technologies, used in this product, are covered and protected by one or more U.S. Patent Numbers: 7,047,380; 7,275,139; 7,281,104; 7,318,135; 7,353,355; 7,366,859; 7,475,282; 7,603,533; 7,636,824; 7,650,473; 7,721,138; 7,779,221; 7,831,789; 7,886,120; 7,895,403; 7,934,064; 7,937,612; 7,949,635; 7,953,948; 7,979,690; 8,005,797; 8,051,044; 8,069,320; 8,073,815; 8,074,035; 8,145,607; 8,180,984; 8,225,133; 8,261,035; 8,296,264; 8,312,259; 8,347,137; 8,484,427; 8,645,748; 8,732,121 and patent pending applications.

Table of contents

1	Introducing Acronis Backup	8
1.1	What's new in Update 6	8
1.2	What's new in Update 5	8
1.3	What's new in Update 4	8
1.4	What's new in Update 3	8
1.5	What's new in Update 2	9
1.6	What's new in Update 1	9
1.7	What's new in Acronis Backup & Recovery 11.5	9
1.8	Acronis Backup components	9
1.8.1	Agent for Linux	10
1.8.2	Management Console	10
1.8.3	Bootable Media Builder	10
1.9	About using the product in the trial mode	11
1.10	Supported file systems	11
1.11	Technical Support	12
2	Getting started	13
2.1	Using the management console	14
2.1.1	"Navigation" pane	15
2.1.2	Main area, views and action pages	16
2.1.3	Console options	19
3	Understanding Acronis Backup	22
3.1	Owners	22
3.2	Credentials used in backup plans and tasks	22
3.3	User privileges on a managed machine	24
3.4	Full, incremental and differential backups	24
3.5	What does a disk or volume backup store?	25
3.6	Backup and recovery of logical volumes and MD devices (Linux)	26
3.6.1	Backing up logical volumes	26
3.6.2	Backing up MD devices	27
3.6.3	Backing up hardware RAID arrays (Linux)	27
3.6.4	Assembling MD devices for recovery (Linux)	28
3.6.5	Recovering MD devices and logical volumes	28
3.7	Support for Advanced Format (4K-sector) hard disks	32
3.8	Support for UEFI-based machines	33
3.9	Support for SNMP	33
4	Backup	35
4.1	Back up now	35
4.2	Creating a backup plan	35
4.2.1	Selecting data to back up	37
4.2.2	Access credentials for source	38
4.2.3	Source files exclusion	38

4.2.4	Backup location selection	40
4.2.5	Access credentials for archive location.....	42
4.2.6	Backup schemes	43
4.2.7	Archive validation.....	52
4.2.8	Backup plan's credentials	53
4.2.9	Label (Preserving machine properties in a backup).....	53
4.2.10	Sequence of operations in a backup plan	55
4.2.11	Why is the program asking for the password?	55
4.3	Simplified naming of backup files	55
4.3.1	The [DATE] variable.....	56
4.3.2	Backup splitting and simplified file naming.....	57
4.3.3	Usage examples.....	57
4.4	Scheduling.....	60
4.4.1	Daily schedule.....	61
4.4.2	Weekly schedule	63
4.4.3	Monthly schedule.....	65
4.4.4	Conditions.....	67
4.5	Replication and retention of backups.....	69
4.5.1	Supported locations	71
4.5.2	Setting up replication of backups.....	71
4.5.3	Setting up retention of backups.....	72
4.5.4	Retention rules for the Custom scheme.....	73
4.5.5	Usage examples.....	74
4.6	How to disable backup cataloging	76
4.7	Default backup options.....	76
4.7.1	Additional settings.....	78
4.7.2	Archive protection.....	79
4.7.3	Backup cataloging	80
4.7.4	Backup performance.....	80
4.7.5	Backup splitting	81
4.7.6	Compression level	82
4.7.7	Disaster recovery plan (DRP).....	83
4.7.8	E-mail notifications.....	84
4.7.9	Error handling	85
4.7.10	Event tracing.....	86
4.7.11	Fast incremental/differential backup	86
4.7.12	File-level backup snapshot	87
4.7.13	LVM snapshotting	87
4.7.14	Media components.....	89
4.7.15	Pre/Post commands.....	89
4.7.16	Pre/Post data capture commands.....	91
4.7.17	Replication/cleanup inactivity time	93
4.7.18	Sector-by-sector backup.....	93
4.7.19	Task failure handling	93
4.7.20	Task start conditions	94
5	Recovery	96
5.1	Creating a recovery task	96
5.1.1	What to recover	97
5.1.2	Access credentials for location	100
5.1.3	Access credentials for destination	101
5.1.4	Where to recover	101
5.1.5	When to recover	107
5.1.6	Task credentials.....	107

5.2	Bootability troubleshooting.....	108
5.2.1	How to reactivate GRUB and change its configuration	109
5.3	Default recovery options	110
5.3.1	Additional settings.....	112
5.3.2	E-mail notifications.....	113
5.3.3	Error handling.....	114
5.3.4	Event tracing.....	114
5.3.5	File-level security.....	115
5.3.6	Pre/Post commands.....	115
5.3.7	Recovery priority	117
6	Conversion to a virtual machine	118
6.1	Conversion methods.....	118
6.2	Conversion to an automatically created virtual machine.....	118
6.2.1	Considerations before conversion	119
6.2.2	Setting up regular conversion to a virtual machine.....	120
6.2.3	Recovery to the "New virtual machine" destination.....	123
6.3	Recovery to a manually created virtual machine	126
6.3.1	Considerations before conversion	126
6.3.2	Steps to perform	127
7	Storing the backed up data.....	128
7.1	Vaults	128
7.1.1	Working with vaults	129
7.1.2	Personal vaults	129
7.2	Acronis Secure Zone	132
7.2.1	Creating Acronis Secure Zone	132
7.2.2	Managing Acronis Secure Zone.....	134
7.3	Removable devices	135
8	Operations with archives and backups.....	137
8.1	Validating archives and backups.....	137
8.1.1	Archive selection	138
8.1.2	Backup selection.....	138
8.1.3	Vault selection	138
8.1.4	Access credentials for source	139
8.1.5	When to validate.....	139
8.1.6	Task credentials.....	140
8.2	Exporting archives and backups	140
8.2.1	Archive selection	142
8.2.2	Backup selection.....	143
8.2.3	Access credentials for source	143
8.2.4	Destination selection	143
8.2.5	Access credentials for destination	144
8.3	Mounting an image.....	145
8.3.1	Archive selection	146
8.3.2	Backup selection.....	146
8.3.3	Access credentials	146
8.3.4	Volume selection.....	147
8.3.5	Managing mounted images.....	147
8.4	Operations available in vaults.....	148
8.4.1	Operations with archives.....	148

8.4.2	Operations with backups.....	149
8.4.3	Converting a backup to full.....	149
8.4.4	Deleting archives and backups.....	150
9	Bootable media	151
9.1	Linux-based bootable media	151
9.1.1	Kernel parameters.....	152
9.1.2	Network settings	154
9.1.3	Network port	155
9.2	Preparing to work under bootable media	155
9.3	Working under bootable media.....	156
9.3.1	Setting up a display mode	156
9.3.2	Configuring iSCSI and NDAS devices	156
9.4	List of commands and utilities available in Linux-based bootable media	157
9.5	Acronis Startup Recovery Manager	158
10	Administering a managed machine.....	160
10.1	Backup plans and tasks.....	160
10.1.1	Actions on backup plans and tasks	160
10.1.2	States and statuses of backup plans and tasks	162
10.1.3	Export and import of backup plans.....	164
10.1.4	Deploying backup plans as files.....	167
10.1.5	Backup plan details	169
10.1.6	Task/activity details.....	170
10.2	Log.....	170
10.2.1	Actions on log entries.....	170
10.2.2	Log entry details	171
10.3	Alerts.....	172
10.4	Changing a license	173
10.5	Collecting system information.....	173
10.6	Adjusting machine options	174
10.6.1	Alerts	174
10.6.2	E-mail settings	174
10.6.3	Event tracing.....	176
10.6.4	Log cleanup rules.....	177
10.6.5	Cloud backup proxy.....	177
11	Cloud backup	179
11.1	Introduction to Acronis Cloud Backup.....	179
11.1.1	What is Acronis Cloud Backup?.....	179
11.1.2	What data can I back up and recover?	179
11.1.3	How long will my backups be kept in the cloud storage?	179
11.1.4	How do I secure my data?	180
11.1.5	Supported operating systems and virtualization products.....	180
11.1.6	Backup and recovery FAQ	181
11.1.7	Initial Seeding FAQ	182
11.1.8	Large Scale Recovery FAQ	188
11.1.9	Subscription lifecycle FAQ	189
11.2	Where do I start?	192
11.3	Choosing a subscription.....	192
11.4	Activating cloud backup subscriptions	193

11.4.1	Activating subscriptions in Acronis Backup	193
11.4.2	Reassigning an activated subscription	193
11.5	Configuring proxy settings	194
11.6	Retrieving files from the cloud storage by using a web browser	195
11.7	Limitations of the cloud storage	196
11.8	Terminology reference	196
12	Glossary.....	199

1 Introducing Acronis Backup

1.1 What's new in Update 6

Improvements added in build 43798

- Support for Linux kernel version 3.17, 3.18, 4.0, and 4.1.
- Support for Ubuntu 15.04, Fedora 22, Oracle Linux 7.1, Debian 8.0 and 8.1.

Cloud backup

- A registered subscription for Server is activated automatically when the first backup to cloud starts, if all subscriptions available in your account are the same type and have the same storage quota.

Linux

- Support for Linux kernel version 3.14, 3.15, and 3.16.
- Support for Red Hat Enterprise Linux 7.x, Ubuntu 14.04 and 14.10, Fedora 21, SUSE Linux Enterprise Server 12, Debian 7.6 and 7.7, CentOS 7.0, and Oracle Linux 7.0.

1.2 What's new in Update 5

- The most recent operating systems remain bootable when recovered to dissimilar hardware. To simplify the recovery procedure, the Acronis Universal Restore setting is removed from the recovery wizard. A separate Acronis Universal Restore tool is introduced for the cases when installing a custom driver is required. The tool can be downloaded from your account on the Acronis website.
- The cataloging performance is improved.
- It is now possible to contact Acronis technical support directly from the error dialog box.

1.3 What's new in Update 4

- Improved backup to Acronis Cloud Storage.
- Support for Linux kernel version 3.13.
- Support for Debian 7.2, 7.4, and 7.5.

1.4 What's new in Update 3

Rebranding

Acronis Backup & Recovery 11.5 is renamed to Acronis Backup.

Licensing

The Universal Restore feature is included in all Acronis Backup licenses. The Universal Restore add-on license is deprecated.

Cloud backup

Cloud backup (p. 179) is now available for machines running Linux.

OS support

- Support for Linux kernel up to version 3.12.
- Support for Fedora 19, Fedora 20, and Debian 7.

1.5 What's new in Update 2

- Support for Ubuntu 13.10.

1.6 What's new in Update 1

Improvements added in build 37975

- Installation of Acronis Backup & Recovery 11.5 in the trial mode without a license key.
- Upgrade from a stand-alone product to the advanced platform without reinstalling the software.
- Support for Linux kernel up to version 3.9.
- Support for Ubuntu 12.10, 13.04, and Fedora 18.

Other

- Support for Oracle Linux 5.x, 6.x – both Unbreakable Enterprise Kernel and Red Hat Compatible Kernel.
- Completely disable backup cataloging (p. 76).
- Save a Disaster Recovery Plan (p. 83) to a local or network folder, in addition to sending it via e-mail.

1.7 What's new in Acronis Backup & Recovery 11.5

The following is a summary of the product's new features and enhancements.

Operating systems and platforms

- **Support for Linux kernel up to 3.6**
- **Support for the following Linux distributions:**
 - Ubuntu 11.04, 11.10, 12.04
 - Fedora 15, 16, 17
 - Debian 6
 - CentOS 6
- **Support for Unified Extensible Firmware Interface (UEFI)** (p. 33)
Back up a UEFI-based machine and recover it to the same or a different UEFI-based machine.

Bootable media

- New Linux kernel version (3.4.5) in Linux-based bootable media. The new kernel makes for better hardware support.

Usability

- Support for 800x600 screen resolution.

1.8 Acronis Backup components

This section contains a list of Acronis Backup 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. Agents require a license to perform operations on each managed machine.

Console

The console provides Graphical User Interface to the agents. Usage of the console is not licensed. 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. Bootable Media Builder is installed together with the agent.

1.8.1 Agent for Linux

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

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 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 directories residing on the machine where the agent is installed or on a network share accessed using the smb or nfs protocol. Files can be recovered to their original location or to another place. It is possible to recover all files and directories that were backed up or select which of them to recover.

Conversion to a virtual machine

Agent for Linux performs the conversion by recovering a disk backup to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, 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 directory you select. You can start the machine using the respective virtualization software or prepare the machine files for further usage.

1.8.2 Management Console

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

1.8.3 Bootable Media Builder

Acronis Bootable Media Builder is a dedicated tool for creating bootable media (p. 201). The media builder that installs on Linux creates bootable media based on Linux kernel.

1.9 About using the product in the trial mode

Before buying an Acronis Backup license, you may want to try the software. This can be done without a license key.

To install the product in the trial mode, run the setup program locally or use the remote installation functionality. Unattended installation and other ways of installation are not supported.

Limitations of the trial mode

When working under bootable media:

- The disk management functionality is not available. You can try the user interface, but there is no option to commit the changes.
- The recovery functionality is available, but the backup functionality is not. To try the backup functionality, install the software in the operating system.

Upgrading to the full mode

After the trial period expires, the product GUI displays a notification requesting you to specify or obtain a license key.

To specify a license key, click **Help > Change License** (p. 173). Specifying the key by running the setup program is not possible.

If you have activated a trial or purchased a subscription for the cloud backup service (p. 179), cloud backup will be available until the subscription period expires, regardless of whether you specify a license key.

1.10 Supported file systems

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

- FAT16/32
- NTFS
- ReFS - volume recovery without the volume resize capability. Supported in Windows Server 2012/2012 R2 only.
- Ext2/Ext3/Ext4
- ReiserFS3 - particular files cannot be recovered from disk backups located on Acronis Backup Storage Node
- ReiserFS4 - volume recovery without the volume resize capability; particular files cannot be recovered from disk backups located on Acronis Backup Storage Node
- XFS - volume recovery without the volume resize capability; particular files cannot be recovered from disk backups located on Acronis Backup Storage Node
- JFS - particular files cannot be recovered from disk backups located on Acronis Backup Storage Node
- Linux SWAP

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

1.11 Technical Support

Maintenance and Support Program

If you need assistance with your Acronis product, please go to <http://www.acronis.com/en-us/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** (<http://www.acronis.com/en-us/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 a license key for Acronis Backup for Linux Server.
- You have the setup program. You can download it from the Acronis website.
- Make sure that the RPM Package Manager (RPM) and the following Linux packages are installed: **gcc**, **make**, and **kernel-devel**. The names of these packages may vary depending on the Linux distribution.

To install Acronis Backup

Run the **AcronisBackupL.i686** or the **AcronisBackupL.x86_64** installation file and follow the on-screen instructions.



Step 2. Running

Log in as root or log in as an ordinary user and then switch user as required. Start the console with the command

```
/usr/sbin/acronis_console
```

 For understanding of the GUI elements see "Using the management console" (p. 14).



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. 151).



Step 4. Backup



Back up now (p. 35)

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. 35)

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. 96)

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. 211). Navigate to the  **Vaults** (p. 128) 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 starts, 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 machine-specific operations.



Acronis Backup Management Console - Welcome screen

Key elements of the console workspace

	Name	Description
1	Navigation pane	Contains the Navigation tree. Lets you navigate to the different views. For details, see Navigation pane (p. 15).
2	Main area	Here you configure and monitor backup, recovery and other operations. The main area displays views and action pages (p. 16) 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. 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.

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.

-  **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.

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. 15).



"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. 17) 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. 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 predefined parameters	Click the appropriate buttons above the table. For example, in the Log view, you can filter the log entries by event type (Error, Warning, Information) or by the period when the event occurred (For last 24 hours, For last week, For last three months, or For custom period).
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 <ol style="list-style-type: none">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 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.

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: **Enabled**.

If the option is enabled, the credentials for various locations that you enter during a console session 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.

If the option is disabled, the credentials are stored only until the console is closed.

To clear the credentials cache for the current user account, click the **Clear credentials cache** button.

2.1.3.3 Fonts

The option defines the fonts to be used in the Graphical User Interface of Acronis Backup. 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

The “Interaction Required” 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 “Interaction Required” 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 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.

Request description when ejecting a tape

This option defines whether to display a prompt for you to describe a tape when you eject it from a tape device by using Acronis Backup. For example, you may describe the physical location where the tape will be kept (recommended). If a tape is ejected automatically according to the **Eject tapes after successful backups** option, no such prompt is displayed.

The preset is: **Enabled**.

To make a selection, select or clear the **Request description when ejecting a tape** check box.

Note *Tape devices are available only if you have upgraded from Acronis Backup & Recovery 10.*

About the task execution results

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.

3 Understanding Acronis Backup

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

This section explains the concept of a backup plan (task) owner and an archive owner.

Plan (task) owner

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

Tasks, belonging to a backup plan, 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.

3.2 Credentials used in backup plans and tasks

This section explains the concept of access credentials, backup plan's credentials and task credentials.

Access credentials

When browsing backup locations, setting up backups, or creating recovery tasks, you may need to provide credentials for accessing various resources, such as the data you are going to back up or the location where the backups are (or will be) stored.

If the **Credentials cache** (p. 19) option is enabled (it is enabled by default), the credentials which you provide during a console session are saved for use during the later sessions. Thus, there is no need to enter the credentials next time. The credentials are cached independently for each user who uses the console on the machine.

Backup plan's credentials

Any backup plan running on a machine runs on behalf of a user.

In Windows

By default, the plan runs under the agent service account, if created by a user having administrative privileges on the machine. If created by a regular user, such as a member of the **Users** group, the plan runs under this user's account.

When creating a backup plan, you are only asked for credentials in specific cases. For example:

- You are scheduling backups as a regular user and did not enter credentials when connecting the console to the machine. This may be the case when the console is installed on the same machine that you are backing up.
- You are backing up a Microsoft Exchange cluster to a storage node.

Specifying the credentials explicitly

You have the option to explicitly specify a user account under which the backup plan will run. To do this, on the backup plan creation page:

1. In the **Plan parameters** section, click **Show plan's credentials, comments, label**.
2. Click **Plan's credentials**.
3. Enter the credentials under which the plan will run. When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).

In Linux

You do not need to specify backup plan's credentials. In Linux, backup plans always run under the root user account.

Task credentials

Like a backup plan, any task runs on behalf of a user.

In Windows

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

- **Manual start**
Every time you manually 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.
- **Scheduled or postponed start**
The task credentials are mandatory. You cannot complete the task creation until you specify the task credentials. Task credentials are specified on the task creation page in a similar manner as the plan's credentials are specified.

In Linux

You do not need to specify task credentials. In Linux, tasks always run under the root user account.

3.3 User privileges on a managed machine

When managing a machine running Linux, the user has or obtains the root privileges, and so can:

- Back up and recover any data or the entire machine, having full control over all Acronis Backup agent operations and log files on the machine.
- Manage local backup plans and tasks owned by any user registered in the operating system.

To avoid routine logging on to the system as root, the root user can log on with the ordinary user credentials and then switch user as required.

3.4 Full, incremental and differential backups

Acronis Backup 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. 200) 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

3.5 What does a disk or volume backup store?

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

Windows

A volume backup stores all 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).

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.

The following items are not included in a disk or volume backup (as well as in a file-level backup):

- The swap file (pagefile.sys) and the file that keeps the RAM content when the machine goes into hibernation (hiberfil.sys). After recovery, the files will be re-created in the appropriate place with the zero size.
- Windows shadow storage. The path to it is determined in the registry value **VSS Default Provider** which can be found in the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup**. This means that in operating systems starting with Windows Vista, Windows Restore Points are not backed up.

Linux

A volume backup stores all files and directories 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.

With the **sector-by-sector (raw mode)** option enabled, a disk backup stores all the disk sectors. The sector-by-sector backup can be used for backing up disks with unrecognized or unsupported file systems and other proprietary data formats.

3.6 Backup and recovery of logical volumes and MD devices (Linux)

This section explains how you would back up and recover volumes managed by Linux Logical Volume Manager (LVM), called logical volumes; and multiple-disk (MD) devices, called Linux Software RAID.

To learn more about LVM please visit <http://tldp.org/HOWTO/LVM-HOWTO/> or http://www.centos.org/docs/5/html/5.1/Deployment_Guide/ch-lvm.html.

3.6.1 Backing up logical volumes

Acronis Backup Agent for Linux can access, back up, and recover logical volumes when running in Linux starting with 2.6.x kernel or under Linux-based bootable media.

Backup

In Acronis Backup GUI, logical volumes appear under **Dynamic volumes** at the end of the list of volumes available for backup. If you select logical volumes for backup, the logical volume structure will be saved to the backup along with the volume contents. This structure can be automatically recreated when you recover these volumes under a Linux-based bootable media.

To back up all available disks, specify all logical volumes plus basic volumes not belonging to them. This is the default choice when you open the **Create backup plan** page.

Basic volumes included in logical volumes are shown in the list with **None** in the **File system** column. If you select such volumes, the program will back them up sector-by-sector. Normally this is not required.

Recovery

When recovering logical volumes, you have two options:

- **Recovering volume contents only.** The type or other properties of the target volume will not change.
This option is available both in the operating system and under bootable media.

This option is useful in the following cases:

- When some data on the volume was lost, but no hard disks were replaced.
- When recovering a logical volume over a basic disk or volume. You can resize the resulting volume in this case.

A system, recovered from a logical volume backup to a basic disk, cannot boot because its kernel tries to mount the root file system at the logical volume. To boot the system, change the loader configuration and `/etc/fstab` so that LVM is not used and reactivate your boot loader (p. 109).

- When recovering a basic or logical volume to an existing logical volume.

If the boot partition (`/boot`) was located on a basic volume, we recommend recovering it to a basic volume, even if your boot loader supports booting from logical volumes.

- **Recovering both the structure of logical volumes and their contents.**

Such is the case when recovering on bare metal or on a machine with different volume structure. The structure of logical volumes can be automatically created at the time of recovery (p. 29).

This option is available only under bootable media.

For detailed instructions on how to recover logical volumes, see [Recovering MD devices and logical volumes](#) (p. 28).

3.6.2 Backing up MD devices

MD devices, known as Linux Software RAID, combine several volumes and make solid block devices (`/dev/md0`, `/dev/md1`, ..., `/dev/md31`). The information about MD devices is stored in `/etc/raidtab` or in dedicated areas of those volumes.

You can back up active (mounted) MD devices in the same way as logical volumes. The MD devices appear at the end of the list of volumes available for backup. If you select MD devices for backup, the structure of the MD devices will be backed up along with their contents.

Backing up volumes included in MD devices does not make sense when an MD device is mounted, as it won't be possible to recover them.

When recovering MD devices under bootable media, the structure of MD devices can be recreated automatically. For detailed information about recovering MD devices under bootable media, see [Recovering MD devices and logical volumes](#) (p. 28).

For information about assembling MD devices when performing recovery in Linux, see [Assembling MD devices for recovery \(Linux\)](#) (p. 28).

3.6.3 Backing up hardware RAID arrays (Linux)

Hardware RAID arrays under Linux combine several physical drives to create a single partitionable disk. The special file related to a hardware RAID array is usually located in `/dev/ataraid`. You can back up hardware RAID arrays in the same way as ordinary hard disks.

Physical drives that are part of hardware RAID arrays may be listed alongside other disks as if they had a bad partition table or no partition table at all. Backing up such disks does not make sense as it won't be possible to recover them.

3.6.4 Assembling MD devices for recovery (Linux)

In Linux, when performing recovery from a disk backup to an existing MD device (also called Linux Software RAID), make sure that this **device is assembled** at the time of recovery.

If the device is not assembled, assemble it by using the **mdadm** utility. Here are two examples:

Example 1. The following command assembles the device `/dev/md0` combined from the volumes `/dev/sdb1` and `/dev/sdc1`:

```
mdadm --assemble /dev/md0 -ayes /dev/sdb1 /sdc1
```

Example 2. The following command assembles the device `/dev/md0` combined from the disks `/dev/sdb` and `/dev/sdc`:

```
mdadm --assemble /dev/md0 -ayes /dev/sdb /dev/sdc
```

If the recovery requires the machine to be rebooted (usually, when the volumes to recover include the boot partition), follow these guidelines:

- If all parts of the MD device are volumes (a typical case, such as in the first example), make sure that each volume type—called partition type or system ID—is **Linux raid automount**; the hexadecimal code of this partition type is `0xFD`. This will guarantee that the device will be automatically assembled following the reboot. To view or change the partition type, use a disk partitioning utility such as **fdisk**.
- Otherwise (such as in the second example), perform the recovery from bootable media. No reboot will be required in that case. In bootable media, you may need to create the MD device manually or automatically, as described in Recovering MD devices and logical volumes (p. 28).

3.6.5 Recovering MD devices and logical volumes

Recovering MD devices and/or volumes created by Logical Volume Manager (logical volumes) assumes that the corresponding volume structure will be recreated.

In Linux-based bootable media, you can choose to recreate the volume structure automatically (p. 29).

This functionality is intended primarily for bare-metal recovery of an entire machine. The software backs up and recreates the entire logical volume structure, even if not all MD devices or logical volumes are being backed up or recovered. Therefore, you need at least as many disks as the original volume structure used.

Do not try to recreate the volume structure automatically in any of the following cases:

- The machine has **data that must be preserved**. The software will destroy all data on the disks that it chooses to recreate the volume structure on.
- The machine has **fewer physical disks** than the original volume structure used. The software will fail to recreate the volume structure even if the capacity of the physical disks is enough to fit all the data being recovered.
- The backup **does not contain the volume structure information**. This information might be absent in backups created by Acronis Backup & Recovery 10, because saving it was optional.

In these cases, create the volume structure manually (p. 29) prior to recovery. You can do this by using the **mdadm** and **lvm** utilities, either in Linux-based bootable media or in Linux.

3.6.5.1 Creating the volume structure automatically

Use the following procedure to automatically recreate the logical volume structure on a machine.

Caution As a result of the following procedure, the current volume structure on the machine will be replaced with the one stored in the backup. This will destroy the data that is currently stored on some or all of the machine's hard disks.

If disk configuration has changed. An MD device or a logical volume resides on one or more disks. If you replaced any of these disks between backup and recovery (or if you are recovering the volumes to a different machine), ensure that the new disk configuration includes at least the same number of disks as the original volume structure did. The capacity of the disks must be enough to fit all the data being recovered.

To create the volume structure automatically

1. Boot the machine from a Linux-based bootable media.
2. Click **Acronis Bootable Agent**. Then, click **Run management console**.
3. In the management console, click **Recover**.
Under the archive contents, Acronis Backup will display a message saying that it detected information about the volume structure.
4. Click **Details** in the area with that message.
5. Review the volume structure, and then click **Apply RAID/LVM** to create it.

3.6.5.2 Creating the volume structure manually

The following is a general procedure for recovering MD devices and logical volumes by using a Linux-based bootable media, and an example of such recovery. You can use a similar procedure in Linux.

To create the volume structure manually

1. Boot the machine from a Linux-based bootable media.
2. Click **Acronis Backup**. Then, click **Run management console**.
3. On the toolbar, click **Actions**, and then click **Start shell**. Alternatively, you can press CTRL+ALT+F2.
4. If necessary, examine the structure of volumes which are stored in the archive, by using the **acrocnd** utility. Also, you can use this utility to mount one or more of these volumes as if they were regular volumes (see "Mounting backup volumes" later in this topic).
5. Create the volume structure according to that in the archive, by using the **mdadm** utility (for MD devices), the **lvm** utility (for logical volumes), or both.

Note: Logical Volume Manager utilities such as **pvcreate** and **vgcreate**, which are normally available in Linux, are not included in the bootable media environment, so you need to use the **Lvm** utility with a corresponding command. For example: **Lvm pvcreate**, **Lvm vgcreate**, and **Lvm lvcreate**.

6. If you previously mounted the backup by using the **acrocnd** utility, use this utility again to unmount the backup (see "Mounting backup volumes" later in this topic).
7. Return to the management console by pressing ALT+F1.
(Do not reboot the machine at this point. Otherwise, you will have to create the volume structure again.)
8. Click **Recover**, then specify the path to the archive and any other required parameters, and then click **OK**.


Note: This procedure will not work if you connect to Acronis Backup Bootable Agent remotely, because the command shell is not available in this case.

Example

Suppose that you previously performed a disk-level backup of a machine with the following disk configuration:

- The machine has two 1-gigabyte and two 2-gigabyte SCSI hard disks, mounted on **/dev/sda**, **/dev/sdb**, **/dev/sdc**, and **/dev/sdd**, respectively.
- The first and second pairs of hard disks are configured as two MD devices; both are in the RAID-1 configuration, and are mounted on **/dev/md0** and **/dev/md1**, respectively.
- A logical volume is based on the two MD devices and is mounted on **/dev/my_volgroup/my_logvol**.

The following picture illustrates this configuration.



Do the following to recover data from this archive.

Step 1: Creating the volume structure

1. Boot the machine from a Linux-based bootable media.
2. In the management console, press CTRL+ALT+F2.
3. Run the following commands to create the MD devices:

```
mdadm --create /dev/md0 --level=1 --raid-devices=2 /dev/sd[ab]
mdadm --create /dev/md1 --level=1 --raid-devices=2 /dev/sd[cd]
```

4. Run the following commands to create the logical volume group:

Caution: The **pvcreate** command destroys all data on the **/dev/md0** and **/dev/md1** devices.

```
lvm pvcreate /dev/md0 /dev/md1
lvm vgcreate my_volgroup /dev/md0 /dev/md1
lvm vgdisplay
```

The output of the **lvm vgdisplay** command will contain lines similar to the following:

```
--- Volume group ---
VG Name      my_volgroup
...
VG Access    read/write
VG Status    resizable
...
VG Size      1.99 GB
...
VG UUID      0qoQ41-Vk7W-yDG3-uF11-Q2AL-C0z0-vMeACu
```

5. Run the following command to create the logical volume; in the **-L** parameter, specify the size given by **VG Size**:

```
lvm lvcreate -L1.99G --name my_logvol my_volgroup
```

6. Activate the volume group by running the following command:

```
lvm vgchange -a y my_volgroup
```

7. Press ALT+F1 to return to the management console.

Step 2: Starting the recovery

1. In the management console, click **Recover**.
2. In **Archive**, click **Change** and then specify the name of the archive.
3. In **Backup**, click **Change** and then select the backup from which you want to recover data.
4. In **Data type**, select **Volumes**.
5. In **Items to recover**, select the check box next to **my_volgroup-my_logvol**.
6. Under **Where to recover**, click **Change**, and then select the logical volume that you created in Step 1. Click the chevron buttons to expand the list of disks.
7. Click **OK** to start the recovery.

For a complete list of commands and utilities that you can use in the bootable media environment, see List of commands and utilities available in Linux-based bootable media (p. 157). For detailed descriptions of the **acrocmd** utility, see the Acronis Backup command-line reference.

Mounting backup volumes

You may want to mount a volume stored in a disk backup, for example, to view some files in it before starting the recovery.

To mount a backup volume

1. Use the **acrocmd list content** command to list the disks and volumes that are stored in the backup. For example, the following command lists the content of the latest backup of the **linux_machine** archive:

```
acrocmd list content --loc=\\server\backups --credentials=user,MyPassWd
--arc=linux_machine
```

The output will contain lines similar to the following:

type: disk	Num	Partition	Flags	Size	Type	GUID
--	--					
	Dyn1	my_volgroup-my_lo...		4 GB	Ext 3	
	Dyn2	md0		2.007 GB	Ext 2	
	Disk 1	sda		16 GB	DT_FIXED	
	1-1	sda1	Act,Pri	203.9 MB	Ext 2	
	1-2	sda2	Pri	11.72 GB	Reiser	
	1-3	sda3	Pri	1.004 GB	Linux swap	
	Disk 2	sdb		8 GB	DT_FIXED	
	2-1	sdb1	Pri	2.007 GB	Ext 2	
	2-2	sdb2	Pri	2.007 GB	None	
	Disk 3	sdc		1 GB	DT_FIXED	
	Disk 4	sdd		8 GB	DT_FIXED	
	4-1	sdd1	Pri	2.007 GB	Ext 2	
	4-2	sdd2	Pri	2.007 GB	None	

- Use the **acrocnd mount** command, specifying the volume's name in the **--volume** parameter. For example:

```
acrocnd mount --loc=\\server\backups --arc=linux_machine --mount_point=/mnt
--volume=DYN1
```

This command mounts the logical volume DYN1 on the mount point /mnt.

To unmount a backup volume

- Use the **acrocnd umount** command, specifying the volume's mount point as a parameter. For example:

```
acrocnd umount --mount_point=/mnt
```

3.7 Support for Advanced Format (4K-sector) hard disks

Acronis Backup can back up hard disks with a sector size of 4 KB (known as Advanced Format disks), as well as traditional hard disks that have 512-byte sectors.

Acronis Backup can recover data from one disk to another as long as *both disks have the same logical sector size*. (This is the sector size presented to the operating system.) Acronis Backup automatically aligns the disk's volumes (p. 105) if necessary. This way, the start of a cluster in the file system always matches the start of a physical sector on the disk.

The disk management functionality of Acronis Backup is not available for disks with a 4-KB logical sector size.

Determining the logical sector size

By disk specification

Development of the Advanced Format technology is coordinated by the International Disk Drive Equipment and Materials Association (IDEMA). For more details, see http://www.idema.org/?page_id=2.

In terms of the logical sector size, IDEMA specifies two types of Advanced Format disks:

- 512 Byte emulation (512e)** disks have 512-byte logical sector size. These disks are supported in Windows starting with Windows Vista, and in modern Linux distributions. Microsoft and Western Digital use the term "Advanced Format" exclusively for this type of disk.

- **4K native (4Kn)** disks have 4-KB logical sector size. Modern operating systems can store data on these disks, but generally cannot boot from these disks. These disks are commonly external drives with USB connection.

By running the appropriate command

To find out the logical sector size of a disk, do the following.

1. Determine the device name of the disk, such as /dev/sdb.
2. Run the following command as the root user, specifying the device name:

```
parted /dev/sdb print
```

3. Examine the first value in the **Sector size (logical/physical)** line. For example, the output may be the following:

```
Sector size (logical/physical): 512B/4096B
```

3.8 Support for UEFI-based machines

Acronis Backup can back up and recover machines that use Unified Extensible Firmware Interface (UEFI) in the same way as it does for machines that use BIOS for booting.

This applies to both physical and virtual machines, no matter if the virtual machines are backed up at a hypervisor level or from inside a guest OS.

Limitations

- Acronis Startup Recovery Manager (ASRM) (p. 199) is not supported on UEFI machines running Linux.
- A machine running Linux cannot be transferred between UEFI and BIOS.

3.9 Support for SNMP

SNMP objects

Acronis Backup 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 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 **supports only TRAP operations**. It is not possible to manage Acronis Backup 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 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. 35) except for the following:

- There are no options to schedule backups and to set up retention rules.
- Simplified naming of backup files (p. 55) is used, if the backup destination supports it. Otherwise, the standard backup naming is used.

The following locations do not support simplified file naming: Acronis Secure Zone and Acronis Cloud Storage.

Due to simplified file naming, an RDX drive or USB flash drive can only be used in the removable media (p. 135) mode.

- 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. 200), please familiarize yourself with the basic concepts used in Acronis Backup.

To create a backup plan, perform the following steps.

What to back up

Items to back up (p. 37)

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. 38)

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

Exclusions (p. 38)

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

Where to back up

Location (p. 40)

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.

Select the mode the removable device will be used in (p. 135)

If the specified location is an RDX drive or USB flash drive, select the device mode:
Removable media or **Fixed drive**.

Backup file naming, access credentials, archive comments

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

File naming (p. 55)

[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 Acronis Secure Zone or Acronis Cloud Storage. When backing up to an RDX drive or USB flash drive, the file naming scheme is determined by the removable device mode (p. 135).

Access credentials (p. 42)

[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. 43)

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).

Replication and retention settings (p. 69)

Not available for removable media or when simplified naming of backup files (p. 55) is chosen.

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

[Optional] To set up replication of backups, select the **Replicate newly created backup to another location** check box. For more information about backup replication, see Setting up replication of backups (p. 71).

Validation, convert to virtual machine

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

When to validate (p. 52)

[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. 120)

[Optional] Applies to: disk or volume backup.

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. 76) 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. 53)

[Optional] Specify the credentials under which the plan will run.

Comments

[Optional] Type a description of the backup plan.

Label (p. 53)

[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. 55).

The plan you have created will be accessible for examination and managing in the **Backup plans and tasks** (p. 160) 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 Agent for Windows or Acronis Backup Agent for Linux is installed.

Select this option to back up entire physical machines or their 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. Also, you can individually recover files and folders. 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 Agent for Windows or Acronis Backup Agent 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.

Notes 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.
- Note for Linux users: Logical volumes and MD devices are shown under **Dynamic volumes**. For more information about backing up such volumes and devices, see "Backup and recovery of logical volumes and MD devices (Linux)" (p. 26).
- Note for Linux users: We recommend that you unmount any volumes that contain non-journaling file systems—such as the ext2 file system—before backing them up. Otherwise, these volumes might contain corrupted files upon recovery; recovery of these volumes with resize might fail.

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 back up.

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.
- **Confirm password.** Re-enter the password.

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, FAT, Ext3, and Ext4 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.

Note: Exclusions override selection of data items to back up. For example, if you select to back up file *MyFile.tmp* and to exclude all *.tmp* files, file *MyFile.tmp* will not be backed up.

To specify which files and folders to exclude, set up any of the following parameters.

Exclude all hidden files and folders

Select this check box to skip files and folders that have the **Hidden** attribute (for file systems that are supported by Windows) or that start with a period (.) (for file systems in Linux such as Ext2 and Ext3). 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 has the **System** attribute, all of its contents (including files that do not have the **System** attribute) will be excluded.

Tip: 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 matching any of the criteria. Use the **Add**, **Edit**, **Remove** and **Remove All** buttons to create the list of criteria.

The criteria are *not* case-sensitive in Windows and Linux. For example, if you choose to exclude all *.tmp* files and the *C:\Temp* folder, also excluded will be all *.Tmp* files, all *.TMP* files, and the *C:\TEMP* folder.

Criteria: full path

Specify the full path to the file or folder, starting with the drive letter (when backing up Windows) or the root directory (when backing up Linux).

Both in Windows and Linux, you can use a forward slash in the file or folder path (as in **C:/Temp** and **C:/Temp/File.tmp**). In Windows, you can also use the traditional backslash (as in **C:\Temp** and **C:\Temp\File.tmp**).

Under a Windows-style bootable media, a volume might have a different drive letter than in Windows. For more information, see "Working under bootable media" (p. 156).

Criteria: name

Specify the name of the file or folder, such as *Document.txt*. All files and folders with that name will be excluded.

Wildcard characters

You can use one or more wildcard characters ***** and **?** in the criterion. These characters can be used both within the full path and in the file or folder name.

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

The question mark (**?**) substitutes for exactly one character in a file name. For example, the criterion *Doc?.txt* covers 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	C:\Finance\F.log	Excludes the file named "F.log" located in the folder C:\Finance
By folder path	C:\Finance\F or C:\Finance\F\	Excludes the folder C:\Finance\F (be sure to specify the full path starting from the drive letter)
Linux		
By file path	/home/user/Finance/F.log	Excludes the file named "F.log" located in the folder (directory) /home/user/Finance
By folder path	/home/user/Finance or /home/user/Finance/	Excludes the folder (directory) /home/user/Finance

4.2.4 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. 41).

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.

4.2.4.1 Selecting backup destinations

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

Destination	Details
 Cloud storage	<p>To back up data to Acronis Cloud Storage, click Log in and specify the credentials to log in to the cloud storage. Then, expand the Cloud storage group and select the account.</p> <p>Prior to backing up to the cloud storage, you need to buy a subscription (p. 192) to the cloud backup service and activate (p. 193) the subscription on the machine(s) you want to back up.</p> <p>Cloud backup is not available under bootable media.</p> <hr/> <p>Note <i>Acronis Cloud Backup might be unavailable in your region. To find more information, click here: http://www.acronis.com/en-us/my/cloud-backup/corporate</i></p>
 Personal	<p>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 to the system.</p>
 Machine	<p>Local machine</p>
 Local folders	<p>To back up data to a local folder of the machine, expand the <Machine name> group and select the required folder.</p>
 CD, DVD, BD	<p>To back up data to optical media such as CD, DVD, or Blu-ray Discs (BD), expand the <Machine name> group, then select the required drive.</p>
 RDX, USB	<p>To back up data to an RDX drive or USB flash drive, expand the <Machine name> group, then select the required drive. For information about using these drives, see the "Removable devices" (p. 135) section.</p>

Destination	Details
 Network folders	<p>To back up data to a network folder, expand the Network folders group, select the required networked machine, and then click the shared folder.</p> <p>If the network share requires access credentials, the program will ask for them.</p> <hr/> <p>Note: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as <i>/mnt/share</i>, select this mount point instead of the network share itself.</p>
 FTP, SFTP	<p>To back up data to FTP or SFTP, type the server name or address in the Path field as follows:</p> <p>ftp://ftp_server:port _number or sftp://sftp_server:port number</p> <p>To establish an active mode FTP connection, use the following notation:</p> <p>aftp://ftp_server:port _number</p> <p>If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.</p> <p>After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.</p> <p>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.</p> <hr/> <p>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.</p>
 NFS folders	<p>To back up data to an NFS share, expand the NFS folders group and click the folder.</p>

4.2.5 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

- 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.
 - Confirm password.** Re-enter the password.
- 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.

4.2.6 Backup schemes

Choose one of the available backup schemes:

- **Simple** – to schedule when and how often to backup data and specify retention rules.
- **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.
- **Initial seeding** – to save locally a full backup whose final destination is Acronis Cloud Storage.

4.2.6.1 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. 60) 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. 72).

Backup type

To access this setting, click **Show backup type, validation, convert to virtual machine**.

Select the backup type.

- **Full** - selected by default for all backup locations (except for Acronis Cloud 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 Cloud Storage.

Note: When the **Incremental** backup type is selected along with retention rules, the archive will be cleaned up using consolidation (p. 203), which is a more time-consuming and resource-intensive operation.

4.2.6.2 Grandfather-Father-Son scheme

At a glance

- Daily ("Son"), weekly ("Father"), 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 on the last Friday of each month, and weekly backups run on all other Fridays. As a result, you will normally obtain 12 monthly backups over a full year.

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 of the week when a backup will be performed. The default value is Workdays .
Weekly/Monthly	Specifies which day of the week (out of the days selected in the Back up on field) you want to reserve for weekly and monthly backups. The default value is Friday . With this value, a monthly backup will run on the last Friday of each month. Weekly backups will run on all other Fridays. If you choose a different day of week, these rules will apply to the day chosen.

Keep backups	<p>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.</p> <p>The default values for each backup type are as follows.</p> <p>Daily: 5 days (recommended minimum)</p> <p>Weekly: 7 weeks</p> <p>Monthly: indefinitely</p> <p>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.</p> <p>We recommend setting a retention period of at least one week for daily backups.</p>
Backup type	<p>Specifies the types of daily, weekly and monthly backups.</p> <ul style="list-style-type: none"> ▪ Always full - all the daily, weekly and monthly backups will always be full. ▪ Full/Differential/Incremental - daily backups are incremental, weekly backups are differential, and monthly backups are full. <p>The first backup is always full. However, this does not mean that it is a monthly backup. It will be kept as a daily, weekly or monthly backup, depending on the day of week it is created.</p>

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 backup, marked with the  icon, for a few days past its expected expiration date.

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.

4.2.6.3 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.

Parameter	Meaning
Clean up archive	<p>Specifies how to get rid of old backups: either to apply retention rules (p. 73) regularly or clean up the archive during a backup when the destination location runs out of space.</p> <p>By default, the retention rules are not specified, which means older backups will not be deleted automatically.</p> <p>Using retention rules</p> <p>Specify the retention rules and when to apply them.</p> <p>This setting is recommended for backup destinations such as shared folders.</p> <p>When there is insufficient space while backing up</p> <p>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:</p> <ul style="list-style-type: none"> ▪ 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 <p>This setting is recommended when backing up to a USB drive or Acronis Secure Zone. This setting is not applicable to FTP and SFTP servers.</p> <p>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.</p>
Apply retention rules (only if the retention rules are set)	<p>Specifies when to apply the retention rules (p. 73).</p> <p>For example, the cleanup procedure can be set up to run after each backup, and also on schedule.</p> <p>This option is available only if you have set at least one retention rule in Retention rules.</p>
Cleanup schedule (only if On schedule is selected)	<p>Specifies a schedule for archive cleanup.</p> <p>For example, the cleanup can be scheduled to start on the last day of each month.</p> <p>This option is available only if you selected On schedule in Apply retention rules.</p>
2nd location, 3rd location, and so on	<p>Specifies where to copy or move (p. 69) the backups from the current location.</p> <p>This option is available only if you selected either the Replicate newly 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.</p>

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. 73).

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. 73).

4.2.6.4 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. 61), weekly (p. 63), or monthly (p. 65) 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 <ul style="list-style-type: none"> ▪ Always full - all levels of backups will be full. ▪ Full/Differential/Incremental - backups of different levels will have different types: <ul style="list-style-type: none"> - Last-level backups are full - Backups of intermediate levels are differential - 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.

4.2.6.5 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 Cloud 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 Cloud Storage.
- **Differential**. At the first time a full backup will be created. The next backups will be differential.

4.2.6.6 Initial seeding

This backup scheme is available when Acronis Cloud Storage is selected as the backup destination. A backup is only successful if you have an Initial Seeding license.

The Initial Seeding service might be unavailable in your region. To find more information, click here:
<http://kb.acronis.com/content/15118>.

Initial seeding enables you to transfer the first backup, which is full and usually the largest, to the cloud 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 cloud storage.

If you back up 500 GB of data or more, initial seeding ensures faster delivery of the backed-up data and lower traffic costs.

Please refer to the "Initial Seeding FAQ (p. 182)" section for more details.

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.

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.

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 the archive will validate all the archive's backups and may take a long time and a lot of system resources.
Validation of the latest backup may also take time, even if this backup is incremental or differential, and small in size. This is because the operation validates not only the data physically contained in the backup, but all of the data recoverable by selecting the backup. This requires access to previously created backups.
3. **Validation schedule** (appears only if you have selected **On schedule** in step 1) - set the schedule of validation. For more information see the Scheduling (p. 60) section.

4.2.8 Backup plan's credentials

Provide the credentials for the account under which the plan will run. By default, the plan runs under the agent service account, if created by a user having administrative privileges on the machine. If created by a regular user, such as a member of the **Users** group, the plan runs under this user's account.

To specify credentials explicitly

1. If you have administrative privileges on the machine, select **Use the following credentials**. Otherwise skip this step.
2. 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.
 - **Confirm password**. Re-enter the password.
3. Click **OK**.

To learn more about operations available depending on the user privileges, see the Users' privileges on a managed machine (p. 24) 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. 96) the machine to a VMware ESX(i) using Agent for VMware, or convert (p. 120) 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. 35) 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 Sequence of operations in a backup plan

If a backup plan contains multiple operations, Acronis Backup performs them in the following order:

1. Cleanup (if configured **Before backup**) and validation (if cleanup has been performed and validation is configured to run **After the retention rules are applied**).
If a backup was moved to a different location during the cleanup, all the operations configured for the subsequent locations are performed before continuing to the following steps in the primary location.
2. Pre-backup command execution.
3. Backup:
 - a. Pre-data capture command execution
 - b. Snapshot creation
 - c. Post-data capture command execution
 - d. Backup process
4. Start of backup cataloging.
Backup cataloging can be a time-consuming process. It is performed in parallel with the following steps.
5. Post-backup command execution.
6. Disaster Recovery Plan (DRP) creation.
7. Conversion to a virtual machine.
8. Backup replication.
9. Cleanup.
If the replication took place, or a backup was moved to a different location during the cleanup, all the operations configured for the subsequent locations are performed before continuing to the following steps in the primary location.
10. Validation.
11. Sending e-mail notification.

4.2.11 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

To use simplified naming of backup files, do either of the following:

- In the welcome screen, click **Create backup plan** (p. 35), expand **Show backup file naming, archive comments**, and then select the **Name backup files using the archive name...** check box.
When you back up to a locally attached RDX drive or USB flash drive, the **Name backup files using the archive name...** check box does not appear. Instead, the removable device mode (p.

135) determines whether the standard or simplified naming scheme will be used. In Linux, the check box appears after you manually mount the device.

- In the welcome screen, click **Back up now** (p. 35). Simplified naming will be used whenever the backup destination supports it (see “Restrictions” below).

When you use simplified file naming

- 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. 58) or all backups created during a week (p. 59). 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. 56) to the archive name.

When you use standard file naming

- 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

Simplified file naming is not available if you back up to Acronis Secure Zone or Acronis Cloud Storage.

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.
- Setting up replication of backups.
- Setting up retention rules.
- Setting up regular conversion of backups to a virtual machine.
- Converting an incremental or differential backup into a full one.

Restrictions on archive names

- The archive name cannot end with a number.
- 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 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. You can also use a script to copy, move, or delete the older backups.

The value of this variable is the current date surrounded by brackets ([]). The date format depends on the regional options on the machine. For example, if the date format is *year-month-day*, the value for January 31, 2012, is **[2012-01-31]**. Characters that are not supported in a file name, such as slashes (/), are replaced with underscores (_).

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, 2012. The archive name is **MyArchive-[DATE]**, the date format is *year-month-day*. Here is the list of backup files after day two:

MyArchive-[2012-01-31].tib (full, created on January 31 at midnight)
MyArchive-[2012-01-31]2.tib (incremental, created on January 31 at noon)
MyArchive-[2012-02-01].tib (full, created on February 1 at midnight)
MyArchive-[2012-02-01]2.tib (incremental, created on February 1 at noon)

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

MyArchive-[2012-01-31].tib (full, created on January 31 at noon)
MyArchive-[2012-02-01].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.

4.3.2 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**.

4.3.3 Usage examples

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

4.3.3.1 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 on a locally attached USB hard drive 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 the USB hard drive as the archive location, 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.

If you choose to back up to a locally attached RDX drive or USB flash drive, you will not see the **Name backup files using the archive name...** check box. Instead, make sure that the removable device mode (p. 135) is set to **Removable media**.

4.3.3.2 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, 2012, January 2, 2012, and so on, are stored respectively as **MyMachine-[2012-01-01].tib**, **MyMachine-[2012-01-02].tib**, and so on.
- Your script can move older backups based on the date stamp.

See also "The [Date] variable" (p. 56).

4.3.3.3 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, 2012, will be stored as **ServerFiles[2012-01-01].tib**, **ServerFiles[2012-01-01]2.tib**, and so on up to **ServerFiles[2012-01-01]24.tib**.
- The following day, the backups will start with the full backup **ServerFiles[2012-01-02].tib**.

See also "The [Date] variable" (p. 56).

4.3.3.4 Example 4. Daily full backups with daily drive swaps

Consider the following scenario:

- You want to perform a daily full backup of your machine.
- You want to store the backup on a locally attached USB hard drive in the file **MyMachine.tib**.
- You have two such drives. You want to swap them 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.
- Create a directory such as **/mnt/backup** and specify it as the archive location. Each time you attach a drive, make sure to mount it to the mount point **/mnt/backup**.
- Select the **Name backup files using the archive name...** check box.
- 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.

4.3.3.5 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 store the backups on a locally attached USB hard drive in the archive **MyMachine**.
- You have two such drives. You want to swap them 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:

- When creating the first backup plan:
 - Specify **MyMachine** as the archive name.
 - Create a directory such as **/mnt/backup** and specify it as the archive location. Each time you attach a drive, make sure to mount it to the mount point **/mnt/backup**.
 - Select the **Name backup files using the archive name...** check box.
 - Select **Full** as the backup type.
 - Schedule the backups to run every week on Monday.
- 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.

Each time you attach a drive, make sure to mount it to the mount point **/mnt/backup**.

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.

4.3.3.6 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, 2012, will be stored as ServerFiles[2012-01-31].tib.
- The 10 differential backups of January 31, 2012, will be stored as ServerFiles[2012-01-31]2.tib, ServerFiles[2012-01-31]3.tib, and so on up to ServerFiles[2012-01-31]11.tib.
- The following day, February 1, the backups will start with the full backup ServerFiles[2012-02-01].tib. The differential backups will start with ServerFiles[2012-02-01]2.tib.

See also “The [Date] variable” (p. 56).

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. 35) with any of the following backup schemes: Simple, Custom or Tower of Hanoi. The schedule also can be set for validation tasks (p. 137).

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 Linux operating system.

Events
Time: Daily, Weekly, Monthly
Time since completion of the last successful backup within the same backup plan (specify the length of time)
System startup

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 Linux operating system.

Condition: run the task only if
Location's host is available
The task run time is within the specified time interval
The specified period of time has passed since the completion of the last successful backup within the same backup plan

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: <...> 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.

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: <...> 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.

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. 47) 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. 47) 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 until 9 PM
- Wednesday: every 3 hours from 9 AM until 9 PM
- Thursday: every 3 hours from 9 AM until 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**.

- Effective:
From: **not set**.
To: **not set**.

Second schedule

- Every **1** week(s) on: **Tue, Wed, Thu**.
- Every **3** hours
From **09:00:00 AM** until **09:00:00 PM**.
- Effective:
From: **not set**.
To: **not set**.

Third schedule

- Every: **1** week(s) on: **Sat, Sun**.
- Once at: **09:00:00 PM**.
- 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: <...> 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.

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. 47) 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) until 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 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. 60)" section for the list of available events), the scheduler checks the specified condition and executes the task if the condition is met.

Conditions are available only when the custom backup scheme (p. 47) is used. You can set conditions for full, incremental and differential backup separately.

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.

Adding multiple conditions

If two or more conditions are specified, the backup will start only when all of them are met.

4.4.4.1 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.

4.4.4.2 Fits the 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 the 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.

4.4.4.3 Time since last backup

Applies to: Windows, Linux

Postpones a backup until the specified time passes since the completion of the last successful backup within the same backup plan.

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. 35), 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. 74)
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. 74)
Delete older backups from a fast storage according to retention rules, in order to not overuse expensive storage space.
- **Using Acronis Cloud Backup to protect data from a natural disaster** (p. 75)
Replicate the archive to the cloud storage by transferring only the data changes outside working hours.
- **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
Manual start (p. 52)	Yes	No	No
Simple (p. 43)	Yes	Yes	Yes
Grandfather-Father-Son (GFS) (p. 44)	Yes	No	Yes
Tower of Hanoi (p. 50)	Yes	No	Yes
Custom (p. 47)	Yes	Yes	Yes
Initial seeding (p. 52)	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. 55), 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 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 drive
- A network folder
- An FTP or SFTP server
- Acronis Cloud Storage
- A removable device (p. 135) used in the **Fixed drive** mode. (You select the removable device mode when creating a backup plan.)

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.
- Copying or moving backups *to and from* removable devices used in the **Removable media** mode is not supported.
- Acronis Cloud Storage can only be the final location. Further copying or moving backups *from* it is not possible.
- 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. 35).

- To set up replication from the primary location, select the **Replicate newly 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.

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

A backup is replicated to the next location as soon as it appears in the previous location. If earlier backups were not replicated (for example, the network connection was lost), the software also replicates all of the backups that appeared after the last successful replication.

4.5.3 Setting up retention of backups

You can set retention rules for backups when creating a backup plan (p. 35). The available retention rules depend on the chosen backup scheme.

Applying retention rules can be restricted by the **Replication/cleanup inactivity time** (p. 93) 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. 50), 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. 73), 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. 73), 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, 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. 47) 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 (marked with the  icon) 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.

This mode is not available for Acronis Cloud Storage when you copy or move backups there. In the cloud storage, all backups are incremental except the first backup of an archive which is always full. This chain cannot be entirely deleted because the most recent backup must always be kept.

- **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.

This mode is not available if you selected the **Archive size greater than** rule for any archive location except for Acronis Cloud Storage.

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 Usage examples

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

4.5.5.1 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. 132) 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 **Location** field, select **Full** in the **Backup type** field, select the **Replicate newly 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.

4.5.5.2 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.

4.5.5.3 Example 3. Replicating backups to the cloud storage

This example assumes that you have activated (p. 193) a cloud backup subscription (p. 179) for the machine that you are backing up.

The following scenario assumes that the amount of data you want to back up is relatively small. For larger backups, see “Replicating large amounts of data to the cloud storage” later in this section.

Consider the following scenario:

- You want to occasionally back up your machine to a local folder.
- You want to keep a copy of the resulting archive off-site in Acronis Cloud Storage.
- No matter when you start the backup, you want the replication to take place outside working hours, when demand on the Internet connection is lower.

In this scenario, create a backup plan with the desired backup scheme. When creating the backup plan, specify a local folder in the **Location** field. Select the **Replicate newly created backup to another location** check box, and then specify the cloud storage in the **2nd location** field.

In **Backup options**, go to **Replication/cleanup inactivity time** (p. 93), and specify the working hours (for example, Monday through Friday from 8:00 until 17:00).

Result:

- After the backup plan starts, the data is backed up to the local folder.
- If the backup finishes outside the working hours, replication starts immediately. Otherwise, replication is postponed until the end of the working hours.

Note: *In the cloud storage, the second and further backups of an archive will always be incremental, no matter what type they are in the original location. This leads to efficient use of storage space for your cloud backup subscription.*

Replicating large amounts of data to the cloud storage

If you are planning to back up 500 GB of data or more, you may want to send the first backup to the cloud storage on a physical hard drive. This option is provided by the Initial Seeding service (p. 182) which you can buy in addition to your cloud backup subscription.

The Initial Seeding service might be unavailable in your region. To find more information, click here: <http://kb.acronis.com/content/15118>.

During the subsequent backups, only changes to the original data will be sent to the cloud storage and will not affect network traffic as much.

In this scenario, create a backup plan with the **Initial seeding** scheme. When creating the backup plan, specify a local folder in the **Location** field. This can be a folder on the hard drive that you are going to send. For more details, see “How to perform initial seeding?” (p. 184).

After you have sent the hard drive and the order status becomes **The data upload has been completed**, edit the backup plan. Change the backup scheme, destination, and replication settings to those previously described in this section.

The updated backup plan will produce backups that will be replicated to the cloud storage outside working hours.

4.6 How to disable backup cataloging

Cataloging a backup adds the contents of the backup to the data catalog as soon as the backup is created. This process can be time-consuming. Therefore, you may want to disable cataloging on a managed machine. To do it, go to **Options > Machine options** and configure the **Backup cataloging** option.

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 (Linux, 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 Linux		Bootable media (Linux-based or PE-based)	
	Disk backup	File backup	Disk backup	File backup
Additional settings (p. 78):				

	Agent for Linux		Bootable media (Linux-based or PE-based)	
	Disk backup	File backup	Disk backup	File backup
Ask for the first media while backing up to removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media
Reset archive bit	-	-	-	+
Restart the machine automatically after backup is finished	-	-	+	+
Archive protection (p. 79) (password + encryption)	+	+	+	+
Backup cataloging (p. 80)	+	+	-	-
Backup performance:				
Backup priority (p. 80)	+	+	-	-
HDD writing speed (p. 81)	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD
Network connection speed (p. 81)	Dest: network share	Dest: network share	Dest: network share	Dest: network share
Backup splitting (p. 81)	+	+	+	+
Compression level (p. 82)	+	+	+	+
Disaster recovery plan (p. 83)	+	+	-	-
E-mail notifications (p. 84)	+	+	-	-
Error handling (p. 85):				
Do not show messages and dialogs while processing (silent mode)	+	+	+	+
Re-attempt if an error occurs	+	+	+	+
Ignore bad sectors	+	+	+	+
Event tracing:				
SNMP (p. 86)	+	+	-	-
Fast incremental/differential backup (p. 86)	+	-	+	-
File-level backup snapshot (p. 87)	-	+	-	-
LVM snapshotting (p. 87)	+	-	-	-
Media components (p. 89)	Dest: removable media	Dest: removable media	-	-

	Agent for Linux		Bootable media (Linux-based or PE-based)	
	Disk backup	File backup	Disk backup	File backup
Pre/Post backup commands (p. 89)	+	+	PE only	PE only
Pre/Post data capture commands (p. 91)	+	+	-	-
Replication/cleanup inactivity time (p. 93)	+	+	-	-
Sector-by-sector backup (p. 93)	+	-	+	-
Task failure handling (p. 93)	+	+	-	-
Task start conditions (p. 94)	+	+	-	-

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: **Disabled**.

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 will reset the archive bits of all files being backed up. Acronis Backup 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 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.

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 specifies whether full or fast cataloging will be performed on a backup as soon as the backup is created.

This option is effective only if backup cataloging is enabled on the backed-up machine or on the storage node.

The preset is: **Full cataloging**.

If you select **Full cataloging**, the backup contents are cataloged to the highest possible level of detail. This means that the following data will be displayed in the catalog:

- For a disk-level backup - disks, volumes, files, and folders.
- For a file-level backup - files and folders.

You may want to select **Fast cataloging** if the full cataloging tends to affect the performance of the managed machine or if your backup window is too narrow. 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 contents of already existing backups to the catalog, you can start the full cataloging manually when appropriate.

For more information about using the data catalog, see the "Data catalog" (p. 99) 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 for 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.

This option is not effective when the backup destination is a managed vault or Acronis Cloud Storage.

The option defines how a backup can be split.

The preset is: **Automatic**

The following settings are available.

Automatic

With this setting, Acronis Backup will act as follows.

- **When backing up to a hard disk or a network share:**

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. This might be the case when the backup is placed on FAT16 and FAT32 file systems that have a 4-GB 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, Blu-Ray Discs, an RDX or USB drive used in the removable device (p. 135) mode):

The task will enter the **Need interaction** state and ask for a new media when the previous one is full.

- **When backing up to an FTP server:**

The backup will automatically be split into files no more than 2 GB in size. Splitting is needed to allow data recovery directly from the FTP server.

- **When backing up to an SFTP server:**

A single backup file will be created. If the destination storage runs out of free space while creating the backup, the task will fail.

When you replicate or move a backup (p. 69) to other locations, these rules apply to each location independently.

Example.

Suppose that the primary location for a 3-GB backup is a hard disk, the second location is an FTP server, and the third location is a network share. In this case, the backup will be stored as a single file in the primary location, as two files in the second location, and as a single file again in the third location.

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 a backup into 2-GB files if you are backing up to a hard disk, but you plan to manually copy the backup to an FTP server later.

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.

This option is not effective for file-level backups.

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.

A DRP is created after the first successful backup is performed by the backup plan. If the **Send disaster recovery plans** option is enabled, the DRP is sent by e-mail to the specified list of users. If the **Save DRP as file** option is enabled, the DRP is saved as a file to the specified location. The DRP will be created again 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.)

You can specify a local folder, a network folder, an FTP or SFTP server as a location to save the DRPs.

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.xml` - in 32-bit Windows
- `%ProgramFiles(x86)%\Acronis\BackupAndRecovery\drp.xml` - in 64-bit Windows
- `/usr/lib/Acronis/BackupAndRecovery/drp.xml` - in Linux

To set up sending DRPs:

1. Select the **Send disaster recovery plans** 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.
4. Enter the parameters of access to the SMTP server. For more detailed information, see E-mail notifications (p. 113).
5. [Optional] Click **Send test e-mail message** to check if the settings are correct.

To set up saving DRPs as files:

1. Select the **Save DRP as file** check box.
2. Click **Browse** to specify a location for the DRP files.

4.7.8 E-mail notifications

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 when user interaction is required.

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.**
 - **When backup fails.**
 - **When user interaction is required.**
3. Select the **Add full log to notification** check box if you want the e-mail notification to include log entries for the operation.
4. In the **E-mail addresses** field, type the destination e-mail address. You can enter several addresses separated by semicolons.
5. In the **Subject** field, type the notification subject.

The subject can include ordinary text and one or more variables. In the received e-mail messages, each variable will be replaced by its value at the time of task execution. The following variables are supported:

- **%description%**

For a machine running Windows, the **%description%** variable will be replaced by the text that is given in the **Computer description** field of the machine. To specify this text, either go to **Control panel > System** or run the following command as an administrator:

```
net config server /srvcomment:<text>
```

For a machine running Linux, the **%description%** variable will be replaced by an empty string ("").

- **%subject%**

The **%subject%** variable will be replaced by the following phrase: *Task <task name> <task result> on machine <machine name>*.

6. In the **SMTP server** field, enter the name of the outgoing mail server (SMTP).
7. In the **Port** field, set the port of the outgoing mail server. By default, the port is set to **25**.

8. If the outgoing mail server requires authentication, enter **User name** and **Password** of the sender's e-mail account.
If the SMTP server does not require authentication, leave the **User name** and **Password** fields blank. If you are not sure whether the SMTP server requires authentication, contact your network administrator or your e-mail service provider for assistance.
9. Click **Additional e-mail parameters...** to configure additional e-mail parameters as follows:
 - a. **From** – type the name of the sender. If you leave this field empty, the messages will contain the sender's e-mail account in the **From** field.
 - 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** and **Password** of the incoming mail server.
 - d. Click **OK**.
10. Click **Send test e-mail message** to check whether e-mail notifications work correctly with the specified settings.

4.7.9 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 are 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.

*If Acronis Cloud Storage is selected as the primary, the second, or a further backup location, the option value is automatically set to **Enabled. Number of attempts: 300**, regardless of the default value.*

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.

4.7.10 Event tracing

It is possible to send log events of the backup operations, performed on the managed machine, to the specified SNMP managers.

4.7.10.1 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, please see "Support for SNMP (p. 33)".

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.

4.7.11 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.

4.7.12 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.

4.7.13 LVM snapshotting

This option is effective only for Linux operating systems. This option is effective for both disk-level and file-level backup of volumes managed by Linux Logical Volume Manager (LVM). Such volumes are also called logical volumes.

This option defines how a snapshot of a logical volume is taken. Acronis Backup can do this on its own or rely on Linux Logical Volume Manager (LVM). Use of a snapshot ensures a time-consistent backup of a volume whose data may change during the backup process.

The preset is: **Acronis Backup**

Select one of the following:

- **Acronis Backup**

Acronis Backup will take the snapshot on its own. With this setting, a backup normally proceeds faster and unallocated space on the volume group is not required. Therefore, we recommend changing the preset only if you are experiencing problems with backing up logical volumes.

▪ Logical volume manager

Acronis Backup will use the snapshot taken by LVM. This snapshot is stored on unallocated space of the volume group. If the unallocated space is missing, Acronis Backup will take the snapshot on its own.

A detailed explanation of LVM snapshotting

Once a volume snapshot is taken, and the data starts changing, the old data has to be kept somewhere until it is saved to the backup.

- Acronis keeps the old data mostly in the RAM. (During a file-level backup, the software may create a temporary file in **/tmp** if the old data size significantly grows.)
- LVM needs a temporary logical volume (a snapshot logical volume) for keeping the old data (see http://tldp.org/HOWTO/LVM-HOWTO/snapshots_backup.html). Writing this data onto the volume causes a large number of disk I/O operations. This is why the backup is usually slower when a snapshot is taken by LVM.

If you choose to use LVM, Acronis Backup creates a snapshot logical volume on its own. The software proceeds as follows:

1. Checks the backed-up volume size (not data size, but volume size).
2. Calculates 10 percent of that size—say, 10 GB.
3. Checks if there is at least 10 GB of unallocated space on the corresponding volume group.
4. If there is, creates a 10-GB snapshot logical volume (using the **lvcreate -s** command) and starts backing up. Otherwise, the software backs up using the Acronis snapshot.
5. Deletes the snapshot volume as soon as it is no longer needed.

Multiple logical volumes are backed up one by one. For each of them, the software creates a separate snapshot volume of the corresponding size. Only one snapshot volume is stored at a time.

If the data changes extremely fast while the software is keeping the snapshot, the snapshot logical volume will run out of space and the backup will fail. That is why you can change the 10 percent default value to any value up to 100 percent (which guarantees success) in

/etc/Acronis/BackupAndRecovery.config.

To change the default size of a snapshot logical volume:

1. Decide how much unallocated space you want to use. If you are backing up two or more logical volumes, base your choice on the size of the biggest of them.

*Tip: To view the amount of unallocated space on a volume group, run the **vgdisplay** command and then examine the **Free PE / Size** line. To view the sizes of logical volumes, run the **lvdisplay** command and then examine the **LV Size** lines.*

2. Open the file **/etc/Acronis/BackupAndRecovery.config** in a text editor.
3. Locate the line **<value name="MMSDirPath" type="TString">**.
4. Add the following fragment directly before that line:

```
<key name="LVMSnapshots">
  <value name="MinimalVGFreeRelative" type="TString">
    20
  </value>
</key>
```

In this example, the new value is 20 percent. The number must be integer.

5. Save the file. The new setting will take effect at the next backup. No service restart is needed.

4.7.14 Media components

This option is effective for both Windows and Linux operating systems, when the backup destination is CD, DVD, or Blue-ray Disc (BD).

When backing up to this media, you can make this media work as regular Linux-based bootable media (p. 201) 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 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. 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.

4.7.15 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.



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.
- Selectively copy backups from an archive to another location. This option may be useful because the replication configured in a backup plan copies *every* backup of an archive to subsequent locations.

Acronis Backup performs the replication *after* executing the post-backup command. For more information see "Sequence of operations in a backup plan" (p. 55).

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**.

4.7.15.1 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 only after the command is successfully executed. Fail the task if the command execution fails.	Perform the backup after the command is executed despite execution failure or success.	N/A	Perform 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.

4.7.15.2 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.

4.7.16 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 at the beginning of the backup procedure.

The following scheme illustrates when the pre/post data capture commands are executed.



If the Volume Shadow Copy Service 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. 89), 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**.

4.7.16.1 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 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.

4.7.16.2 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 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.

4.7.17 Replication/cleanup inactivity time

This option is effective only if you set up replication or retention rules (p. 69) 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.

4.7.18 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.

4.7.19 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.

N=3: 2nd attempt succeeded



N=3: none of attempts succeeded



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.

4.7.20 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. 60) and Conditions (p. 67).

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

Time interval > waiting for condition



Time interval < waiting for condition



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.

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 managed machine's **operating system fails to start** or you need to **recover data to bare metal**, boot the machine from the bootable media (p. 201) or using Acronis Startup Recovery Manager (p. 158). Then, create a recovery task.

The most recent operating systems remain bootable when recovered to dissimilar hardware, or the VMware or Hyper-V platforms. If a recovered operating system does not boot, use the Acronis Universal Restore tool to update the drivers and modules that are critical for the operating system startup. For details, see the Acronis Universal Restore page at <http://www.acronis.com/en-us/aur/>.

For detailed information about recovering Linux Software RAID devices and volumes created by Logical Volume Manager (LVM), see "Recovering MD devices and logical volumes" (p. 28).

Acronis Backup Agent for Linux has the ability to recover a disk (volume) backup to a new virtual machine. See the "Recovery to the "New virtual machine" destination" (p. 123) section for more details.

5.1 Creating a recovery task

To create a recovery task, perform the following steps

What to recover

Select data (p. 97)

Select data to recover.

Access credentials (p. 100)

[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. 101)

Volumes (p. 103)

Files (p. 106)

Access credentials (p. 101)

[Optional] Provide credentials for the destination if the task credentials do not enable recovery of the selected data. To access this option, select the **Advanced view** check box.

When to recover

Recover (p. 107)

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. 110) 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 option, click **Show task credentials**.

After you complete all the required steps, click **OK** to create 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. 98).

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.

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. 99).

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 not displayed (for example, if the archive metadata has been lost), click **Refresh**.

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 recovers Track 0, which does not affect the target disk's partition table and partition layout. Acronis Backup 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
 Cloud storage	If the archive is stored in Acronis Cloud Storage, click Log in and specify the credentials to log in to the cloud storage. Then, expand the Cloud storage group and select the account. <hr/> <i>Exporting and mounting are not supported for backups stored in Acronis Cloud Storage.</i>
 Personal	If the archive is stored in a personal vault, expand the Personal group and click the required vault.
 Machine name	Local machine
 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, BD	If the archive is stored on optical media such as CD, DVD, or Blu-ray Discs (BD), expand the <Machine name> group, then select the required drive. First insert the last disc. Then insert the discs in order starting from the first one when the program prompts.
 RDX, USB	If the archive is stored on an RDX drive or USB flash drive, expand the <Machine name> group, then select the required drive. For information about using these drives, see the "Removable devices" (p. 135) section.
 Network folders	If the archive is stored on 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. <hr/> <i>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.</i>

Location	Details
 FTP, SFTP	<p>If the archive is stored on an FTP or SFTP server, type the server name or address in the Path field as follows:</p> <p>ftp://ftp_server:port _number or sftp://sftp_server:port number</p> <p>To establish an active mode FTP connection, use the following notation:</p> <p>aftp://ftp_server:port _number</p> <p>If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.</p> <p>After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.</p> <p>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.</p> <hr/> <p><i>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.</i></p>
 NFS folders	If the archive is stored on an NFS share, expand the NFS folders group and click the folder.

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.

Acronis Backup may upload data catalog files from a vault to a local cache folder. By default, this folder is located on the disk where the operating system is installed. For information about changing the default cache folder, refer to the "Changing the default cache folder for catalog files" section.

Selecting the backed-up data for recovery

- To access the **Data view** tab, navigate to **Vaults** view, and click the required vault.
- 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.
- In the **Display data backed up for** field, specify the time period for which the backed-up data will be displayed.
- 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. Select the required data and click **OK** to return to the **Data view**.
- 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 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 disabled or fast cataloging is turned on

If the data is displayed partially or is not displayed at all, most likely cataloging was disabled or the fast cataloging (p. 80) was turned on during backup.

Solutions:

- If cataloging is disabled, enable it in the **Backup cataloging** option (**Options > Machine options**).
- Run the full cataloging manually by clicking **Catalog now**. 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. 129) 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, RDX or USB devices.
- Data backed up to Acronis Cloud 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.

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 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.

Disks/volumes

Map automatically

Acronis Backup attempts to map the selected disks to the target disks as described in the "How the automatic mapping works" (p. 103) 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. 102)

Select the destination disk for each of the source disks.

NT signature (p. 102)

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 will generate a new NT signature for the target hard disk.

- **Recover from backup**

Acronis Backup 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 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.
- **Keep existing**
The program will leave the NT signature of the target hard disk untouched.

How the automatic mapping works

Acronis Backup 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 (p. 28).

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 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.

Disks/volumes

Map automatically

Acronis Backup attempts to map the selected volumes to the target disks as described in the "How the automatic mapping works" (p. 103) 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. 104)

Choose the disk to recover the Master Boot Record to.

NT signature: (p. 102)

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. 104)

[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 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 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 a hard 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 with 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.

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.

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 files** – this will give the file in the backup priority over the file on the hard disk.
- **Overwrite an 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 files** – 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. 106)

Specify files and folders you do not wish to be recovered.

Recovery exclusions

Set up exclusions for the specific files and folders you do not wish to recover.

Note: Exclusions override selection of data items to recover. For example, if you select to recover file *MyFile.tmp* and to exclude all *.tmp* files, file *MyFile.tmp* will not be recovered.

Use the **Add**, **Edit**, **Remove** and **Remove All** buttons to create the list of files and folders to exclude. Specify the name of the file or folder, such as *Document.txt*.

The names are *not* case-sensitive in Windows and Linux. For example, if you choose to exclude all *.tmp* files and the Temp folders, also excluded will be all *.Tmp* files, all *.TMP* files, and the *TEMP* folders.

You can use one or more wildcard characters *** and *?*:

- The asterisk (***) substitutes for zero or more characters. For example, *Doc*.txt* covers files such as *Doc.txt* and *Document.txt*.
- The question mark (*?*) substitutes for exactly one character. For example, *Doc?.txt* covers files such as *Doc1.txt* and *Docs.txt*, but not the files *Doc.txt* or *Doc11.txt*.

Exclusion examples

Criterion	Example	Description
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"

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:

- **Use current user credentials**

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, see the "Credentials used in backup plans and tasks" (p. 22) section.

To learn more about operations available depending on the user privileges, see the "User privileges on a managed machine" (p. 24) section.

5.2 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 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: The most recent operating systems remain bootable when recovered to dissimilar hardware, or the VMware or Hyper-V platforms. If a recovered operating system does not boot, use the Acronis Universal Restore tool to update the drivers and modules that are critical for the operating system startup. For details, see the Acronis Universal Restore page at <http://www.acronis.com/en-us/aur/>.
- **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 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.

5.2.1 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 website.

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.

5.3 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 (Linux, 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 Linux		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. 112):				
Validate backup archive before recovery	+	+	+	+
Restart the machine automatically if it is required for recovery	+	+	-	-
Restart the machine automatically after recovery is finished	-	-	+	+
Check file system after recovery	+	-	+	-
Set current date and time for recovered files	-	+	-	+
E-mail notifications (p. 113)	+	+	-	-
Error handling (p. 114):				
Do not show messages and dialogs while processing (silent mode)	+	+	+	+
Re-attempt if an error occurs	+	+	+	+
Event tracing:				
SNMP (p. 114)	+	+	-	-
File-level security (p. 115):				

	Agent for Linux		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)
Recover files with their security settings	-	+	-	+
Pre/Post recovery commands (p. 115)	+	+	PE only	PE only
Recovery priority (p. 117)	+	+	-	-

5.3.1 Additional settings

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

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.

Validate backups 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.

Check file system after recovery

This option is effective only when recovering disks or volumes.

The preset is **Disabled**.

This option defines whether to check the integrity of the file system after a disk or volume recovery. The check takes place either immediately after recovery or after the machine boots into the recovered operating system.

Restart the 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 the 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.

5.3.2 E-mail notifications

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 when user interaction is required.

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.**
 - **When recovery fails.**
 - **When user interaction is required.**
3. In the **E-mail addresses** field, type the destination e-mail address. You can enter several addresses separated by semicolons.
4. In the **Subject** field, type the notification subject.

The subject can include ordinary text and one or more variables. In the received e-mail messages, each variable will be replaced by its value at the time of task execution. The following variables are supported:

- **%description%**

For a machine running Windows, the **%description%** variable will be replaced by the text that is given in the **Computer description** field of the machine. To specify this text, either go to **Control panel > System** or run the following command as an administrator:

```
net config server /srvcomment:<text>
```

For a machine running Linux, the **%description%** variable will be replaced by an empty string ("").

- **%subject%**

The **%subject%** variable will be replaced by the following phrase: *Task <task name> <task result> on machine <machine name>*.

5. In the **SMTP server** field, enter the name of the outgoing mail server (SMTP).
6. In the **Port** field, set the port of the outgoing mail server. By default, the port is set to **25**.
7. If the outgoing mail server requires authentication, enter **User name** and **Password** of the sender's e-mail account.

If the SMTP server does not require authentication, leave the **User name** and **Password** fields blank. If you are not sure whether the SMTP server requires authentication, contact your network administrator or your e-mail service provider for assistance.
8. Click **Additional e-mail parameters...** to configure additional e-mail parameters as follows:
 - a. **From** – type the name of the sender. If you leave this field empty, the messages will contain the sender's e-mail account in the **From** field.

- 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** and **Password** of the incoming mail server.
 - d. Click **OK**.
9. Click **Send test e-mail message** to check whether e-mail notifications work correctly with the specified settings.

5.3.3 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 are 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.

5.3.4 Event tracing

It is possible to send log events of the recovery operations, performed on the managed machine, to the specified SNMP managers.

5.3.4.1 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 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, please see "Support for SNMP (p. 33)".

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.

5.3.5 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, you can choose whether to recover the permissions or let the files inherit the NTFS permissions from the folder to which they are recovered.

5.3.6 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**.

5.3.6.1 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 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.

5.3.6.2 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.

5.3.7 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 Conversion to a virtual machine

Acronis Backup offers a number of ways of converting a disk backup into a virtual machine. This section helps you choose the method that best fits your needs and provides step-by-step instructions for conversion.

6.1 Conversion methods

Depending on your needs, you can choose among the following conversion methods:

a) **Make the conversion a part of a backup plan**

When to use.

- If you want the backup and the conversion to be executed on a schedule. This helps you maintain a standby virtual server ready to power on in case your physical server fails.
- If you do not need to adjust the resulting virtual machine settings.

How to perform. When creating a backup plan (p. 35), enable conversion of a backup to a virtual machine (p. 120).

b) **Recover the backed-up disks or volumes to the "New virtual machine" destination**

When to use.

- If you want to do the conversion once or occasionally, as required.
- If you want to do a lossless physical to virtual migration. In this case, you boot the original machine from bootable media, back up the machine in the off-line state and immediately recover the machine from the resulting backup.
- If you need to adjust the resulting virtual machine settings. You may want to add or remove disks, choose the disk provisioning mode, change the volume sizes and their location on the disks, and more.

How to perform. Follow the steps described in "Recovery to the 'New virtual machine' destination" (p. 123).

c) **Recover the backed-up disks or volumes to a manually created virtual machine by using bootable media**

When to use.

- If you want to create a machine directly on a virtualization server, rather than import it.
Tip. With Agent for VMware or Agent for Hyper-V, a virtual machine can be created directly on a respective virtualization server by using methods (a) and (b).
- If you need to recreate dynamic volumes on a Windows machine.
- If you need to recreate logical volumes or software RAID on a Linux machine.

How to perform. Follow the steps described in "Recovery to a manually created virtual machine" (p. 126).

6.2 Conversion to an automatically created virtual machine

This section describes the conversion methods (p. 118) in which Acronis Backup automatically creates a new virtual machine:

- During conversion which is part of a backup plan (p. 120), the software creates the virtual machine in addition to creating the backup. The virtual machine has the same configuration as the original machine.
- During recovery to the "New virtual machine" destination (p. 123), the software creates the virtual machine from a backup you already have. You can change the configuration of the virtual machine.

Depending on the agent that performs the conversion, Acronis Backup can create a virtual machine of any of these formats:

Agent for Windows, Agent for Linux

- VMware Workstation
- Microsoft Virtual PC (includes Windows Virtual PC)
- Citrix XenServer OVA (only during recovery to the "New virtual machine" destination)
- Kernel-based Virtual Machine
- Red Hat Enterprise Virtualization (RAW format)

Agent for VMware

- VMware ESX(i)

Agent for Hyper-V

- Microsoft Hyper-V

6.2.1 Considerations before conversion

Converting a UEFI-based machine

Virtual machines that use Unified Extensible Hardware Interface (UEFI) are currently supported in VMware ESXi 5 only. If the target virtualization platform is ESXi 5, Acronis Backup creates a UEFI-based machine. Otherwise, the resulting machine will use the BIOS boot firmware.

Acronis Backup adjusts the Windows boot mode to the BIOS boot firmware and ensures that Windows remains bootable.

For Linux operating systems, changing the boot mode from UEFI to BIOS is not supported. When converting a UEFI-based machine running Linux, make sure that it uses GRUB version 1 and that the target virtualization platform is ESXi 5. For more details, see "Support for UEFI-based machines" (p. 33).

Logical and dynamic volumes

The resulting machine will have basic volumes, even if Linux logical volume structure is present in the backup. The same applies to dynamic volumes used in Windows systems. If you want to recreate logical or dynamic volumes on the machine, perform the conversion as described in "Recovery to a manually created virtual machine" (p. 126).

Custom loader reactivation

- During conversion, the disk interfaces may be changed as a result of migration to a different platform or just manually. The software sets the system-disk interface to be the same as the default interface for the new platform. The default interface is SCSI for VMware and IDE for other supported platforms. If the system disk interface changes, the name of the boot device also changes, while the boot loader still uses the old name.

- Conversion of logical volumes to basic ones may also prevent the system from booting up.

For these reasons, 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 normally not needed because Acronis Backup does this automatically. Should the need arise, use the procedure described in "How to reactivate GRUB and change its configuration" (p. 109).

For more considerations about physical to virtual machine conversion, see the "Backing up virtual machines" document.

6.2.2 Setting up regular conversion to a virtual machine

When creating a backup plan (p. 35), 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 of a backup from the following locations is not available: CD, DVD, Blu-Ray Discs, and Acronis Cloud Storage.
- Conversion to a Citrix XenServer virtual machine is not available as a part of the backup plan. As an alternative, use methods (b) and (c) as described in "Conversion methods" (p. 118).
- Microsoft Virtual PC does not support virtual disks larger than 127 GB. During a conversion to a Virtual PC machine, the size of every disk that exceeds 127 GB will be reduced to this value. If the disk resize is not possible, the conversion will fail. If you need larger virtual disks in order to connect them to a Hyper-V machine, use methods (b) and (c) as described in "Conversion methods" (p. 118).

6.2.2.1 Conversion settings

This section provides information that helps you make the appropriate conversion settings.

The settings are specified in the **Convert to virtual machine** section of the **Create backup plan** page.

Convert to virtual machine

Convert from

If you are copying or moving backups to other locations (p. 69), select the location where the backup will be taken from. Conversion locations which are not available (p. 120), such as Acronis Cloud Storage, are not listed.

By default, conversion will be performed from the primary location.

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. 121)** if required.

Target host... (p. 121)

Select the resulting virtual machine type and location. Available options depend on the agent that will perform conversion. This may be the agent that performs the backup (by default) or 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.

To specify another agent, click **Change** and select a machine where Agent for VMware, Agent for Hyper-V, Agent for Windows, or Agent for Linux is installed.

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.

6.2.2.2 Setting up a conversion schedule

A disk backup (p. 204) 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.

6.2.2.3 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 VMware** is installed on the ESX(i) host
You can only create a virtual machine on the ESX(i) host where the agent is running.
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 only create a virtual machine on the Hyper-V server.
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 intended to be backed up.
- **Agent for Windows** or **Agent for Linux** is installed on the machine

You can choose the virtual machine type: VMware Workstation, Microsoft Virtual PC, Red Hat Kernel-based Virtual Machine (KVM) or Red Hat Enterprise Virtualization (RHEV).

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

For VMware, Hyper-V and Virtual PC, 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.

For KVM or RHEV, disks of the resulting virtual machine will have the raw format. This means that virtual disk size is always equal to the original disk capacity. Assuming that the original disk size is 100 GB, the corresponding virtual disk will occupy 100 GB even if the disk stores 10 GB of data.

6.2.2.4 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, the old machine is temporarily renamed. Then, a new virtual machine is created that has the previous name of the old machine. If this operation succeeds, the old machine is deleted. If this operation fails, the new machine is deleted and the old machine is given its previous name. This way, the conversion always ends up with a single machine. However, extra storage space is required during conversion to store the old 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, 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.

6.2.3 Recovery to the "New virtual machine" destination

Rather than converting a TIB file to a virtual disk file, which requires additional operations to bring the virtual disk into use, Acronis Backup 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 Agent for Windows** or **Agent for Linux**, you can create a new virtual machine in a local or network folder. You can start the machine using the respective virtualization software or prepare the machine files for further usage. The following table summarizes the available virtual machine formats and the actions you can take to add the machine to a virtualization server.

VM format	Further action and tool to use	Target virtualization platform
VMware Workstation	Export using VMware Workstation; or Convert to OVF using VMware OVF tool > Deploy OVF template using vSphere Client	ESX(i)
Microsoft Virtual PC*	Add the VHD file to a Hyper-V machine	Hyper-V
Citrix XenServer OVA	Import using Citrix XenCenter	XenServer
Kernel-based Virtual Machine (Raw format)	Move the virtual machine files to a machine running Linux and run the virtual machine by using Virtual Machine Manager	-
Red Hat Enterprise Virtualization (RHEV) (Raw format)	Import using RHEV Manager	RHEV

*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.

With **Acronis Backup Agent for Hyper-V** or **Agent for VMware**, you can create a new virtual machine directly on the respective virtualization server.

6.2.3.1 Steps to perform

To perform a recovery to a new virtual machine

1. Connect the console to the management server, to a machine where an agent is installed, or to a machine booted from a bootable media.
2. Click **Recover** to open the **Recover data** (p. 96) page.
3. Click **Select data** (p. 97). 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. Click **Browse**. In the **VM/VS Selection** (p. 124) window, select the resulting virtual machine type or the virtualization server where to create the machine.
6. [Optional] In **Storage**, you can view or select the storage where the virtual machine will be created.
7. [Optional] In **Virtual machine settings** (p. 125), you can change the name of the virtual machine, the disk provisioning mode, the allocated memory, and other settings.

Machines of the same type and with the same name cannot be created in the same folder. If you get an error message caused by identical names, change either the VM name or the path.

8. The destination disk for each of the source disks or source volumes and the 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**, 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.

In the **Backup plans and tasks** view, you can examine the state and progress of the recovery task.

6.2.3.2 Virtual machine type / virtualization server selection

Select the resulting virtual machine type or the virtualization server where the machine will be created.

The available options depend on the agent(s) installed on the machine the console is connected to. If the console is connected to the management server, you can choose any registered machine that is able to perform the required operation.

To select the virtualization server where the new virtual machine will be created

1. Choose the **Create a new virtual machine on the server** 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.

[Only if the console is connected to the management server] If multiple agents manage the selected ESX(i) host, you can choose the agent that will perform recovery. If no agent manages the ESX(i) and automatic deployment is turned on, Agent for VMware (Virtual Appliance) will be

deployed immediately after you click **OK**. Recovery will be performed by that agent. It will take a license.

3. Click **OK** to return to the **Recover data** page.

To select the virtual machine type

1. Choose the **Save the virtual machine as a set of files** 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.

[Only if the console is connected to the management server] You can select the machine that will perform recovery. This can be any registered machine where Agent for Windows or Agent for Linux is installed.

3. Click **OK** to return to the **Recover data** page.

6.2.3.3 Virtual machine settings

The following virtual machine settings can be configured.

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 software 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.

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.

Name

Initial setting: if not contained in the backup, **New virtual machine**.

Enter the name for the new virtual machine. If the backup was created by Agent for VMware or Agent for Hyper-V, the software takes the name from the virtual machine configuration contained in the backup.

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.

6.3 Recovery to a manually created virtual machine

This section describes the conversion method (p. 118) in which you create a virtual machine yourself and perform a recovery to it as if it were a physical machine.

6.3.1 Considerations before conversion

Converting a UEFI-based machine

If the original machine uses Unified Extensible Firmware Interface (UEFI) for booting, consider creating a virtual machine that is also UEFI-based.

If your virtualization product does not support UEFI, you can create a BIOS-based machine, provided that the original machine is running Windows. Acronis Backup adjusts the Windows boot mode to the BIOS boot firmware and ensures that Windows remains bootable.

For Linux operating systems, changing the boot mode from UEFI to BIOS is not supported. Acronis Backup can convert a UEFI-based machine running Linux only when the machine uses GRUB version 1 and the target machine is also UEFI-based. For more details, see "Support for UEFI-based machines" (p. 33).

Choosing the disk interface

When creating the virtual machine, you may want its disks to have a different interface than those of the original machine.

- You may want to change all disk interfaces from IDE to SCSI when migrating a machine to ESX(i), because SCSI is a default disk interface for ESX(i) and it provides better performance.
- You need to change the system disk interface from SCSI to IDE when migrating a machine to Hyper-V, because Hyper-V does not support booting from SCSI disks.

If the original machine uses a custom boot loader, either recover the system disk to a disk with the same interface, or manually configure the boot loader. The reason is that when the interface of the system disk changes, the name of the boot device also changes; however, the boot loader still uses the old name. Configuring GRUB is normally not needed because Acronis Backup does this automatically.

6.3.2 Steps to perform

To perform a recovery to a manually created virtual machine

1. [Optional] When recovering Windows, prepare Windows drivers that correspond to the target virtualization platform.
For machines running Linux, the necessary drivers are normally already present in the operating system.
2. Create a bootable media (p. 151) by using Acronis Bootable Media Builder.
3. Create a virtual machine by using your virtualization product's native tools.
4. Boot the virtual machine from the bootable media.
5. [When recovering Windows] If you need dynamic volumes, create a volume group by using the disk management functionality.
6. Select **Actions > Recover**.
7. [When recovering Linux] If you need logical volumes, click **Apply RAID/LVM**. The LVM structure will be automatically recreated (p. 29) during the recovery.
8. Configure other recovery settings and perform a recovery in the same way as onto a physical machine.

The most recent operating systems remain bootable when recovered to dissimilar hardware, or the VMware or Hyper-V platforms. If a recovered operating system does not boot, use the Acronis Universal Restore tool to update the drivers and modules that are critical for the operating system startup. For details, see the Acronis Universal Restore page at <http://www.acronis.com/en-us/aur/>.

7 Storing the backed up data

7.1 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 or detachable media.

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 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. 130) section.

 **Personal vaults.** These vaults are 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. 129) and to take actions on archives (p. 148) and backups (p. 149) stored in there.

7.1.1 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. 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.

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. 99).

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. 148).
- Operations with backups (p. 149).
- Sorting, filtering and configuring table items (p. 17).

What does the icon mean?

When browsing archives on the **Archive view** tab, you may encounter a backup with the  icon. This icon means that the backup is marked for deletion but cannot be deleted immediately because other backups depend on it and consolidation is either not possible or disabled by retention rules.

You cannot perform any operation on backups marked for deletion. They disappear from the **Archive view** after they are physically deleted. This happens when all of the dependent backups are also deleted, or at next cleanup after you enable consolidation in the retention rules.

7.1.2 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 or removable media, Acronis Cloud Storage, 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.

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. 22).

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.

7.1.2.1 Actions on personal vaults

To access actions on personal vaults, click **Vaults > Personal** in the **Navigation** pane.

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 . The procedure of creating personal vaults is described in-depth in the Creating a personal vault (p. 131) section.
Edit a vault	1. Select the vault. 2. Click  Edit . 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 . In the appearing dialog box, provide the credentials required for accessing the vault.
Create Acronis Secure Zone	Click  Create Acronis Secure Zone . The procedure of creating the Acronis Secure Zone is described in-depth in the Creating Acronis Secure Zone (p. 132) section.
Explore a vault's content	Click  Explore . In the appearing Explorer window, examine the selected vault's content.
Validate a vault	Click  Validate . You will be taken to the Validation (p. 137) page, where this vault is already pre-selected as a source. The vault validation checks all the archives stored in the

To	Do
	vault.
Delete a vault	Click  Delete . 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 . 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 Cloud Storage, or on a hard drive local to the machine.
4. 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. 160).
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. 160).
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*.

7.2 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. 71).

Limitations

- Acronis Secure Zone cannot be organized on a dynamic disk.

7.2.1 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. 132)

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. 133)

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. 133)

[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. 133) window, review the expected layout and click OK to start creating the zone.

7.2.1.1 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.

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. 133) window.
3. Click **OK**.

7.2.1.2 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.

7.2.1.3 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**.

7.2.1.4 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.

7.2.2 Managing Acronis Secure Zone

Acronis Secure Zone is considered as a personal vault (p. 211). Once created on a managed machine, the zone is always present in the list of **Personal vaults**.

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. 148).

7.2.2.1 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**.

7.2.2.2 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 the free space after the zone is decreased.
If you select several volumes, the space will be distributed to each partition equally. If you do not select any volumes, the freed space becomes unallocated.
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**.

7.2.2.3 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 to each partition equally. If you do not select any volumes, the freed space becomes unallocated.

After you click **OK**, Acronis Backup will start deleting the zone.

7.3 Removable devices

This section describes peculiarities of backing up to removable devices.

By a removable device, we mean an RDX drive or USB flash drive. A USB hard disk drive is not considered to be a removable device unless it is recognized as such by the operating system.

In Linux, an RDX drive or USB flash drive is considered to be a removable device if it is specified by its name (for example, **sdf:/**). If a device is specified by its mount point (for example, **/mnt/backup**), it behaves as a fixed drive.

The method of working with removable disk libraries (multi-cartridge devices) depends on the device type, brand, and configuration. Therefore, each case should be considered individually.

Vaults on removable devices

Before backing up a machine to a removable device, you can create a personal vault (p. 131). If you do not want to, the software will automatically create a personal vault in the drive folder selected for backing up.

Limitation

- Vaults created on removable devices do not have the **Data view** (p. 99) tab.

Usage modes of removable devices

When creating a backup plan (p. 35), you can choose whether to use your removable device as a fixed drive or as removable media. The **Fixed drive** mode presumes that the removable device will always be attached to the machine. The **Removable media** mode is selected by default.

When you back up using the **Back up now** feature or under bootable media, the removable device is always used in the **Removable media** mode.

The difference between the two modes is mostly related to retention and replication of backups.

Functionality	Fixed drive	Removable media
If there is insufficient space to continue backing up, the software will prompt you to...	...manually free up disk space.	...insert new media.
You can set retention rules (p. 72) for backups stored on the device.	Yes	No
You can set the option to clean up the archive " When there is insufficient space while backing up " within the Custom (p. 47) backup scheme.	Yes	No
Simplified naming (p. 55) of backup files...	...is unavailable.	...is always used.
You can replicate backups (p. 71) <i>to</i> the removable device.	Yes	No
You can replicate backups <i>from</i> the removable device.	No	No
An archive with several full backups can be created.	Yes	No. Before creating a new full backup, the software will delete the entire archive and start a new one.
You can delete any backup of an archive.	Yes	No. You can delete only a backup that does not have dependent backups.

Since the removable device mode determines the naming scheme for backup files, the **Name backup files using the archive name...** check box does not appear when the backup destination is a removable device.

8 Operations with archives and backups

8.1 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 the 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.

Limitation

You cannot validate archives and backups in Acronis Cloud Storage (p. 179). However, an initial seeding backup (p. 182) is automatically validated immediately after its creation.

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. 35).

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. 142) - in this case, you need to specify the archive.

Backup (p. 138) - specify the archive first. Then, select the desired backup in this archive.

Vault (p. 138) - select a vault (or other location), to validate archives from.

Credentials (p. 139)

[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. 139)

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. 140)

[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.

8.1.1 Archive selection

To specify an archive to validate

1. Enter the full path to the archive location in the **Path** field, or select the required location in the tree (p. 98).
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**.

8.1.2 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**.

8.1.3 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 personal vault, expand the **Personal** group and click the appropriate vault.
 - To select a local folder, 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 folders** 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.

2. Click **OK**.

8.1.4 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.

8.1.5 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. 60) section.

8.1.6 Task credentials

Provide credentials for the account under which the task will run.

To specify credentials

1. Select one of the following:

- **Use current user credentials**

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, see the Owners and credentials (p. 22) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine (p. 24) section.

8.2 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.

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.

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. 142) or the **Backups** (p. 143).

Show access credentials (p. 143)

[Optional] Provide credentials for accessing the source if the task account does not have enough privileges to access it.

Where to export

Browse (p. 143)

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.

Full cataloging/Fast cataloging

Not available under bootable media or for locations that do not support cataloging

Select whether full or fast cataloging will be performed on the exported backups. For more information about the cataloging, see "Backup cataloging" (p. 80).

Show access credentials (p. 144)

[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.

8.2.1 Archive selection

To specify an archive to export

1. Enter the full path to the archive location in the **Path** field, or select the required location in the tree (p. 98).

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**.

8.2.2 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**.

8.2.3 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.

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.

8.2.4 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 personal vault, expand the **Personal** 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.

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.

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.

8.2.5 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.

8.3 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. In this mode, the software creates an incremental backup containing the changes you make to the backup content. Please be aware that none of the subsequent backups will contain these changes.

You can mount volumes if the disk backup is stored in a local folder (except optical disks), 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. 146)

Specify the path to the archive location and select the archive containing disk backups.

Backup (p. 146)

Select the backup.

Access credentials (p. 146)

[Optional] Provide credentials for the archive location.

Mount settings

Volumes (p. 147)

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.

8.3.1 Archive selection

To select an 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 a personal vault located in a local folder, Acronis Secure Zone, or on a network share, expand the **Personal** group and click the required vault.
- If the archive is stored in a local folder on the machine, expand the **Local folders** group and click the required folder.

Mounting is not available if the archive is stored on optical media such as CD, DVD, or Blu-ray Discs (BD).

- If the archive is stored on 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 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 NFS share, expand the **NFS folders** group and click the 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**.

8.3.2 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**.

8.3.3 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.

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.

8.3.4 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 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**.

8.3.5 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**.

8.4 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. 148)
 - Operations with backups (p. 149)

8.4.1 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	<p>Click  Validate.</p> <p>The Validation (p. 137) page will be opened with the pre-selected archive as a source.</p> <p>Validation of an archive will check all the archive's backups.</p>
Export an archive	<p>Click  Export.</p> <p>The Export (p. 140) 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.</p>
Delete a single archive or multiple archives	<ol style="list-style-type: none"> 1. Select one or more archives you want to delete. 2. Click  Delete. <p>The program duplicates your selection in the Backups deletion (p. 150) 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.</p>
Delete all archives in the vault	<p>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.</p> <p>Click  Delete all.</p> <p>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.</p>

8.4.2 Operations with backups

To perform any operation with a backup

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  View content . In the Backup Content window, examine the backup content.
Recover	Click  Recover . The Recover data (p. 96) page will be opened with the pre-selected backup as a source.
Validate a backup	Click  Validate . The Validation (p. 137) 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  Export . The Export (p. 140) 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  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. 149) for more information.
Delete a single or multiple backups	Select one of the backups you want to delete, then click  Delete . The program duplicates your selection in the Backups deletion (p. 150) 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. Click  Delete all . The program duplicates your selection in the Backups deletion (p. 150) window that has check boxes for each archive and each backup. Review the selection and correct if need be, then confirm the deletion.

8.4.3 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.

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**.

Tips on usage

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. 140) operation.

When you mount an image (p. 145) in the read/write mode, the software creates an incremental backup containing the changes you make to the backup content. The subsequent backups do not contain these changes. Naturally, if you convert any of the subsequent backups to full, none of these changes will appear in the resulting full backup.

Limitations

Conversion is not allowed for the following backups:

- Backups stored on CD/DVD or in Acronis Cloud Storage.
- Backups that have simplified names (p. 55).
- Backups of Microsoft Exchange Server data.

8.4.4 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.

9 Bootable media

Bootable media

Bootable media is physical media (CD, DVD, USB flash drive or other removable media supported by a machine BIOS as a boot device) that boots on any PC-compatible machine and enables you to run Acronis Backup Agent in a bootable environment based on Linux kernel, without the help of an operating system. The agent can boot and perform operations on any PC-compatible hardware, including bare metal and machines with corrupted or non-supported file systems.

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. This server 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 using the same wizard.

A list of hardware supported by bootable media is available in the following Acronis Knowledge Base article: <https://kb.acronis.com/content/55310>.

9.1 Linux-based bootable media

To create a Linux-based bootable media

1. Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media**, or as a separate component.
2. Select **Bootable media type: Default (Linux-based media)**.
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.
3. Follow the wizard steps to specify the following:
 - a. [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. 152).
 - b. The Acronis bootable components to be placed on the media.
You can select 32-bit and/or 64-bit components. The 32-bit components can work on 64-bit hardware. However, you need 64-bit components to boot a machine that uses Unified Extensible Firmware Interface (UEFI).

To use the media on different types of hardware, select both types of components. When booting a machine from the resulting media, you will be able to select 32-bit or 64-bit components on the boot menu.

- c. [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.
- d. [Optional] Remote logon settings:
 - User name and password to be entered on the console side at the connection to the agent. If you leave these boxes empty, the connection will be enabled without specifying credentials.
- e. [Optional] Network settings (p. 154):
 - TCP/IP settings to be assigned to the machine network adapters.
- f. [Optional] Network port (p. 155):
 - The TCP port that the bootable agent listens for incoming connection.
- g. Licenses to be used:

If Agent for Windows or Agent for Linux is *not* installed on the machine, specify the license key or the license server with licenses. The licenses will not get assigned or reassigned. They determine which functionality to enable for the created media. Without a license, you can create media only for recovery from the cloud storage.

If Agent for Windows or Agent for Linux *is* installed on the machine, the media inherits its functionality.
- h. 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.
- i. Path to the media ISO file or the name or IP and credentials for PXE or WDS/RIS.

9.1.1 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 can also 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.

9.1.2 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.

9.1.3 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 Management Console.

9.2 Preparing to work under bootable 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. You can change the network settings before or after starting the work.

To start working, click **Manage this machine locally**.

Configuring network settings

To change the network settings for a current session, click **Configure network** in the startup window. The **Network Settings** window that appears will allow you to configure network settings for each network interface card (NIC) of the machine.

Changes made during a session will be lost after the machine reboots.

Adding VLANs

In the **Network Settings** window, you can add virtual local area networks (VLANs). Use this functionality if you need access to a backup location that is included in a specific VLAN.

VLANs are mainly used to divide a local area network into segments. A NIC that is connected to an *access* port of the switch always has access to the VLAN specified in the port configuration. A NIC connected to a *trunk* port of the switch can access the VLANs allowed in the port configuration only if you specify the VLANs in the network settings.

To enable access to a VLAN via a trunk port

1. Click **Add VLAN**.
2. Select the NIC that provides access to the local area network that includes the required VLAN.
3. Specify the VLAN identifier.

After you click **OK**, a new entry appears in the list of network adapters.

If you need to remove a VLAN, click the required VLAN entry, and then click **Remove VLAN**.

9.3 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. Under a Windows-style bootable media, a volume has the same drive letter as in Windows. Volumes that do not have drive letters in Windows (such as the **System Reserved** volume) are assigned free letters in order of their sequence on the disk.

If the bootable media cannot detect Windows on the machine or detects more than one of them, all volumes, including those without drive letters, are assigned letters in order of their sequence on the disk. This way, the volume letters may differ from those seen in Windows. For example, the D: drive under the bootable media 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. 55). 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.

9.3.1 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. 151) section for details).

9.3.2 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**.

9.4 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

<code>chroot</code>	<code>kpartx</code>	<code>sed</code>
<code>cp</code>	<code>ln</code>	<code>sg_map26</code>
<code>dd</code>	<code>ls</code>	<code>sh</code>
<code>df</code>	<code>lspci</code>	<code>sleep</code>
<code>dmesg</code>	<code>lvm</code>	<code>ssh</code>
<code>dmraid</code>	<code>mdadm</code>	<code>sshd</code>
<code>e2fsck</code>	<code>mkdir</code>	<code>strace</code>
<code>e2label</code>	<code>mke2fs</code>	<code>swapoff</code>
<code>echo</code>	<code>mknod</code>	<code>swapon</code>
<code>egrep</code>	<code>mkswap</code>	<code>sysinfo</code>
<code>fdisk</code>	<code>more</code>	<code>tar</code>
<code>fsck</code>	<code>mount</code>	<code>tune2fs</code>
<code>fxload</code>	<code>mtx</code>	<code>udev</code>
<code>gawk</code>	<code>mv</code>	<code>udevinfo</code>
<code>gpm</code>	<code>pccardctl</code>	<code>udevstart</code>
<code>grep</code>	<code>ping</code>	<code>umount</code>
<code>growisofs</code>	<code>pktsetup</code>	<code>uuidgen</code>
<code>grub</code>	<code>poweroff</code>	<code>vconfig</code>
<code>gunzip</code>	<code>ps</code>	<code>vi</code>
<code>halt</code>	<code>raidautorun</code>	<code>zcat</code>
<code>hexdump</code>	<code>readcd</code>	
<code>hotplug</code>	<code>reboot</code>	

9.5 Acronis Startup Recovery Manager

Acronis Startup Recovery Manager is a modification of the bootable agent (p. 201), 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).

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. This section also covers supplementary operations that can be performed on a managed machine, such as changing a license, adjusting **Machine options**, and collecting system information.

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. 162). A backup plan execution state is a cumulative state of the plan's most recent activities. The status of a backup plan (p. 163) helps you to estimate whether the data is successfully protected.

To keep track of a task's current progress, examine its state (p. 163). Check a task status (p. 164) 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. 17).
- 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. 160).
- 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. 169) and **Task details** (p. 170) 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.

To	Do
Create a new backup plan or task	Click  New , then select one of the following: <ul style="list-style-type: none">▪ Backup plan (p. 35)▪ Recovery task (p. 96)▪ Validation task (p. 137)

To	Do
View details of a plan/task	Click  Details . In the respective Plan Details (p. 169) or Task Details (p. 170) window, review the plan or task details.
View plan's/task's log	Click  Log . You will be taken to the Log (p. 170) view containing the list of the log entries grouped by the plan/task-related activities.
Run a plan/task	<p><u>Backup plan</u></p> <ol style="list-style-type: none"> 1. Click  Run. 2. In the drop-down list, select the plan's task you need run. <p>Running the backup plan starts the selected task of that plan immediately in spite of its schedule and conditions.</p> <p><u>Task</u></p> <p>Click  Run.</p> <p>The task will be executed immediately in spite of its schedule and conditions.</p>
Stop a plan/task	<p>Click  Stop.</p> <p><u>Backup plan</u></p> <p>Stopping the running backup plan stops all its tasks. Thus, all the task operations will be aborted.</p> <p><u>Task</u></p> <p>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.</p> <p>What will happen if I stop the recovery task?</p> <ul style="list-style-type: none"> ▪ 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. ▪ 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. ▪ 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	<p>Click  Edit.</p> <p>Backup plan editing is performed in the same way as creation (p. 35), except for the following limitations:</p> <p>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).</p> <ol style="list-style-type: none"> 1. It is not possible to change the scheme to Grandfather-Father-Son or Tower of Hanoi. 2. If the Tower of Hanoi scheme is used, it is not possible to change the number of levels. <p>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.</p>
Clone a backup plan	<p>Click  Clone.</p> <p>The clone of the original backup plan will be created with default name "<i>Clone of <original_plan_name></i>". 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.</p>
Enable a plan	<p>Click  Enable.</p> <p>The previously disabled backup plan will run again as scheduled.</p>
Disable a plan	<p>Click  Disable.</p> <p>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.</p>
Export a plan	<p>Click  Export.</p> <p>Specify the path and name of the resulting file. See Export and import of backup plans (p. 164) for more information.</p>
Import a plan	<p>Click  Import.</p> <p>Specify the path and name of the file that contains a previously exported plan. See Export and import of backup plans (p. 164) for more information.</p>
Delete a plan/task	<p>Click  Delete.</p>

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	<p>At least one task needs user interaction.</p> <p>Otherwise, see 2.</p>	<p>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).</p>

	State	How it is determined	How to handle
2	Running	At least one task is running. Otherwise, see 3.	No action is required.
3	Waiting	At least one task is waiting. 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.) 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. 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. 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: <ul style="list-style-type: none"> ▪ Remove the reason of the failure -> [optionally] Start the failed task manually ▪ Edit the local plan to prevent its future failure if a local plan has failed
2	Warning	At least one task has succeeded with warnings. 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: <ul style="list-style-type: none">▪ Remove the reason of the failure -> [optionally] Start the failed task manually▪ 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.

You can edit plans in the Acronis Backup 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. 165) 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 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. 167) to the other machines.

Adjusting credentials

Before exporting a backup plan that will further be imported to a different machine, check the user account under which the plan runs (**Edit > Plan parameters > Show task credentials, comments, label > Plan's credentials**).

The plan will successfully run on a different machine if the **Plan's credentials** value is either **Acronis service credentials** or **Run as: ... (current user)**. If the **Plan's credentials** parameter contains a specific user account, the plan will start only if there is an identical account on that machine. Therefore, you may need to do one of the following:

- Create an account with identical credentials on the machine where the plan will be imported.
- Edit credentials in the export file before importing. For details, see *Editing the export file* (p. 165).
- Edit credentials after importing the plan.

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. Confirm your choice.
4. If you need to edit the newly imported backup plan, select it in the **Backup plans and tasks** view, then click  **Edit**. Make the necessary changes and click **Save**.

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 Management Console or Acronis Backup command-line utility (**acrocml**) is installed. The path to the utility is as follows:

- In a 32-bit version of Windows: **%CommonProgramFiles%\Acronis\Utils**
- In a 64-bit version of Windows: **%CommonProgramFiles(x86)%\Acronis\Utils**
- In Linux: **/usr/sbin**

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 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. 165) 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. 164).
3. [Optional] Edit the export file. See Editing the export file (p. 165) 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** (in Windows Vista and later versions of Windows) or **%ALLUSERSPROFILE%\Application Data\Acronis\BackupAndRecovery\import** (in versions of Windows earlier than Windows Vista).

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 **Acronis Managed Machine Service**.

In Linux

The default path to the dedicated folder is **/usr/lib/Acronis/BackupAndRecovery/import**.

The path is stored in the file **/etc/Acronis/MMS.config**.

To change the path, edit the **/usr/lib/Acronis/BackupAndRecovery/import** value in the following tag:

```
<key name="Import">
  <value name="FolderPath" type="TString">
    "/usr/lib/Acronis/BackupAndRecovery/import"
  </value>
</key>
```

The change will be applied after a restart of the agent. To restart the agent, run the following command as the root user:

```
/etc/init.d/acronis_mms restart
```

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. 162) of the backup plan.
- **Status** - status (p. 163) 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 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. 17).

Select the activity or log entry to take an action on log entries. For details, see "Actions on log entries" (p. 170) and "Log entry details" (p. 171).

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. The Information pane will show log entries for the selected activity.
Select a single log entry	Click on it.

To	Do
Select multiple log entries	<ul style="list-style-type: none"> ▪ <i>non-contiguous</i>: hold down CTRL and click the log entries one by one ▪ <i>contiguous</i>: 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	<ol style="list-style-type: none"> 1. Select a log entry. 2. Do one of the following: <ul style="list-style-type: none"> ▪ Double click the selection. ▪ Click  Details. <p>The log entry's details will be displayed. See Log entry details for details of the log entry's operations.</p>
Save the selected log entries to a file	<ol style="list-style-type: none"> 1. Display Activities and select activities or display Events and select log entries. 2. Click  Save selected to file. 3. In the opened window, specify a path and a name for the file. <p>All log entries of the selected activities or selected log entries will be saved to the specified file.</p>
Save all the log entries to a file	<ol style="list-style-type: none"> 1. Make sure, that the filters are not set. 2. Click  Save all to file. 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	<ol style="list-style-type: none"> 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. <p>All log entries in the list will be saved to the specified file.</p>
Delete all the log entries	<p>Click  Delete all.</p> <p>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.</p>

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 the event has occurred. It is an integer number that may be used by Acronis support service to solve the problem.
- **Owner** - The user name of the backup plan owner (p. 22).
- **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 take responsibility 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 happened. 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. 19) - specify the alert types to display in the **Alerts** view.
- **Notifications** (p. 175) - set up e-mail notifications about alerts.
- **Settings** (p. 174) - 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 Changing a license

By changing the license, you switch a product from trial mode to full mode or switch to a different product. The following table summarizes the available options.

Switching a license	Why you may need it
Trial > Full	After trying the product, you decided to buy a license.
Full > Full, different product	<ul style="list-style-type: none">You want to upgrade from Acronis Backup to Acronis Backup Advanced in order to use the centralized management capability. For more information refer to the "Upgrading from Acronis Backup to Acronis Backup Advanced" section of the installation documentation.You used a server license (for example, Acronis Backup for Windows Server) for a workstation. Now you want to assign the workstation a workstation license (Acronis Backup for PC). After that, you can revoke the server license and use it for a server.
Backing up to the cloud storage* > Full	After backing up to the cloud storage only, you decided to buy a license to obtain greater functionality.
Trial > Backing up to the cloud storage*	After trying the product, you decided to back up to the cloud storage only.

*Prior to backing up to the cloud storage, you need to activate a subscription for the cloud backup service on the machine(s) you want to back up. For more information refer to the "Cloud backup" (p. 179) section.

To change a license

1. Click **Help > Change license**.
2. Click **Change** or **Specify** near your current license, click **Change**, and then click **Use the following license keys**.
3. Enter the new license key.

Managing a cloud backup subscription

The **Acronis Cloud** block of the **Licenses** window requires that you sign in to your Acronis account. After that, it shows the cloud backup subscription activated on the machine. If no subscription is activated, this block allows you to request for a subscription, to enter the registration code that you received after the subscription purchase, and to activate the subscription.

10.5 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 or Agent for Linux 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.

10.6 Adjusting machine options

The machine options define the general behavior of all Acronis Backup 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.

10.6.1 Alerts

10.6.1.1 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.

10.6.1.2 Time-based alerts

Last backup

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.

10.6.2 E-mail settings

The option enables you to configure e-mail settings to send notifications about alerts which occurred on the managed machine.

The notification schedule and the types of alerts to send are configured in **Machine options > E-mail settings > Alert notifications** (p. 175).

The preset is: **Disabled**.

Note: Alerts warn only about problems. Therefore, e-mail notifications about successful backup or recovery operations will not be sent. These e-mail notifications are configured in *Backup options > Notifications > E-mail* (p. 84) and in *Recovery options > Notifications > E-mail* (p. 113) respectively.

To configure e-mail notification

1. In the **E-mail addresses** field, type the destination e-mail address. You can enter several addresses separated by semicolons.
2. In the **Subject** field, type the notification subject or leave the default value. Variables are not supported in this field.
3. In the **SMTP server** field, enter the name of the outgoing mail server (SMTP).
4. In the **Port** field, set the port of the outgoing mail server. By default, the port is set to **25**.
5. If the outgoing mail server requires authentication, enter **User name** and **Password** of the sender's e-mail account.

If the SMTP server does not require authentication, leave the **User name** and **Password** fields blank. If you are not sure whether the SMTP server requires authentication, contact your network administrator or your e-mail service provider for assistance.

6. Click **Additional e-mail parameters...** to configure additional e-mail parameters as follows:
 - a. **From** – type the name of the sender. If you leave this field empty, the messages will contain the sender's e-mail account in the **From** field.
 - 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** and **Password** of the incoming mail server.
 - d. Click **OK**.
7. Click **Send test e-mail message** to check whether e-mail notifications work correctly with the specified settings.

10.6.2.1 Alert notifications

This option enables you to specify when to send e-mail notifications about alerts which occurred on the managed machine and to select the types of alerts to send.

When using this option, make sure that the e-mail settings are properly configured in **Machine options > E-mail settings** (p. 174).

The preset is: **Disabled**.

To configure alert notifications

1. Select when to send alert notifications:
 - **As soon as an alert appears** – to send a notification every time a new alert occurs.
Click **Select the types of alerts...** to specify the types of alerts to send notifications about.
 - **On schedule send notification about all current alerts** – to send a cumulative alert notification including all alerts which occurred over a time interval you specify.
Click **Select the types of alerts...** to specify the types of alerts to send notifications about.
Set up the notification **Frequency** and **Time**.
2. Click **OK**.

10.6.3 Event tracing

It is possible to send log events generated by the agent(s), operating on the managed machine, 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.

10.6.3.1 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 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, please see "Support for SNMP (p. 33)".

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. 177).

10.6.3.2 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 Immib2.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.

10.6.4 Log cleanup rules

This option specifies how to clean up the Acronis Backup agent log.

This option defines the maximum size of the agent log file (`/var/lib/Acronis/BackupAndRecovery/MMS/events.db3`).

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.

10.6.5 Cloud backup proxy

This option is effective only for backup to and recovery from Acronis Cloud Storage over the Internet.

This option defines whether the Acronis agent will connect to the Internet through a proxy server.

Note *The proxy server must be configured to redirect both HTTP/HTTPS and TCP traffic.*

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 **Settings**, click **Show advanced settings**. Under **Network**, click **Change proxy settings**.

11 Cloud backup

This section provides details about using the Acronis Cloud Backup service. This service enables you to back up your data to Acronis Cloud Storage.

Acronis Cloud Backup might be unavailable in your region. To find more information, click here:
<http://www.acronis.com/en-us/my/cloud-backup/corporate>

To configure backup to the cloud storage or recovery from the storage, follow the regular steps described in the corresponding sections:

Creating a backup plan (p. 35)

Creating a centralized backup plan

Recovering data (p. 96)

The main difference is that you select the cloud storage as the backup destination.

11.1 Introduction to Acronis Cloud Backup

This section contains a brief overview of Acronis Cloud Backup and answers questions that may arise during evaluation and usage of this product.

11.1.1 What is Acronis Cloud Backup?

Acronis Cloud Backup is a service that enables you to back up data to Acronis Cloud 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 cloud backup service will be available to you.

Examples of subscriptions:

- A 1 TB volume subscription means that you can back up data from an unlimited number of physical and/or virtual machines, for a period of one year. The backups can occupy no more than one terabyte.
- A subscription for PC means that you can back up data from a machine running a non-server Windows operating system, for a period of one year. The storage quota is unlimited.

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 cloud backup solutions, Acronis Cloud Backup enables bare metal recovery directly from the cloud 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 cloud storage?

Your backups remain in the cloud storage until you delete them or until the subscription expires. Recovering data from the cloud 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 do I 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 Supported operating systems and virtualization products

Acronis Cloud Backup supports the following operating systems and virtualization platforms.

Server operating systems

Windows

- Windows Server 2003/2003 R2 – Standard and Enterprise editions (x86, x64)
- Windows Small Business Server 2003/2003 R2
- Windows Server 2008 – Standard, Enterprise, Datacenter, and Web editions (x86, x64)
- Windows Small Business Server 2008
- Windows Server 2008 R2 – Standard, Enterprise, Datacenter, Foundation, and Web editions
- Windows MultiPoint Server 2010/2011/2012
- Windows Small Business Server 2011 – all editions
- Windows Server 2012/2012 R2 – all editions
- Windows Storage Server 2003/2008/2008 R2/2012/2012 R2

Linux

- Linux with kernel from 2.4.20 to 4.1 and glibc 2.3.2 or later
- Various x86 and x86_64 Linux distributions, including:
 - Red Hat Enterprise Linux 4.x, 5.x, 6.x, and 7.x
 - Ubuntu 9.10, 10.04, 10.10, 11.04, 11.10, 12.04, 12.10, 13.04, 13.10, 14.04, 14.10, and 15.04
 - Fedora 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, and 22
 - SUSE Linux Enterprise Server 10 and 11
 - SUSE Linux Enterprise Server 12 – supported on file systems, except for Btrfs
 - Debian 4, 5, 6, 7.0, 7.2, 7.4, 7.5, 7.6, 7.7, 8.0, and 8.1
 - CentOS 5.x, 6.x, and 7.0
 - Oracle Linux 5.x, 6.x, 7.0, and 7.1 – both Unbreakable Enterprise Kernel and Red Hat Compatible Kernel
 - CloudLinux 6.x

Before installing the product on a system that does not use RPM Package Manager, such as an Ubuntu system, you need to install this manager manually; for example, by running the following command (as the root user): **apt-get install rpm**

Workstation operating systems

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)

Windows 8/8.1 – all editions except for the Windows RT editions (x86, x64)

Virtualization products (host-based backup of virtual machines)

VMware ESX(i) 5.0, 5.1, 5.5, and 6.0

(Host-based backup is available only for paid licenses of VMware ESXi.)

Windows Server 2008 (x64) with Hyper-V

Windows Server 2008 R2 with Hyper-V

Microsoft Hyper-V Server 2008/2008 R2

Windows Server 2012/2012 R2 with Hyper-V

Microsoft Hyper-V Server 2012/2012 R2

Windows 8, 8.1 (x64) with Hyper-V

Windows 10 – Pro, Education, and Enterprise editions with Hyper-V

11.1.6 Backup and recovery FAQ

This section answers questions related to backup and recovery processes.

11.1.6.1 What backup methods are available?

Full and incremental backup methods are available through several backup schemes. Regardless of the backup scheme, the first task run produces a full backup; subsequent task runs produce incremental backups. The following backup schemes are available:

- **Manual start** (postponed start). You can run the task again manually.
- **Simple** (start on schedule). With this backup scheme, you can set up a retention rule to automatically delete old backups.
- **GFS (Grandfather-Father-Son)** (start on schedule). You specify which of the daily backups to consider as weekly backups and monthly backups. You can set up separate retention rules for daily, weekly, and monthly backups.
- **Tower of Hanoi** (start on schedule). You set up the number of levels. This is the number of backups stored at a time. The excessive backups will be deleted in a manner that leaves more recovery points for recent dates and fewer recovery points for older dates.
- An additional backup scheme that is available only for cloud backups is **Initial seeding**. With this scheme, the backup starts immediately to a local destination and using the full backup method. To use this scheme, you need a license for the Initial Seeding (p. 52) service.

11.1.6.2 What recovery methods are available?

There are two methods to recover your data from Acronis Cloud Storage:

- Recovering disks or files by using Acronis Backup GUI or command line interface. This method enables you to use a wide range of Acronis Backup functionality.

- Retrieving files (p. 195) from file-level backups by using a web browser. To do this, you only need a machine with Internet access.

11.1.6.3 Is the cloud storage available under Acronis bootable media?

Recovery from Acronis Cloud Storage is available but backup to the storage is not.

11.1.6.4 What if a network connection is lost during cloud backup or recovery?

The software will try to reach the cloud storage every 30 seconds. The attempts will be stopped as soon as the connection is resumed OR a certain number of attempts are performed, depending on which comes first. The default number of attempts is 300 when backing up and 30 when recovering.

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.

11.1.6.5 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 webpage near the machine. This means you have to free some space for future backups. Or, you can consider increasing the storage quota. You may also want to set or edit the retention rule (p. 179) so that an overflow does not occur in future. Once the occupied space reaches the limit, the backups will cease to run.

11.1.6.6 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 cloud storage delete them. Since the deletion is performed on the cloud storage side, it does not take your machine's CPU resource.

The cleanup task runs after every cloud 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 cloud storage?" (p. 179)

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.

11.1.6.7 How do I 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. 193) 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.

11.1.7 Initial Seeding FAQ

This section explains what Initial Seeding is, why you would want to use it and provides some usage details.

11.1.7.1 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 cloud 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. 188) option.

11.1.7.2 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 cloud storage.

11.1.7.3 Is Initial Seeding a paid service?

Yes, you need to buy one Initial Seeding license per machine.

11.1.7.4 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.

11.1.7.5 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.

11.1.7.6 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.

11.1.7.7 How to buy an Initial Seeding license?

You can buy an Initial Seeding license from an Acronis partner or from the Acronis online store.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Log in to your Acronis account and enter the registration code in the product registration section. The registered license appears on the **Initial Seeding / Recovery** tab on your account management webpage.

A license purchased from the Acronis online store appears on the **Initial Seeding / Recovery** tab immediately after the payment is processed.

11.1.7.8 How do I perform initial seeding?

Preparing

1. Ensure that you have activated an Acronis Cloud Backup subscription on the machine where you will do initial seeding (skip this step if you have a volume subscription).
2. If you are currently using a trial subscription, ensure 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.
3. Decide on the media (p. 183) that you will send.
4. Choose the shipping company that you will use to ship your media. We strongly recommend that you use a recognized provider, such as UPS, FedEx, or DHL.

Creating the initial full backup

1. 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.
2. Start Acronis Backup, 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 **Cloud 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**.

3. [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.

Packaging and sending

1. Package the media along with a prepaid return shipping label.

Important Carefully follow the instructions in the "How to package a hard drive for shipment?" (p. 184) section, to ensure that your hard drive is safely delivered and returned.

2. Send the media to Acronis by physical mail.
3. On your account management webpage, mark the order as "shipped" and track (p. 187) the order status.

Creating subsequent backups

Once you observe that the backup has been uploaded to the cloud storage, you can edit the backup plan to do incremental backups:

- In **Backup scheme**, select the desired backup scheme and specify its settings.
- Click **Save**.

When started manually or on a schedule, your backup plan will add incremental backups to the initial backup stored in the cloud storage.

11.1.7.9 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

Using the website of the shipping company that you chose, 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.



11.1.7.10 How do I track an Initial Seeding order status?

On the Acronis website, 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. 184) 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 Cloud Storage has started.
- **The data upload has been completed** – The initial full backup has been successfully uploaded to the cloud storage. You can do incremental cloud 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.

11.1.8 Large Scale Recovery FAQ

This section explains what Large Scale Recovery is, why you would want to use it and provides some usage details.

11.1.8.1 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 cloud 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.

11.1.8.2 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.

11.1.8.3 Do I need to perform initial seeding to be able to use Large Scale Recovery?

No, these services are independent.

11.1.8.4 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.

11.1.8.5 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.

11.1.8.6 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.

11.1.8.7 How to buy a Large Scale Recovery license?

You can buy a Large Scale Recovery license from an Acronis partner or from the Acronis online store.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Log in to your Acronis account and enter the registration code in the product registration section. The registered license appears on the **Initial Seeding / Recovery** tab on your account management webpage.

A license purchased from the Acronis online store appears on the **Initial Seeding / Recovery** tab immediately after the payment is processed.

11.1.8.8 How do I track a Large Scale Recovery order status?

On the Acronis website, 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 website.

11.1.8.9 How to perform large scale recovery?

The recovery procedure is the same as when recovering from the cloud storage. Just specify the path to the location where your backups are. For detailed information about recovery please refer to the context-sensitive help.

11.1.9 Subscription lifecycle FAQ

This section explains a subscription lifecycle and subscription operations that you can perform on your account management webpage.

11.1.9.1 How do I access my account management webpage?

Go to <http://www.acronis.com/en-us/my/cloud-backup/corporate> and log in to your account (create one if you are not registered).

To access this webpage from Acronis Backup:

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).

11.1.9.2 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. Go to your account management webpage, click **Enter new registration code**, and then enter the registration codes. The subscriptions will appear in the list of available subscriptions on the **Manage subscriptions** tab.

Another way to register the subscriptions is to enter the registration codes during local installation of Acronis Backup in Windows.

Subscriptions that are purchased from the Acronis website are available on this tab immediately.

11.1.9.3 When does my subscription begin?

The subscription period of a **volume** subscription begins at the time of purchase.

The subscription period of a **machine** subscription begins as soon as the subscription is activated on a machine.

11.1.9.4 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 webpage near the machine. This means you need to renew (p. 190) 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 cloud storage for five days following the expiration date. You will be able to recover data from the cloud storage for 30 days following the expiration date.

11.1.9.5 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

You can renew an activated subscription to a subscription with the same or a larger storage quota.

To renew a **volume** subscription, you need a volume subscription. To renew a **machine** subscription, you need a machine subscription of the same type or a volume subscription.

A subscription for virtual machines (now deprecated) can be renewed to a server subscription or to a volume subscription.

Volume subscriptions

To renew a **volume** subscription, go to the account management webpage, click **Renew** next to the volume subscription, and then follow the on-screen instructions.

The new expiration date will appear in the **Expires** column.

- If the new subscription has the same storage quota as the old one, the subscription periods will be added together.
- If the new subscription has a larger storage quota, the resulting subscription period will be recalculated.

Machine subscriptions

To renew a **machine** subscription, go to the account management webpage, find the machine that you want to renew the subscription for, click **Renew** next to the machine, and then follow the on-screen instructions.

The new subscription will appear in the **Next subscription** column for the 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 webpage. 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.

11.1.9.6 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.

11.1.9.7 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. 193) it to a different machine in Acronis Backup GUI.

11.1.9.8 Can I cancel my subscription?

Just wait until the subscription expires. Refunds are not available for the cloud backup subscriptions.

11.2 Where do I start?

Go to <http://www.acronis.com/en-us/my/cloud-backup/corporate> and log in to your account (create one if you are not registered). This is your *account management webpage*. Here you can get a trial subscription, locate an Acronis partner or buy subscriptions online. The newly obtained subscriptions are listed as available subscriptions on 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 start backing up to Acronis Cloud Storage.

11.3 Choosing a subscription

Volume subscriptions

A **volume** subscription enables you to back up an unlimited number of machines running any supported operating system (p. 180). All backed up machines share a common storage quota. The subscription period begins at the time of purchase.

Machine subscriptions

A **machine** subscription enables you to back up a single machine. The storage quota applies only to this machine. The subscription period begins when the subscription is activated on the machine.

Choose the subscription for **server** or for **PC** based on the operating system that the machine is running. If you doubt whether a machine is a server or a workstation, refer to the list of the supported operating systems (p. 180).

If your backups are likely to exceed the storage quota for the subscription, you may want to use a server subscription on a workstation. The inverse usage is not possible. You cannot back up a server by using a subscription for PC.

Trial subscriptions

You can get one free subscription per account. The trial subscription enables you to back up a single machine. The subscription period is limited to one month.

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 a subscription, and then renew the trial subscription specifying the paid subscription. Your backed up data will be kept. 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:

- On the account management webpage, click the cloud backup trial link, and then select the necessary subscription type.
- Install Acronis Backup on the machine you want to back up, start the product, click **Back up now** or **Create backup plan**, click **Location**, and then 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 cloud backup subscriptions

Activating a subscription on a machine means allowing the machine to back up to the cloud storage.

A **volume** subscription is activated automatically as soon as you start backing up the machine to Acronis Cloud Storage.

A **machine** subscription (for Server or for PC) is also activated automatically if all subscriptions available in your account are the same type and have the same storage quota. If various subscriptions are available in your account, choose the one to be activated either when creating a backup plan or by manual activation (described in this section). The subscription period of a machine subscription starts at the moment of activation.

Important Before activating the first subscription for your account, check the country selected in your profile. Depending on this country, the service determines the data center where your backups will be sent. Make sure to select the country where all or most of the machines you want to back up to the cloud storage are located. Otherwise, the data may unnecessarily travel a long way. Later, you will not be able to change the data center even if you change the country in your profile. To access your profile, go to the Acronis website, log in to your account, and then click **Personal Profile**.

11.4.1 Activating subscriptions in Acronis Backup

To activate a subscription

1. Start Acronis Backup.
2. On the **Actions** menu, click **Activate cloud backup subscription**.
3. Specify the credentials to log in to the cloud storage.
4. Select the subscription that you want to activate for the machine.
5. Click **Activate now**.

11.4.2 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 on a machine and want to resume its cloud 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 cloud 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. Click **Reassign an already used subscription**.
3. Select the machine whose subscription you want to reassign to the current machine.
4. Click **Reassign 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 to use the proxy server.

1. Start Acronis Backup.

2. On the **Options** menu, click **Machine options**.
3. Click **Cloud backup proxy**.
4. Enter the proxy server settings. For detailed information (p. 177) about the settings please refer to the context-sensitive help.

11.6 Retrieving files from the cloud storage by using a web browser

By using a web browser, you can browse Acronis Cloud Storage, view contents of file-level archives, and download the selected files and folders.

The following browsers support these operations:

- Internet Explorer 7 or later
- Mozilla Firefox 3.5 or later
- Google Chrome 10 or later
- Safari 5.0.5 or later

To retrieve files from the cloud storage:

1. Go to the account management webpage (p. 190) and click **Recover files from Acronis Cloud**. You will see the list of the machines backed up by using the specified account. The list of machines that share a volume subscription appears when you select this subscription.
2. Click the name of the machine whose data you want to retrieve. The software displays both file-level and disk-level archives of this machine's data.

Note for users of the Initial Seeding (p. 52) service. While an initial seeding backup is being uploaded from your hard drive to Acronis Cloud Storage, the backup is visible but its data is not retrievable.

3. Click the required file-level archive. If prompted, enter the archive password. The software displays all of files and folders that were ever backed up to this archive.
4. If necessary, browse to the required folder or use search to obtain the list of the required files and folders.

Details. The search string can contain one or more wildcard characters * and ?.

5. Do one of the following:
 - To get the latest version of a single file or folder, simply click its name.
 - To get the latest versions of multiple files and folders, select the check boxes to the left of them and click the **Download** button.
 - To get an earlier version of a single file, click the  icon to the right of it and choose **View versions**. This opens a window with the version list. In this window, select the required version by its date and time and click **Download**.
 - [Not available if you used search] To get earlier versions of multiple files and folders, select the required date and time from the **Versions** list. Select the check boxes to the left of the files and folders and click the **Download** button.

Details. You will get the file and folder versions that were backed up prior and closest to the selected point in time.

6. To download the selected files, click **Save**.

Details. If you selected a single file, it is downloaded as is. Otherwise, the selected data will be archived into a .zip file (named AcronisArchive.zip by default).

7. Select the location to save the data to, and then click **Save**.

11.7 Limitations of the cloud storage

Unlike other types of storage available in Acronis Backup, the cloud storage has the following limitations.

Operations

The following operations are not possible.

Backup operations:

- Backing up from bootable media
- Backing up with Agent for Exchange
- Creating differential backups
- Using the **Custom** backup scheme
- Simplified naming of backup files
- Simultaneous host-based backup of multiple virtual machines
- Setting up regular conversion of backups to a virtual machine

Operations with backups:

- Validating a backup*
- Exporting a backup
- Mounting a backup
- Replicating or moving backups *from* the cloud storage
- Converting an incremental backup to full

Operation with archives (an archive is a set of backups):

- Validating an archive
- Exporting an archive

These limitations also apply to backing up data using Initial Seeding and to recovering data using Large Scale Recovery.

* An initial seeding backup is automatically validated immediately after its creation.

Backup and recovery options

Some backup and recovery options are not supported by cloud 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.

11.8 Terminology reference

The following is the list of terms related to the Acronis Cloud Backup service.

Activate a subscription

Allow the machine to use the cloud storage according to the subscription.

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 in order to renew its current subscription.

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 cloud 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 cloud storage. After that, you can add incremental backups to this full backup, either manually or on a schedule.

The Initial Seeding service might be unavailable in your region. To find more information, click here: <http://kb.acronis.com/content/15118>.

Large Scale Recovery

An extra service that enables you to obtain a copy of the backups you have in the cloud storage. You can then recover data from this copy.

The Large Scale Recovery service might be unavailable in your region. To find more information, click here: <http://kb.acronis.com/content/15118>.

License

Not to be confused with Acronis Backup product license.

Permission for a machine to use an extra service of Acronis Cloud Backup.

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 webpage, and these subscriptions and licenses become available for use.

Renew a subscription

Assign a subscription of the same type and with the same or a larger 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 can be occupied according to the subscription.

Subscription

Permission for a machine or for multiple machines to use a specific amount of space in the cloud 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. 204) 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
- GPT disks and the UEFI boot mode are not supported.

Acronis Plug-in for WinPE

A modification of Acronis Backup Agent for Windows that can run in the preinstallation environment. The plug-in can be added to a WinPE (p. 212) image using Bootable Media Builder. The resulting bootable media (p. 201) can be used to boot any PC-compatible machine and perform, with certain limitations, most of the direct management (p. 204) 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. 203).

Acronis Secure Zone

A secure volume for storing backup archives (p. 200) within a managed machine (p. 208).

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. 205).

Acronis Secure Zone is considered as a personal vault (p. 209).

Acronis Startup Recovery Manager (ASRM)

A modification of the bootable agent (p. 201), 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.

Activity

An action performed by Acronis Backup 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. 211) always causes one or more activities.

Agent (Acronis Backup Agent)

An application that performs data backup and recovery and enables other management operations on the machine (p. 208), 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 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. 200).

B

Backup

A backup is the result of a single backup operation (p. 200). 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 have a TIB extension. The TIB files which are the result of a backup export (p. 207) or consolidation (p. 203) are also called backups.

Backup archive (Archive)

A set of backups (p. 200) created and managed by a backup plan (p. 200). An archive can contain multiple full backups (p. 207) as well as incremental (p. 207) and differential backups (p. 204). Backups belonging to the same archive are always stored in the same location. If the backup plan includes replication (p. 210) 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. 208) 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. 200), 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. 200).

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. 200) name and location
- the backup scheme (p. 201). This includes the backup schedule and [optionally] the retention rules (p. 210)
- [optionally] additional operations to perform with the backups (replication (p. 210), validation (p. 211), conversion to a virtual machine)
- the backup options (p. 200).

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. 211) executed on a managed machine (p. 208).

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. 202)).

Backup scheme

A part of the backup plan (p. 200) that includes the backup schedule and [optionally] the retention rules and the cleanup (p. 203) schedule. For example, perform a full backup (p. 207) monthly on the last day of the month at 10:00AM and an incremental backup (p. 207) 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 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 Agent (p. 200). Bootable agent is based on Linux kernel. A machine (p. 208) can be booted into a bootable agent using either bootable media (p. 201) or Acronis PXE Server. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 203).

Bootable media

A physical media (CD, DVD, USB flash drive or other media supported by a machine (p. 208) as a boot device) that contains the bootable agent (p. 201) or Windows Preinstallation Environment (WinPE) (p. 212) with the Acronis Plug-in for WinPE (p. 199). 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. 206) 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. 209).

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. 200) adds the contents of the backup to the data catalog (p. 203). Backups are cataloged automatically as soon as they are created. Backups stored on a storage node (p. 210) are cataloged by the node. Backups stored anywhere else are cataloged by the agent (p. 200). In the backup options (p. 200), a user can choose between full and fast cataloging. Full cataloging can also be started manually.

Centralized backup plan

A backup plan (p. 200) that is deployed to a managed machine (p. 208) from the management server (p. 209). Such plan can be modified only by editing the original backup plan on the management server.

Centralized management

Management of the Acronis Backup infrastructure through a central management unit known as Acronis Backup Management Server (p. 209). The centralized management operations include:

- creating centralized backup plans (p. 202) for the registered machines (p. 209) and groups of machines
- creating and managing static (p. 210) and dynamic groups (p. 206) of machines (p. 208)
- managing the tasks (p. 211) existing on the machines
- creating and managing centralized vaults (p. 202) for storing archives
- managing storage nodes (p. 210)
- monitoring activities of the Acronis Backup components, creating reports, viewing the centralized log and more.

Centralized task

A task (p. 211) propagated to a machine from the management server (p. 209). Such task can be modified only by editing the original task or centralized backup plan (p. 202) on the management server.

Centralized vault

A networked location allotted by the management server (p. 209) administrator to serve as storage for the backup archives (p. 200). A centralized vault can be managed by a storage node (p. 210) 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. 209) on the server. The shortcut to the vault appears on the machines in the **Vaults** list. Any backup plan (p. 200) 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. 200) from a backup archive (p. 200) 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 of applying retention rules (p. 210) to an archive. The retention rules are set by the backup plan (p. 200) 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 Management Console)

A tool for remote or local access to Acronis agents (p. 200) and Acronis Backup Management Server (p. 209).

Having connected the console to the management server, the administrator sets up centralized backup plans (p. 202) and accesses other management server functionality, that is, performs centralized management (p. 202). Using the direct console-agent connection, the administrator performs direct management (p. 204).

Consolidation

Combining two or more subsequent backups (p. 200) belonging to the same archive (p. 200) into a single backup.

Consolidation might be needed when deleting backups, either manually or during cleanup (p. 203). For example, the retention rules require to delete a full backup (p. 207) that has expired but retain the next incremental (p. 207) 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. 208), users can view and search data in any vault (p. 211) accessible from this machine.

The centralized catalog available on the management server (p. 209) contains all data stored on its storage nodes (p. 210).

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. 208) in which deduplication (p. 204) is enabled.

Deduplication

A method of storing different duplicates of the same information only once.

Acronis Backup can apply the deduplication technology to backup archives (p. 200) stored on storage nodes (p. 210). 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. 207). 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. 208) using the direct console (p. 203)-agent (p. 200) connection (as opposed to centralized management (p. 202) when the operations are configured on the management server (p. 209) and propagated by the server to the managed machines).

The direct management operations include:

- creating and managing local backup plans (p. 208)
- creating and managing local tasks (p. 208) such as recovery tasks
- creating and managing personal vaults (p. 209) and archives stored there
- viewing the state, progress and properties of the centralized tasks (p. 202) 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. 201).

Disaster recovery plan (DRP)

A document 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. 200) is enabled, a DRP is created after the first successful backup is performed by the backup plan, and also after any change to the list of data items or the DRP parameters. A DRP can be sent to the specified e-mail addresses or saved as a file to a local or network folder.

Disk backup (Image)

A backup (p. 200) 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 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. 205) 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. 208) 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.)



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. 208) which is populated automatically by the management server (p. 209) according to membership criteria specified by the administrator. Acronis Backup 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. 205), or more precisely, on a disk group (p. 205). 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. 200) 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. 200), each backup belonging to the archive is encrypted by the agent (p. 200) before saving the backup to its destination.

Encrypted vault

A managed vault (p. 208) to which anything written is encrypted and anything read is decrypted transparently by the storage node (p. 210), 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. 206) will be encrypted over the encryption performed by the agent (p. 200).

Export

An operation that creates a copy of an archive (p. 200) 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. 200) or to your choice of backups belonging to the same archive. An entire vault (p. 211) can be exported by using the command line interface.

F

Full backup

A self-sufficient backup (p. 200) 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. 201) aimed to maintain the optimal balance between a backup archive (p. 200) size and the number of recovery points (p. 209) 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 (p. 44).

I

Image

The same as Disk backup (p. 204).

Incremental backup

A backup (p. 200) that stores changes to the data against the latest backup. You need access to other backups to recover data from an incremental backup.

Indexing

An activity (p. 200) performed by a storage node (p. 210) after a backup (p. 200) has been saved to a deduplicating vault (p. 204).

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. 200) during the backup operation (p. 200). A user can suspend and resume indexing.

L

Local backup plan

A backup plan (p. 200) created on a managed machine (p. 208) using direct management (p. 204).

Local task

A task (p. 211) created on a managed machine (p. 208) using direct management (p. 204).

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 Agent (p. 200) 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. 201).

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. 208), either physical or virtual, where at least one Acronis Backup Agent (p. 200) is installed.

Managed vault

A centralized vault (p. 202) managed by a storage node (p. 210). Archives (p. 200) 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. 203) and validation (p. 211) for each archive stored in the managed vault. An administrator can specify additional operations that the storage node will perform (deduplication (p. 204), encryption).

Management server (Acronis Backup Management Server)

A central server that drives data protection within the enterprise network. Acronis Backup Management Server provides the administrator with:

- a single entry point to the Acronis Backup infrastructure
- an easy way to protect data on numerous machines (p. 208) using centralized backup plans (p. 202) and grouping
- enterprise-wide monitoring and reporting functionality
- the ability to create centralized vaults (p. 202) for storing enterprise backup archives (p. 200)
- the ability to manage storage nodes (p. 210)
- the centralized catalog (p. 203) 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. 201).

P

Personal vault

A local or networked vault (p. 211) created using direct management (p. 204). 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. 200).

R

Recovery point

Date and time to which the backed-up data can be reverted.

Registered machine

A machine (p. 208) managed by a management server (p. 209). A machine can be registered on only one management server at a time. A machine becomes registered as a result of the registration (p. 210) procedure.

Registration

A procedure that adds a managed machine (p. 208) to a management server (p. 209).

Registration sets up a trust relationship between the agent (p. 200) 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. 200) 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. 200) that specifies when and how to delete or move the backups (p. 200) created by the plan.

S

Single-pass backup

A single-pass backup (aka application-aware backup) is a disk backup containing metadata of VSS-aware applications that are present on the disk. This metadata enables browsing and recovery of the backed-up application data without recovering the entire disk or volume.

Static group

A group of machines which a management server (p. 209) 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 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. 208). Storage Node enables the administrator to:

- use a single centralized catalog (p. 203) of data stored in the managed vaults

- relieve managed machines (p. 208) of unnecessary CPU load by performing cleanup (p. 203), validation (p. 211) and other operations with backup archives (p. 200) which otherwise would be performed by agents (p. 200)
- drastically reduce backup traffic and storage space taken by the archives (p. 200) by using deduplication (p. 204)
- prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted vaults (p. 207).

T

Task

A set of actions to be performed by Acronis Backup 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. 201) aimed to maintain the optimal balance between a backup archive (p. 200) size and the number of recovery points (p. 209) available from the archive. Unlike the GFS (p. 207) 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. 50)".

U

Unmanaged vault

Any vault (p. 211) that is not a managed vault (p. 208).

V

Validation

An operation that checks the possibility of data recovery from a backup (p. 200).

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. 200). 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. 203), but the total size of archives stored in the vault is limited by the storage size only.

Virtual machine

On Acronis Backup Management Server (p. 209), a machine (p. 208) is considered virtual if it can be backed up from the virtualization host without installing an agent (p. 200) 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 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. 199) enables running the Acronis Backup Agent (p. 200) in the preinstallation environment.