

Acronis® Backup & Recovery® 10 Server for Windows

Update 5

User Guide

Copyright © Acronis, Inc., 2000-2011. All rights reserved.

“Acronis” and “Acronis Secure Zone” are registered trademarks of Acronis, Inc.

“Acronis Compute with Confidence”, “Acronis Startup Recovery Manager”, “Acronis Active Restore” and the Acronis logo are trademarks of Acronis, Inc.

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>

Table of contents

1	Introducing Acronis® Backup & Recovery® 10	7
1.1	Acronis Backup & Recovery 10 overview	7
1.2	Getting started	7
1.2.1	Using the management console	8
1.3	Acronis Backup & Recovery 10 components	14
1.3.1	Agent for Windows	14
1.3.2	Management Console	15
1.3.3	Bootable Media Builder	15
1.4	Supported file systems	15
1.5	Supported operating systems	16
1.6	System requirements	17
1.7	Technical Support	17
2	Understanding Acronis Backup & Recovery 10	18
2.1	Basic concepts	18
2.2	Full, incremental and differential backups	22
2.3	User privileges on a managed machine	24
2.4	Owners and credentials	24
2.5	GFS backup scheme	25
2.6	Tower of Hanoi backup scheme	29
2.7	Retention rules	31
2.8	Backing up dynamic volumes (Windows)	34
2.9	Tape support	36
2.9.1	Tape compatibility table	36
2.9.2	Using a single tape drive	37
2.10	Support for SNMP	38
2.11	Proprietary Acronis technologies	39
2.11.1	Acronis Secure Zone	39
2.11.2	Acronis Startup Recovery Manager	40
2.11.3	Universal Restore (Acronis Backup & Recovery 10 Universal Restore)	40
2.11.4	Acronis Active Restore	42
3	Options	45
3.1	Console options	45
3.1.1	Startup page	45
3.1.2	Pop-up messages	45
3.1.3	Time-based alerts	46
3.1.4	Number of tasks	46
3.1.5	Fonts	46
3.2	Machine options	47
3.2.1	Event tracing	47
3.2.2	Log cleanup rules	49
3.2.3	Online backup proxy	49
3.2.4	Customer Experience Program	50

3.3	Default backup and recovery options.....	50
3.3.1	Default backup options	50
3.3.2	Default recovery options.....	71
4	Vaults	80
4.1	Personal vaults.....	81
4.1.1	Working with the "Personal vault" view.....	81
4.1.2	Actions on personal vaults	82
4.2	Common operations	84
4.2.1	Operations with archives stored in a vault.....	84
4.2.2	Operations with backups	85
4.2.3	Deleting archives and backups.....	86
4.2.4	Filtering and sorting archives	86
5	Scheduling	87
5.1	Daily schedule	88
5.2	Weekly schedule	90
5.3	Monthly schedule	92
5.4	At Windows Event Log event.....	94
5.5	Upon an Acronis Drive Monitor alert	96
5.6	Conditions	97
5.6.1	User is idle.....	97
5.6.2	Location's host is available	98
5.6.3	Fits time interval.....	98
5.6.4	User logged off	99
5.6.5	Time since last backup	100
6	Direct management	101
6.1	Administering a managed machine	101
6.1.1	Dashboard	101
6.1.2	Backup plans and tasks	103
6.1.3	Log.....	114
6.2	Creating a backup plan	116
6.2.1	Why is the program asking for the password?.....	119
6.2.2	Backup plan's credentials.....	119
6.2.3	Source type.....	119
6.2.4	Items to back up	120
6.2.5	Access credentials for source.....	121
6.2.6	Exclusions.....	122
6.2.7	Archive	123
6.2.8	Simplified naming of backup files	124
6.2.9	Access credentials for archive location.....	128
6.2.10	Backup schemes	129
6.2.11	Archive validation	138
6.2.12	Setting up regular conversion to a virtual machine	139
6.3	Recovering data	140
6.3.1	Task credentials	142
6.3.2	Archive selection	143
6.3.3	Data type.....	144
6.3.4	Content selection	144
6.3.5	Access credentials for location.....	145

6.3.6	Destination selection.....	145
6.3.7	Access credentials for destination	152
6.3.8	When to recover.....	153
6.3.9	Universal Restore	153
6.3.10	How to convert a disk backup to a virtual machine	154
6.3.11	Bootability troubleshooting	156
6.3.12	Recovering a vast number of files from a file backup	159
6.4	Validating vaults, archives and backups	160
6.4.1	Task credentials	161
6.4.2	Archive selection	161
6.4.3	Backup selection.....	162
6.4.4	Location selection.....	162
6.4.5	Access credentials for source.....	163
6.4.6	When to validate	163
6.5	Mounting an image.....	164
6.5.1	Archive selection	165
6.5.2	Backup selection.....	166
6.5.3	Access credentials	166
6.5.4	Volume selection	166
6.6	Managing mounted images	167
6.7	Exporting archives and backups	167
6.7.1	Task credentials	170
6.7.2	Archive selection	170
6.7.3	Backup selection.....	171
6.7.4	Access credentials for source.....	171
6.7.5	Location selection.....	172
6.7.6	Access credentials for destination	173
6.8	Acronis Secure Zone	174
6.8.1	Creating Acronis Secure Zone	174
6.8.2	Managing Acronis Secure Zone.....	176
6.9	Acronis Startup Recovery Manager	177
6.10	Bootable media.....	178
6.10.1	How to create bootable media	179
6.10.2	Connecting to a machine booted from media	186
6.10.3	Working under bootable media.....	187
6.10.4	List of commands and utilities available in Linux-based bootable media	188
6.10.5	Recovering MD devices and logical volumes.....	189
6.11	Disk management	193
6.11.1	Basic precautions.....	194
6.11.2	Running Acronis Disk Director Lite.....	194
6.11.3	Choosing the operating system for disk management	194
6.11.4	"Disk management" view.....	195
6.11.5	Disk operations	195
6.11.6	Volume operations.....	202
6.11.7	Pending operations	208
6.12	Collecting system information	209
7	Online backup	210
7.1	Introduction to Acronis Backup & Recovery 10 Online	210
7.1.1	What is Acronis Backup & Recovery 10 Online?	210
7.1.2	What data can I back up and recover?	210
7.1.3	How long will my backups be kept in the online storage?.....	211

7.1.4	How to secure my data?	211
7.1.5	Backup and recovery FAQ.....	211
7.1.6	Initial Seeding FAQ.....	212
7.1.7	Large Scale Recovery FAQ.....	217
7.1.8	Subscription lifecycle FAQ.....	219
7.2	Where do I start?	221
7.3	Choosing a subscription.....	221
7.4	Activating online backup subscriptions	222
7.4.1	Activating subscriptions	223
7.4.2	Reassigning an activated subscription	223
7.5	Configuring proxy settings	224
7.6	Limitations of the online storage.....	225
7.7	Terminology reference	225
8	Glossary.....	228

1 Introducing Acronis® Backup & Recovery® 10

1.1 Acronis Backup & Recovery 10 overview

Based on Acronis' patented disk imaging and bare metal restore technologies, Acronis Backup & Recovery 10 succeeds Acronis True Image Echo as the next generation disaster recovery solution.

Acronis Backup & Recovery 10 Server for Windows inherits the benefits of the Acronis True Image Echo product family:

- Backup of an entire disk or volume, including the operating system, all applications, and data
- Bare metal recovery to any hardware
- File and folder backup and recovery.

Acronis Backup & Recovery 10 Server for Windows offers new benefits that help organizations meet challenging Recovery Time Objectives while reducing both capital expense and software maintenance costs.

- **Leveraging existing IT infrastructure**
Backward compatibility and an easy upgrade from Acronis True Image Echo
- **Highly automated data protection**
All-round planning of data protection (backup, retention and validation of backups) within a backup policy
Built-in Tower of Hanoi and Grandfather-Father-Son backup schemes with customizable parameters
A variety of events and conditions can be chosen to trigger a backup
- **Easy work with virtual environments**
Conversion of a backup to a fully configured VMware, Microsoft, Parallels, Citrix or Red Hat KVM virtual machine
- **Redesigned GUI**
Dashboard for quick operational decision making
Overview of all configured and running operations with color-coding for successful and failed operations
- **Enterprise level of security**
Controlling user rights to perform operations and access backups
Running services with minimal user rights.

1.2 Getting started

Direct management

1. Install Acronis Backup & Recovery 10 Management Console and Acronis Backup & Recovery 10 Agent.
2. Start the console.

Windows

Start the console by selecting it from the start menu.

3. Connect the console to the machine where the agent is installed.

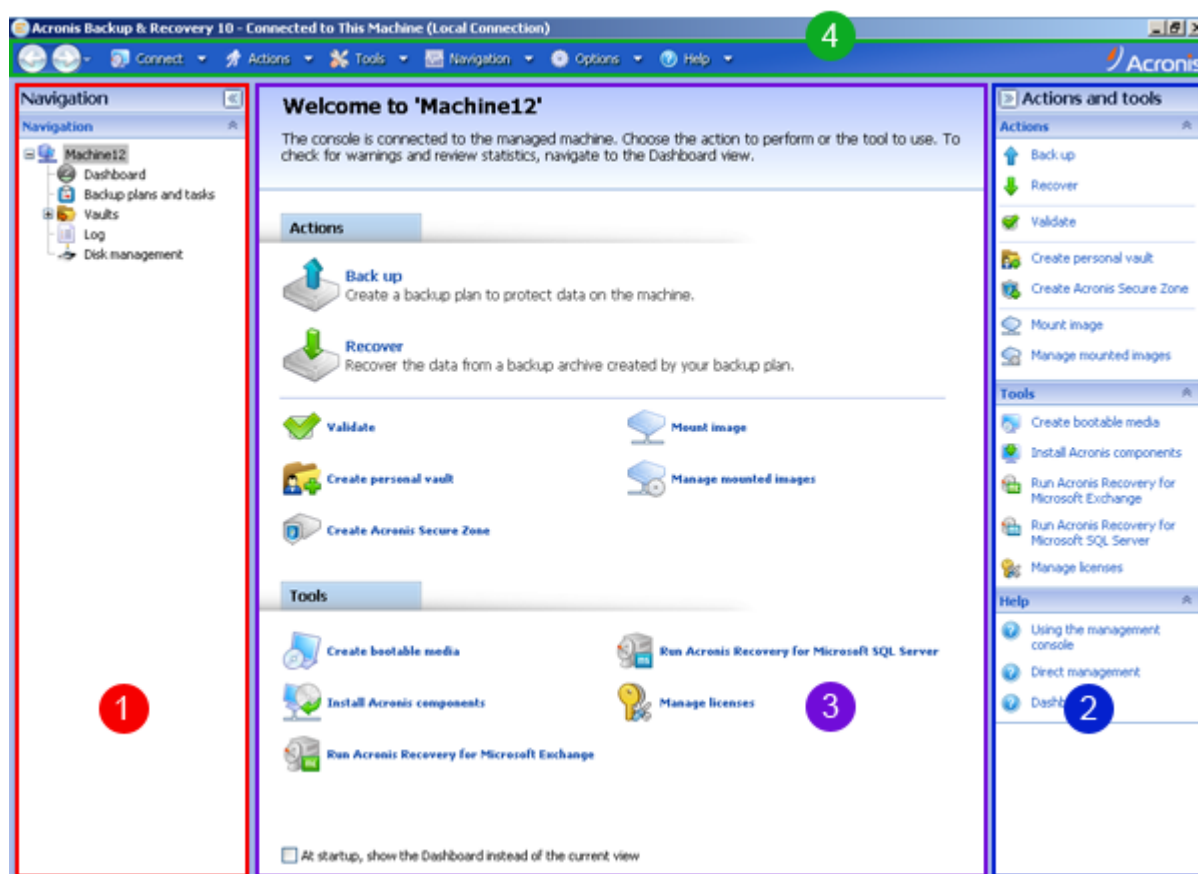
Where to go from here

For what to do next see "Basic concepts (p. 18)".

For understanding of the GUI elements see the next section.

1.2.1 Using the management console

As soon as the console connects to a managed machine (p. 237) or to a management server (p. 238), the respective items appear across the console's workspace (in the menu, in the main area with the **Welcome** screen, the **Navigation** pane, the **Actions and tools** pane) enabling you to perform agent-specific or server-specific operations.



Acronis Backup & Recovery 10 Management Console - Welcome screen

Key elements of the console workspace

	Name	Description
1	Navigation pane	Contains the Navigation tree and the Shortcuts bar and lets you navigate to the different views (see the Navigation pane (p. 9) section.)
2	Actions and tools pane	Contains bars with a set of actions that can be performed and tools (see the Actions and Tools pane (p. 10) section).
3	Main area	The main place of working, where you create, edit and manage backup plans, policies, tasks and perform other operations. Displays the different views and action pages (p. 11) depending on items selected in the menu, Navigation

		tree, or on the Actions and Tools pane.
4	Menu bar	Appears across the top of the program window and lets you perform all the operations, available on both panes. Menu items change dynamically.

1024x768 or higher display resolution is required for comfortable work with the management console.

1.2.1.1 "Navigation" pane







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

Navigation tree

The **Navigation** tree enables you to navigate across the program views. Views depend on whether the console is connected to a managed machine or to the management server.

Views for a managed machine

When the console is connected to a managed machine, the following views are available in the navigation tree.

-  **[Machine name]**. Root of the tree also called a **Welcome** view. 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.
 -  **Dashboard**. Use this view to estimate at a glance whether the data is successfully protected 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 states and statuses, monitor plans.
 -  **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, mount backups as virtual drives, etc.
 -  **Log**. Use this view to examine information on operations performed by the program on the managed machine.
 -  **Disk management**. Use this view to perform operations on the machine's hard disk drives.

Shortcuts bar

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

To add a shortcut to a machine

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

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

1.2.1.2 "Actions and tools" pane


The **Actions and tools** pane enables you to easily and efficiently work with Acronis Backup & Recovery 10. The pane's bars provide quick access to program's operations and tools. All items of the **Actions and tools** bar are duplicated in the program menu.

Bars

'[Item's name]' actions

Contains a set of actions that can be performed on the items selected in any of the navigation views. Clicking the action opens the respective action page (p. 12). Items of different navigation views have their own set of actions. The bar's name changes in accordance with the item you select. For example, if you select the backup plan named *System backup* in the **Backup plans and tasks** view, the actions bar will be named as **'System backup' actions** and will have the set of actions typical to backup plans.

All actions can also be accessed in the respective menu items. A menu item appears on the menu bar when you select an item in any of the navigation views.

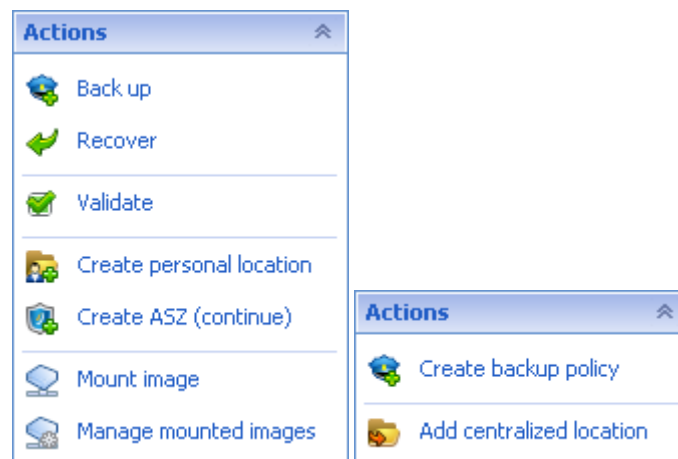


Examples of "'Item name' actions" bars

Actions

Contains a list of common operations that can be performed on a managed machine or on a management server. Always the same for all views. Clicking the operation opens the respective action page (see the Action pages (p. 12) section.)

All the actions can also be accessed in the **Actions** menu.




"Actions" bar on a managed machine and on a management server

Tools

Contains a list of the Acronis tools. Always the same across all the program views.

All the tools can also be accessed in the **Tools** menu.



"Tools" bar

Help

Contains a list of help topics. Different views and action pages of Acronis Backup & Recovery 10 provided with lists of specific help topics.

1.2.1.3 Operations with panes

How to expand/minimize panes

By default, the **Navigation** pane appears expanded and the **Actions and Tools** - minimized. You might need to minimize the pane in order to free some additional workspace. To do this, click the chevron (◀ - for the **Navigation** pane; ▶ - for the **Actions and tools** pane). 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.

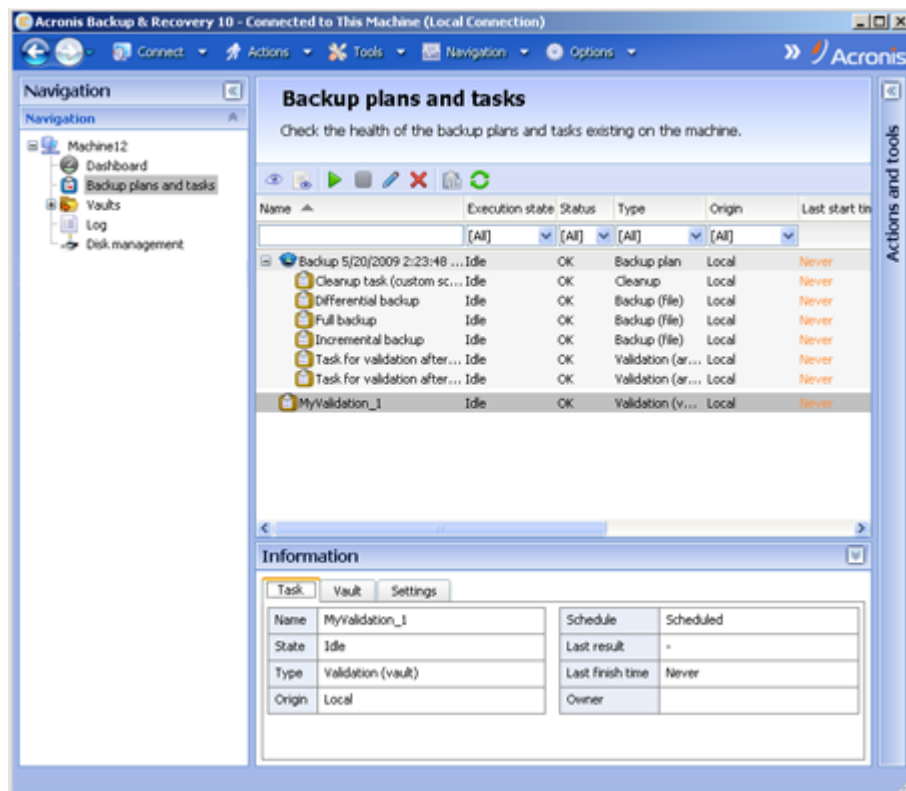
The management console "remembers" the way the panes' borders are set. When you run the management console next time, all the panes' borders will have the same position that was set previously.

1.2.1.4 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, policies, tasks and perform other operations. The main area displays different views and action pages according the items you select in the menu, **Navigation** tree, or on the **Actions and Tools** pane.

Views

A view appears on the main area when clicking any item in the **Navigation** tree in the Navigation pane (p. 9).



"Tasks" view

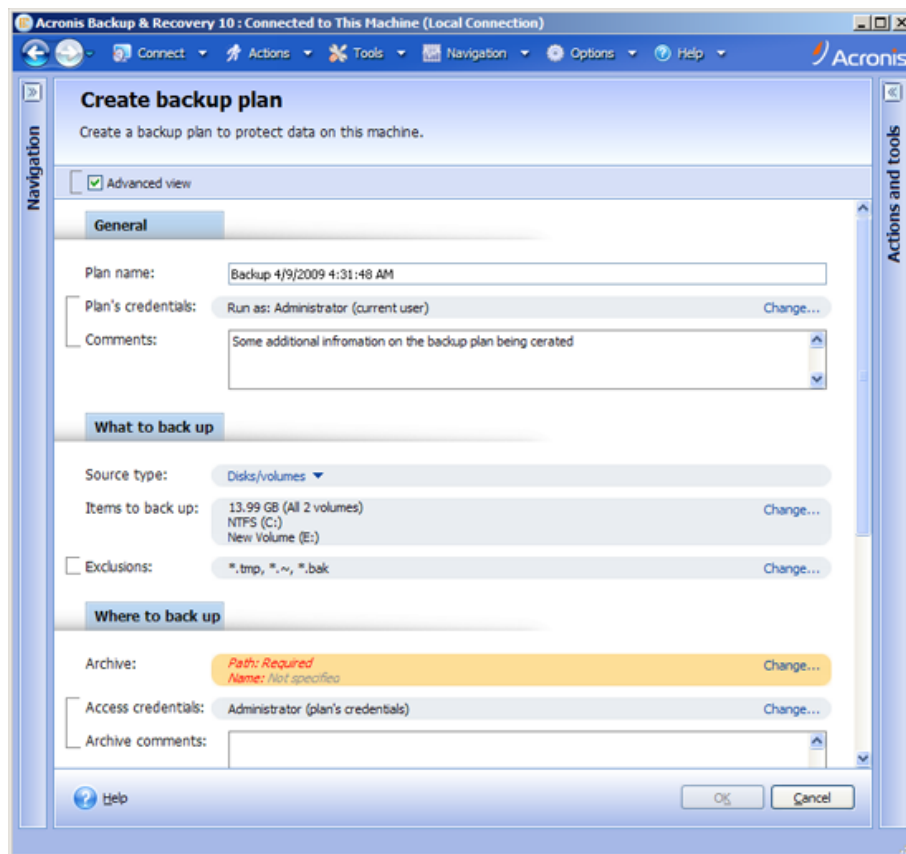
Common way of working with views

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

- Use filtering and sorting 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
- 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 clicking in the items in the **[Item's name] Actions** bar (on the **Actions and Tools** pane);
 - By selecting the items in the **Actions** menu;
 - By right-clicking the item and selecting the operation in the context menu.

Action pages

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




Action page - Create backup plan

Using controls and specifying settings


The action pages offer two ways of representation: basic and advanced. The basic representation hides such fields as credentials, comments, etc. When the advanced representation is enabled, all the available fields are displayed. You can switch between the views by selecting the **Advanced view** check box at the top of the action page.

Most settings are configured by clicking the respective **Change...** links to the right. Others are selected from the drop-down list, or typed manually in the page's fields.



Action page - Controls

Acronis Backup & Recovery 10 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

1.3 Acronis Backup & Recovery 10 components

This section contains a list of Acronis Backup & Recovery 10 components with a brief description of their functionality.

Components for a managed machine (agents)

These are applications that perform data backup, recovery and other operations on the machines managed with Acronis Backup & Recovery 10. Agents require a license to perform operations on each managed machine. Agents have multiple features, or add-ons, that enable additional functionality and so might require additional licenses.

Console

The console provides Graphical User Interface and remote connection to the agents. Usage of the console is not licensed.

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. Availability of the agent add-ons in a rescue environment depends on whether an add-on is installed on the machine where the media builder is working.

1.3.1 Agent for Windows

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

Disk backup

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

File backup

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

Other operations

Conversion to a virtual machine

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

Disk management

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

1.3.1.1 Universal Restore

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

1.3.2 Management Console

Acronis Backup & Recovery 10 Management Console is an administrative tool for local access to Acronis Backup & Recovery 10 agent.

1.3.3 Bootable Media Builder

Acronis Bootable Media Builder is a dedicated tool for creating bootable media (p. 231). The media builder that installs on Windows can create bootable media based on either Windows Preinstallation Environment, or Linux kernel.

The Universal Restore (p. 15) add-on enables you to create bootable media with the restore to dissimilar hardware functionality. Universal Restore handles differences in devices that are critical for Windows start-up, such as storage controllers, motherboard or chipset.

1.4 Supported file systems

Acronis Backup & Recovery 10 can back up and recover the following file systems with the following limitations:

- FAT16/32
- NTFS
- Ext2/Ext3/Ext4
- ReiserFS3 - particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 10 Storage Node

- ReiserFS4 - volume recovery without the volume resize capability; particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 10 Storage Node
- XFS - volume recovery without the volume resize capability; particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 10 Storage Node
- JFS - particular files cannot be recovered from disk backups located on Acronis Backup & Recovery 10 Storage Node
- Linux SWAP

Acronis Backup & Recovery 10 can back up and recover corrupted or non-supported file systems using the sector-by-sector approach.

1.5 Supported operating systems

Acronis Backup & Recovery 10 Management Console, Acronis Backup & Recovery 10 Agent for Windows

- Windows XP Professional SP2+ (x86, x64)
- Windows 2000 SP4 – all editions except for the Datacenter edition
- Windows Server 2003/2003 R2 - the Standard and Enterprise editions (x86, x64)
- Windows Small Business Server 2003/2003 R2 (x86)
- Windows Vista - all editions (x86, x64)
- Windows 7 SP1 - all editions (x86, x64)
- Windows Server 2008 - the Standard and Enterprise editions (x86, x64)
- Windows Small Business Server 2008 (x64)
- Windows Small Business Server 2011
- Windows Server 2008 R2 SP1 - the Standard, Enterprise, Datacenter, Foundation editions
- Windows MultiPoint Server 2010

Remote connection to the agent is not possible.

Acronis products do not support systems with Extensible Firmware Interface (EFI). Although it is possible to restore a GPT partition with Acronis if Windows is installed on it, the restored system will not be bootable. Acronis Backup & Recovery 10 can back up and recover operating systems if they are installed in BIOS/MBR mode, even if they run on EFI-capable servers. Most servers have BIOS settings that allow booting the installation CD in BIOS/MBR mode instead of EFI mode. The MBR mode ensures that after installation the boot disk is partitioned in MBR standard, not GPT.

1.6 System requirements

The components installed in Windows

Edition name	Memory (above the OS and running applications)	Disk space required during installation or update	Disk space occupied by the component(s)	Additional
Server for Windows	80 MB	1 GB	500 MB	Screen resolution 1024*768 pixels or higher
Workstation	80 MB	1 GB	500 MB	
Bootable Media Builder (Windows)	80 MB	700 MB	300 MB	

Bootable media

Media type	Memory	ISO image size	Additional
Based on Windows PE	512 MB	300 MB	
Linux-based	256 MB	130 MB	

1.7 Technical Support

Maintenance and Support Program

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

Product Updates

You can download the latest updates for all your registered Acronis software products from our website at any time after logging into your **Account** (<https://www.acronis.eu/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 Understanding Acronis Backup & Recovery 10

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.

2.1 Basic concepts

Please familiarize yourself with the basic notions used in the Acronis Backup & Recovery 10 graphical user interface and documentation. Advanced users are welcome to use this section as a step-by-step quick start guide. The details can be found in the context help.

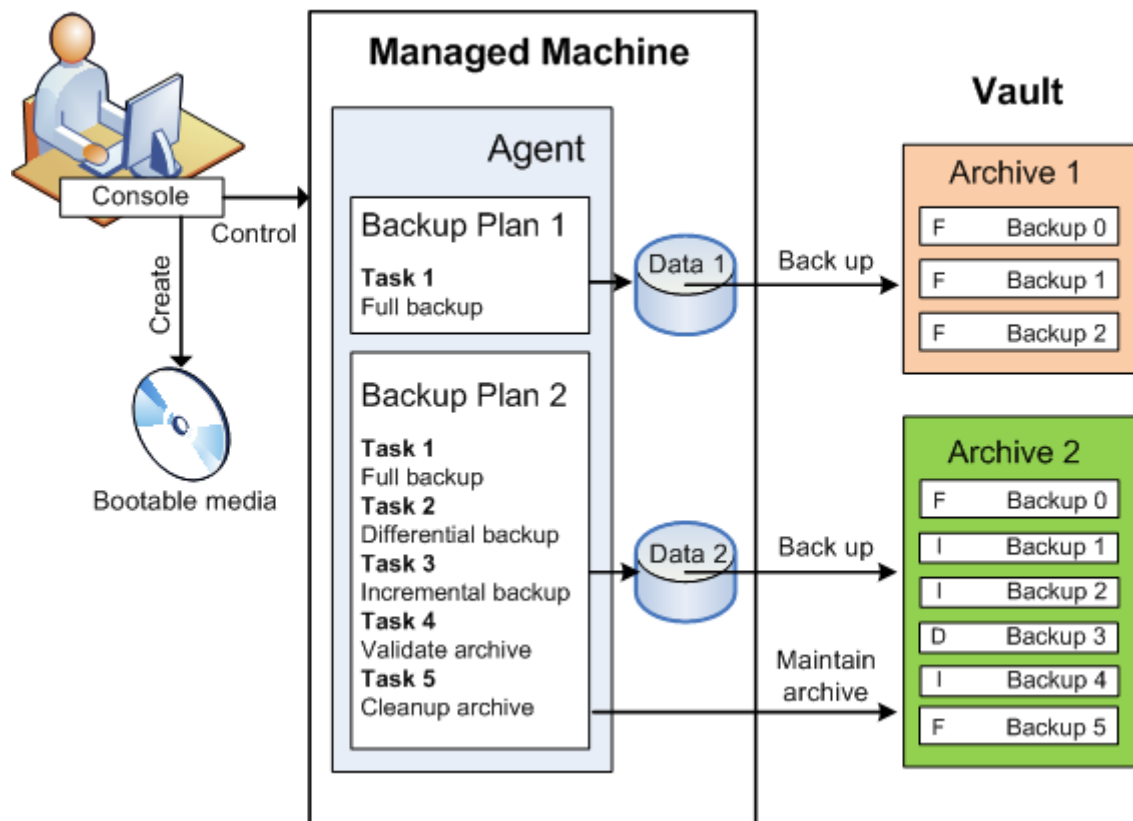
Backup under operating system

1. To protect data on a machine, install Acronis Backup & Recovery 10 agent (p. 229) on the machine which becomes a managed machine (p. 237) from this point on.
2. To be able to manage the machine using Graphical User Interface, install Acronis Backup & Recovery 10 Management Console (p. 233) on the same machine or any machine from which you prefer to operate. If you have the standalone product edition, skip this step since in your case the console installs with the agent.
3. Run the console. To be able to recover the machine's operating system if the system fails to start, create bootable media (p. 231).
4. Connect the console to the managed machine.
5. Create a backup plan (p. 230).

To do so, you have to specify, at the very least, the data to be protected and the location where the backup archive (p. 229) will be stored. This will create a minimal backup plan consisting of one task (p. 240) that will create a full backup (p. 229) of your data every time the task is manually started. A complex backup plan might consist of multiple tasks which run on schedule; create full, incremental or differential backups (p. 22); perform archive maintenance operations such as backup validation (p. 241) or deleting outdated backups (archive cleanup (p. 232)). You can customize backup operations using various backup options, such as pre/post backup commands, network bandwidth throttling, error handling or notification options.


6. Use the **Backup plans and tasks** page to view information about your backup plans and tasks and monitor their execution. Use the **Log** page to browse the operations log.
7. The location where you store backup archives is called a vault (p. 241). Navigate to the **Vaults** page to view information about your vaults. Navigate further to the specific vault to view archives and backups and perform manual operations with them (mounting, validating, deleting, viewing contents). You can also select a backup to recover data from it.

The following diagram illustrates the notions discussed above. For more definitions please refer to the Glossary.



Backup using bootable media

You can boot the machine using the bootable media, configure the backup operation in the same way as a simple backup plan and execute the operation. This will help you extract files and logical volumes from a system that failed to boot, take an image of the offline system or back up sector-by-sector an unsupported file system.



Recovery under operating system

When it comes to data recovery, you create a recovery task on the managed machine. You specify the vault, then select the archive and then select the backup referring to the date and time of the backup creation, or more precisely, to the time when the creation has started. In most cases, the data will be reverted to that moment.


Examples of exceptions to this rule:

Recovering a database from a backup that contains the transaction log (a single backup provides multiple recovery points and so you can make additional selections).

Recovering multiple files from a file backup taken without snapshot (each file will be reverted to the moment when it was actually copied to the backup).

You also specify the destination where to recover the data. You can customize the recovery operation using recovery options, such as pre/post recovery commands, error handling or notification options.


The following diagram illustrates data recovery under the operating system (online). No backup can proceed on the machine while the recovery operation is taking place. If required, you can connect the console to another machine and configure a recovery operation on that machine. This ability (remote parallel recovery) first appeared in Acronis Backup & Recovery 10; the previous Acronis products do not provide it.



Recovery using bootable media

Recovery over a volume locked by the operating system, such as the volume where the operating system resides, requires a reboot to the bootable environment which is a part of the agent. After the recovery is completed, the recovered operating system goes online automatically.

If the machine fails to boot or you need to recover data to bare metal, you boot the machine using the bootable media and configure the recovery operation in the same way as the recovery task. The following diagram illustrates the recovery using the bootable media.



2.2 Full, incremental and differential backups

Acronis Backup & Recovery 10 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. 229) 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

2.3 User privileges on a managed machine

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

Regular users

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

- Perform file-level backup and recovery of the files that the user has permissions to access—but without using a file-level backup snapshot.
- Create backup plans and tasks and manage them.
- View—but not manage—backup plans and tasks created by other users.
- View the local event log.

Administrative users

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

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

Members of the Administrators group also can:

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

2.4 Owners and credentials

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

Plan (task) owner

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

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

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

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

Managing a plan (task) owned by another user

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

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

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

Archive owner

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

Plan's credentials and task credentials

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

Manual start

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

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

1. Select the **Advanced view** check box.
2. Select **General -> Plan's (Task) credentials -> Change**.
3. Enter the credentials under which the plan (task) will run.

Scheduled or postponed start

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

Why does the program compel me to specify credentials?

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

2.5 GFS backup scheme

This section covers implementation of the Grandfather-Father-Son (GFS) backup scheme in Acronis Backup & Recovery 10.

With this backup scheme you are not allowed to back up more often than once a day. The scheme enables you to mark out the daily, weekly and monthly cycles in your daily backup schedule and set the retention periods for the daily, monthly and weekly backups. The daily backups are referred to as

“sons”; weekly backups are referred to as “fathers”; the longest lived monthly backups are called “grandfathers”.

GFS as a tape rotation scheme

GFS was initially created and is often referred to as a tape rotation scheme. Tape rotation schemes, as such, do not provide automation. They just determine:

- how many tapes you need to enable recovery with the desired resolution (time interval between recovery points) and roll-back period
- which tapes you should overwrite with the forthcoming backup.

Tape rotation schemes enable you to get by with the minimal number of cartridges and not to be buried in used tapes. A lot of Internet sources describe varieties of the GFS tape rotation scheme. You are free to use any of the varieties when backing up to a locally attached tape device.

GFS by Acronis

With Acronis Backup & Recovery 10, you can easily set up a backup plan that will regularly back up data and clean up the resulting archive according to the GFS scheme.

Create the backup plan as usual. For the backup destination, choose any storage device where automatic cleanup can be performed, such as an HDD-based storage device or robotic tape library. (Since the space freed on the tape after cleanup cannot be reused until all the tape becomes free, take into account additional considerations when using GFS on a tape library.)

The following is an explanation of the settings that are specific for the GFS backup scheme.

GFS-related settings of the backup plan

Start backup at:

Back up on:

This step creates the total backup schedule, that is, defines all the days you need to back up on.

Assume you select backing up at 8:00 PM on workdays. Here is the total schedule you have defined.

“B” stands for “backup”.

Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat
	B	B	B	B	B			B	B	B	B	B			B	B	B	B	B			B	B	B	B	B	

The total schedule.

Schedule: Workdays at 8:00 PM

Weekly/Monthly

This step forms the daily, weekly and monthly cycles in the schedule.

Select a day of the week from the days selected in the previous step. Each 1st, 2nd and 3rd backup created on this day of the week will be considered as a weekly backup. Each 4th backup created on this day of the week will be considered as a monthly backup. Backups created on the other days will be considered as daily backups.

Assume you select Friday for Weekly/Monthly backup. Here is the total schedule marked out according to the selection.

“D” stands for the backup that is considered Daily. “W” stands for the backup that is considered Weekly. “M” stands for the backup that is considered Monthly.

	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat						
Total schedule	D	D	D	D	W			D	D	D	D	W			D	D	D	D	W			D	D	D	D	W			D	D	D	D	M	

The schedule marked out according to the GFS scheme.

Schedule: Workdays at 8:00 PM

Weekly/Monthly: Friday

Acronis uses incremental and differential backups that help save storage space and optimize the cleanup so that consolidation is not needed. In terms of backup methods, weekly backup is differential (Dif), monthly backup is full (F) and daily backup is incremental (I). The first backup is always full.

The Weekly/Monthly parameter splits the total schedule into daily, weekly and monthly schedules.

Assume you select Friday for Weekly/Monthly backup. Here is the real schedule of the backup tasks that will be created.

	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat	Su	Mo	Tue	We	Thu	Fri	Sat
Total schedule	D	D	D	D	W			D	D	D	D	W			D	D	D	D	W			D	D	D	D	M		
Daily task	F	I	I	I				I	I	I	I				I	I	I	I				I	I	I	I			
Weekly task	_____ Dif _____						_____ Dif _____						_____ Dif _____															
Monthly task	_____ F _____																											

Backup tasks created according to the GFS scheme by Acronis Backup & Recovery 10.

Schedule: Workdays at 8:00 PM

Weekly/Monthly: Friday

Keep backups: Daily

This step defines the retention rule for daily backups. The cleanup task will run after each daily backup and delete all daily backups that are older than you specify.

Keep backups: Weekly

This step defines the retention rule for weekly backups. The cleanup task will run after each weekly backup and delete all weekly backups that are older than you specify. The weekly backups' retention period cannot be less than the daily backups' retention period. It is usually set several times longer.

Keep backups: Monthly

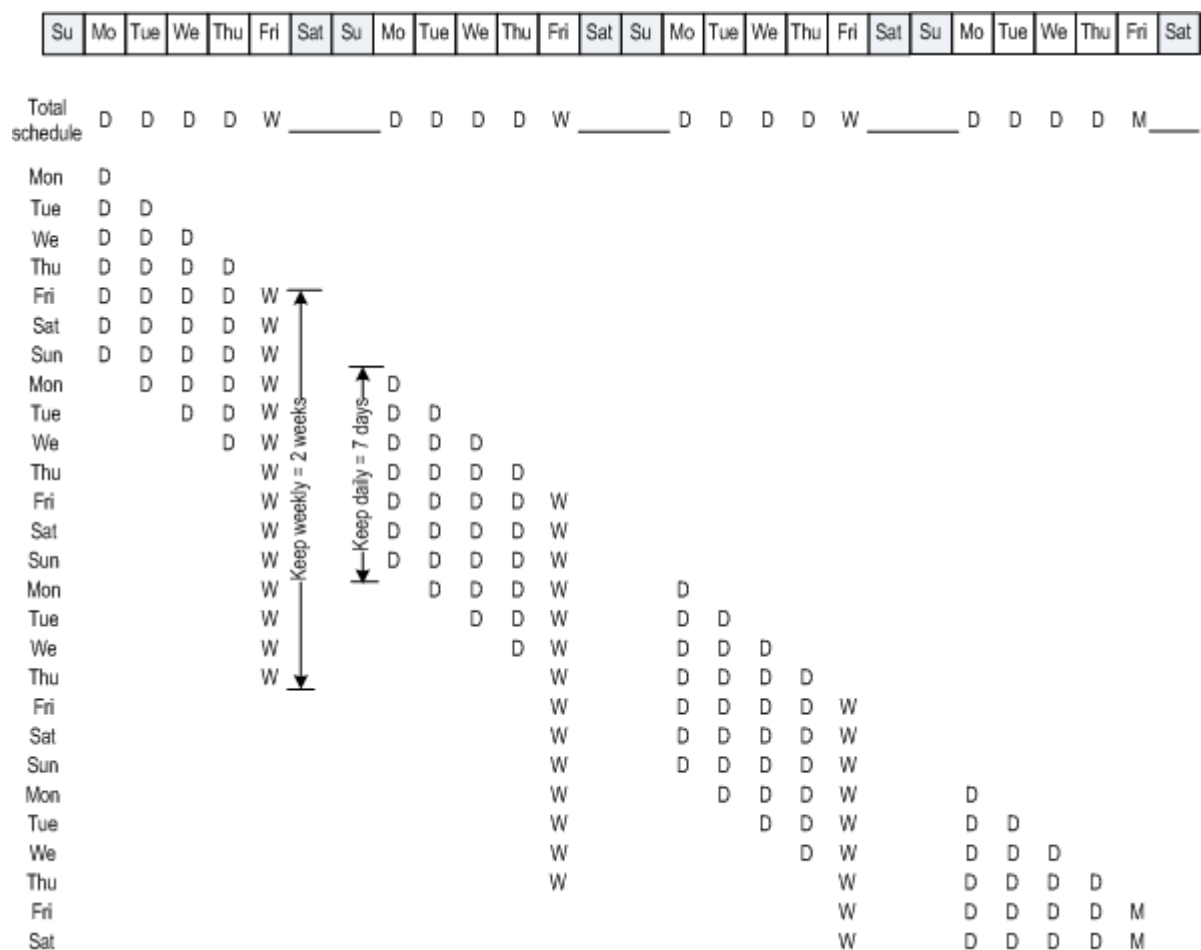
This step defines the retention rule for monthly backups. The cleanup task will run after each monthly backup and delete all monthly backups that are older than you specify. The monthly backups' retention period cannot be less than the weekly backups' retention period. It is usually set several times longer. You have the option to keep the monthly backups indefinitely.

The resulting archive: ideal

Assume you select to keep daily backups for 7 days, weekly backups for 2 weeks and monthly backups for 6 months. Here is how your archive would appear after the backup plan is launched if all the backups were full and so could be deleted as soon as the scheme requires.

The left column shows days of the week. For each day of the week, the content of the archive after the regular backup and the subsequent cleanup is shown.

“D” stands for the backup that is considered Daily. “W” stands for the backup that is considered Weekly. “M” stands for the backup that is considered Monthly.



An ideal archive created according to the GFS scheme.

Schedule: Workdays at 8:00 PM

Weekly/Monthly: Friday

Keep daily backups: 7 days

Keep weekly backups: 2 weeks

Keep monthly backups: 6 months

Starting from the third week, weekly backups will be regularly deleted. After 6 months, monthly backups will start to be deleted. The diagram for weekly and monthly backups will look similar to the week-based timescale.


The resulting archive: real

In reality, the archive content will somewhat differ from the ideal scheme.

When using the incremental and differential backup methods, you cannot delete a backup as soon as the scheme requires if later backups are based on this backup. Regular consolidation is unacceptable because it takes too much system resources. The program has to wait until the scheme requires the deletion of all the dependent backups and then deletes the entire chain.

Here is how the first month of your backup plan will appear in real life. “F” stands for full backup. “Dif” stands for differential backup. “I” stands for incremental backup.

The backups that outlive their nominal lifetime because of dependencies are marked pink. The initial full backup will be deleted as soon as all differential and incremental backups based on this backup are deleted.



An archive created according to the GFS scheme by Acronis Backup & Recovery 10.

Schedule: Workdays at 8:00 PM

Weekly/Monthly: Friday

Keep daily backups: 7 days

Keep weekly backups: 2 weeks

Keep monthly backups: 6 months

2.6 Tower of Hanoi backup scheme

The need to have frequent backups always conflicts with the cost of keeping such backups for a long time. The Tower of Hanoi (ToH) backup scheme is a useful compromise.

Tower of Hanoi overview

The Tower of Hanoi scheme is based on a mathematical puzzle of the same name. In the puzzle a series of rings are stacked in size order, the largest on the bottom, on one of three pegs. The goal is to move the ring series to the third peg. You are only allowed to move one ring at a time, and are prohibited from placing a larger ring above a smaller ring. The solution is to shift the first ring every other move (moves 1, 3, 5, 7, 9, 11...), the second ring at intervals of four moves (moves 2, 6, 10...), the third ring at intervals of eight moves (moves 4, 12...), and so on.

For example, if there are five rings labeled A, B, C, D, and E in the puzzle, the solution gives the following order of moves:

Move \ Ring	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	A		A		A		A		A		A		A		A		A		A		A		A		A		A		A		A
2		B				B				B				B				B				B				B				B	
3				C								C									C								C		
4								D																D							
5																E															

The Tower of Hanoi backup scheme is based on the same patterns. It operates with **Sessions** instead of **Moves** and with **Backup levels** instead of **Rings**. Commonly an N-level scheme pattern contains (N-th power of two) sessions.

So, the five-level Tower of Hanoi backup scheme cycles the pattern that consists of 16 sessions (moves from 1 to 16 in the above figure).

The table shows the pattern for the five-level backup scheme. The pattern consists of 16 sessions.

Session \ Backup level	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	A		A		A		A		A		A		A		A	
2		B				B				B				B		
3				C								C				
4								D								
5																E

The Tower of Hanoi backup scheme implies keeping only one backup per level. All the outdated backups have to be deleted. So the scheme allows for efficient data storage: more backups accumulate toward the present time. Having four backups, you can recover data as of today, yesterday, half a week ago, or a week ago. For the five-level scheme you can also recover data backed up two weeks ago. So every additional backup level doubles the maximal roll-back period for your data.

Tower of Hanoi by Acronis

The Tower of Hanoi backup scheme is generally too complex to mentally calculate the next media to be used. But Acronis Backup & Recovery 10 provides you with automation of the scheme usage. You can set up the backup scheme while creating a backup plan.

Acronis implementation for the scheme has the following features:

- up to 16 backup levels
- incremental backups on first level (A) - to gain time and storage savings for the most frequent backup operations; but data recovery from such backups takes longer because it generally requires access to three backups
- full backups on the last level (E for five-level pattern) - the rarest backups in the scheme, take more time and occupy more space in storage
- differential backups on all intermediate levels (B, C and D for five-level pattern)

- the pattern starts with a full backup since the very first backup cannot be incremental
- the scheme forces every backup level to keep only the most recent backup, other backups from the level have to be deleted; however backup deletion is postponed in cases where the backup is a base for another incremental or differential one
- an old backup on a level is kept until a new backup has been successfully created on the level.

The table shows the pattern for the five-level backup scheme. The pattern consists of 16 sessions.

Session Backup level	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1 (Incremental)		A		A		A		A		A		A		A		A
2 (Differential)			B				B				B				B	
3 (Differential)					C								C			
4 (Differential)									D							
5 (Full)	E															

As a result of using incremental and differential backups the situation may arise when an old backup deletion must be postponed as it still is a base for other backups. The table below indicates the case when deletion of full backup (E) created at session 1 is postponed at session 17 until session 25 because the differential backup (D) created at session 9 is still actual. In the table all cells with deleted backups are grayed out:

Session Backup level	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1 (Incremental)		A		A		A		A		A		A		A		A		A		A		A		A	
2 (Differential)			B				B				B				B			B				B			
3 (Differential)					C							C								C					
4 (Differential)									D																D
5 (Full)	E																E								

Differential backup (D) created at session 9 will be deleted at session 25 after creation of a new differential backup is completed. This way, a backup archive created in accordance with the Tower of Hanoi scheme by Acronis sometimes includes up to two additional backups over the classical implementation of the scheme.

For information about using Tower of Hanoi for tape libraries, see Using the Tower of Hanoi tape rotation scheme.


2.7 Retention rules

The backups produced by a backup plan make an archive. The two retention rules described in this section enable you to limit the archive size and set the lifetime (retention period) of the backups.

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. 135) if you accept the risk of losing the last backup.


1. Delete backups older than

This is a time interval counted back from the moment when the retention rules are applied. Every time a retention rule is applied, the program calculates the date and time in the past corresponding to this interval and deletes all backups created before that moment. None of the backups created after this moment will be deleted.




2. Keep the archive size within

This is the maximum size of the archive. Every time a retention rule is applied, the program compares the actual archive size with the value you set and deletes the oldest backups to keep the archive size within this value. The diagram below shows the archive content before and after the deletion.

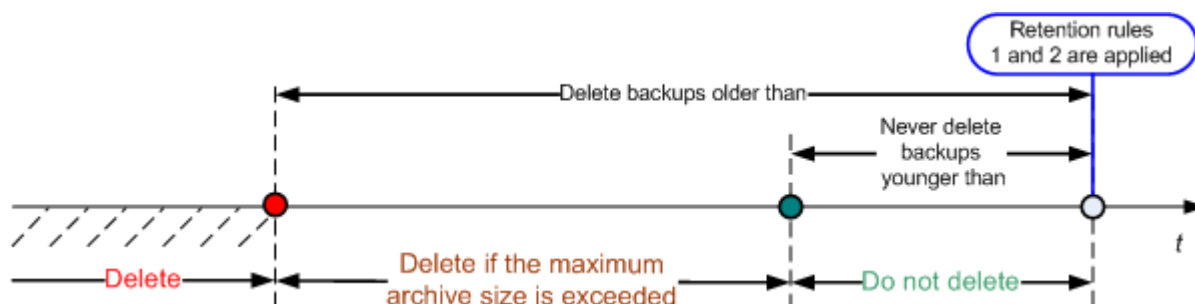


There is a certain risk that all but one backup will be deleted if the maximum archive size is set improperly (too small) or a regular backup turns out to be too large. To protect the recent backups from deletion, select the **Never delete backups younger than** check box and specify the maximum age of backups that must be retained. The diagram below illustrates the resulting rule.



Combination of rules 1 and 2

You can limit both the backups' lifetime and the archive size. The diagram below illustrates the resulting rule.



Example

Delete backups older than = 3 Months

Keep the archive size within = 200GB

Never delete backups younger than = 10 Days

- Every time the retention rules are applied, the program will delete all backups created more than 3 months (or more exactly, 90 days) ago.
- If after the deletion the archive size is more than 200GB, and the oldest backup is older than 10 days, the program will delete that backup.
- Then, if necessary, the next old backup will be deleted, until the archive size decreases to the preset limit or the oldest backup age reaches 10 days.

Deleting backups with dependencies

Both retention rules presume deleting 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 of a backup affects other backups, one of the following rules is applied:

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

The outdated backup will be kept until all backups that depend on it also become outdated. Then all the chain will be deleted at once during the regular cleanup. 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 can exceed the values you specify.

- **Consolidate the backup**

The program will consolidate the backup that is subject to deletion with the next dependent backup. For example, the retention rules require to 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 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 backups' age are within the bounds you specify. The consolidation, however, may take a lot of time and system resources. And you still need some extra space in the vault for temporary files created during consolidation.

What you need to know about consolidation

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

Backups resulting from consolidation always have maximum compression. This means that all backups in an archive may acquire the maximum compression as a result of repeated cleanup with consolidation.

Best practices

Maintain the balance between the storage device capacity, the restrictive parameters you set and the cleanup frequency. The retention rules logic assumes that the storage device capacity is much more than the average backup size and the maximum archive size does not come close to the physical storage capacity, but leaves a reasonable reserve. Due to this, exceeding the archive size that may occur between the cleanup task runs will not be critical for the business process. The rarer the cleanup runs, the more space you need to store backups that outlive their lifetime.

The Vaults (p. 80) page provides you with information about free space available in each vault. Check this page from time to time. If the free space (which in fact is the storage device free space) approaches zero, you might need to toughen the restrictions for some or all archives residing in this vault.

2.8 Backing up dynamic volumes (Windows)

This section explains in brief how to back up and recover dynamic volumes (p. 236) using Acronis Backup & Recovery 10. Basic disks that use the GUID Partition Table (GPT) are also discussed.

Dynamic volume is a volume located on dynamic disks (p. 235), or more exactly, on a disk group (p. 234). Acronis Backup & Recovery 10 supports the following dynamic volume types/RAID levels:

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

Acronis Backup & Recovery 10 can back up and recover dynamic volumes and, with minor limitations, basic GPT volumes.

Backing up dynamic volumes

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

Command line examples

```
trueimagecmd /create /partition:DYN1,DYN2 /asz
```

This will back up DYN1 and DYN2 volumes to the Acronis Secure Zone.

```
trueimagecmd /create /harddisk:DYN /asz
```

This will back up all dynamic volumes in the system to the Acronis Secure Zone.

The boot code on basic GPT volumes is not backed up or recovered.

Recovering dynamic volumes

A dynamic volume can be recovered

- over any type of existing volume
- to unallocated space of a disk group
- to unallocated space of a basic disk.

Recovery over an existing volume

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

Recovery to disk group unallocated space

When a dynamic volume is recovered to disk group unallocated space, both the type and the content of the resulting volume are recovered. The unallocated space size has to be enough to accommodate the backup content. The way unallocated space is distributed among the disks is also important.

Example

Striped volumes consume equal portions of space on each disk.

Assume you are going to recover a 30GB striped volume to a disk group consisting of two disks. Each disk has volumes and a certain amount of unallocated space. The total size of unallocated space is 40GB. The recovery will always result in a striped volume if the unallocated space is distributed evenly among the disks (20GB and 20GB).

If one of the disks has 10GB and the other has 30GB of unallocated space, then the recovery result depends on the size of the data being recovered.

- If the data size is less than 20GB, then one disk can hold, say, 10GB; the other will hold the remaining 10GB. This way, a striped volume will be created on both disks and 20GB on the second disk will remain unallocated.
- If the data size is more than 20GB, the data cannot be distributed evenly between the two disks, but can fit into a single simple volume. A simple volume accommodating all the data will be created on the second disk. The first disk will remain untouched.

	Backed up (source):		
Recovered to:	Dynamic volume	Basic MBR volume	Basic GPT volume
Dynamic volume	Dynamic volume Type as of the target	Dynamic volume Type as of the target	Dynamic volume Type as of the target
Unallocated space (disk group)	Dynamic volume Type as of the source	Dynamic volume Simple	N/A
Basic MBR volume	Basic MBR volume	Basic MBR volume	Basic MBR volume
Basic GPT volume	Basic GPT volume	Basic GPT volume	Basic GPT volume
Unallocated space (basic MBR disk)	Basic MBR volume	Basic MBR volume	Basic MBR volume

Unallocated space (basic GPT disk)	Basic GPT volume	Basic GPT volume	Basic GPT volume
------------------------------------	------------------	------------------	------------------

Moving and resizing volumes during recovery

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

Preparing disk groups and volumes

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

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

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

Acronis Backup & Recovery 10 includes a handy disk management utility which enables you to perform the above operations both under the operating system and on bare metal. To find out more about Acronis Disk Director Lite, see the Disk management (p. 193) section.

2.9 Tape support

Acronis Backup & Recovery 10 supports tape libraries, autoloaders, SCSI and USB tape drives as storage devices. A tape device can be locally attached to a managed machine (in this case, the Acronis Backup & Recovery 10 Agent writes and reads the tapes) or accessed through the Acronis Backup & Recovery 10 Storage Node. Storage nodes ensure fully automatic operation of tape libraries and autoloaders.

Backup archives created using different ways of access to tape have different formats. A tape written by a storage node cannot be read by an agent.

Linux-based and PE-based bootable media allow for backup and recovery using both local access and access through the storage node. Backups created using the bootable media can be recovered with the Acronis Backup & Recovery 10 Agent running in the operating system.

2.9.1 Tape compatibility table

The following table summarizes the readability of tapes written by Acronis True Image Echo and Acronis True Image 9.1 product families in Acronis Backup & Recovery 10. The table also illustrates the compatibility of tapes written by various components of Acronis Backup & Recovery 10.

			...is readable on a tape device attached to a machine with...			
			ABR10 Bootable Media	ABR10 Agent for Windows	ABR10 Agent for Linux	ABR10 Storage Node
Tape written on a locally attached tape	Bootable Media	ATIE 9.1	+	+	+	+
		ATIE 9.5	+	+	+	+
		ATIE 9.7	+	+	+	+
		ABR10	+	+	+	+

device (tape drive or tape library) by...	Agent for Windows	ATIE 9.1	+	+	+	+
		ATIE 9.5	-	-	-	+
		ATIE 9.7	-	-	-	+
		ABR10	+	+	+	+
	Agent for Linux	ATIE 9.1	+	+	+	+
		ATIE 9.5	+	+	+	+
		ATIE 9.7	+	+	+	+
		ABR10	+	+	+	+
Tape written on a tape device through...	Backup Server	ATIE 9.1	+	+	+	+
		ATIE 9.5	-	-	-	+
		ATIE 9.7	-	-	-	+
	Storage Node	ABR10	-	-	-	+

2.9.2 Using a single tape drive

A tape drive that is locally attached to a managed machine can be used by local backup plans as a storage device. The functionality of a locally attached autoloader or tape library is limited to the ordinary tape drive. This means that the program can only work with the currently mounted tape and you have to mount tapes manually.

Backup to a locally attached tape device

When creating a backup plan, you are able to select the locally attached tape device as the backup destination. An archive name is not needed when backing up to a tape.

An archive can span multiple tapes but can contain only one full backup and an unlimited number of incremental backups. Every time you create a full backup, you start with a new tape and create a new archive. As soon as the tape is full, a dialog window with a request to insert a new tape will appear.

The content of a non-empty tape will be overwritten on prompt. You have an option to disable prompts, see Additional settings (p. 69).

Workaround

In case you want to keep more than one archive on the tape, for example, back up volume C and volume D separately, choose incremental backup mode instead of a full backup when you create an initial backup of the second volume. In other situations, incremental backup is used for appending changes to the previously created archive.

You might experience short pauses that are required to rewind the tape. Low-quality or old tape, as well as dirt on the magnetic head, might lead to pauses that can last up to several minutes.

Limitations

1. Multiple full backups within one archive are not supported.
2. Individual files cannot be recovered from a disk backup.
3. Backups cannot be deleted from a tape either manually or automatically during cleanup. Retention rules and backup schemes that use automatic cleanup (GFS, Tower of Hanoi) are disabled in the GUI when backing up to a locally attached tape.
4. Personal vaults cannot be created on tape devices.
5. Because the presence of an operating system cannot be detected in a backup located on a tape, Acronis Universal Restore (p. 240) is proposed at every disk or volume recovery, even when recovering a Linux or non-system Windows volume.

6. Acronis Active Restore (p. 228) is not available when recovering from a tape.

Recovery from a locally attached tape device

Before creating a recovery task, insert or mount the tape containing the backup you need to recover. When creating a recovery task, select the tape device from the list of available locations and then select the backup. After recovery is started, you will be prompted for other tapes if the tapes are needed for recovery.

2.10 Support for SNMP

SNMP objects

Acronis Backup & Recovery 10 provides the following Simple Network Management Protocol (SNMP) objects to SNMP management applications:

- Type of event
Object identifier (OID): 1.3.6.1.4.1.24769.100.200.1.0
Syntax: OctetString
The value may be "Information", "Warning", "Error" and "Unknown". "Unknown" is sent only in the test message.
- Text description of the event
Object identifier (OID): 1.3.6.1.4.1.24769.100.200.2.0
Syntax: OctetString
The value contains the text description of the event (it looks identical to messages published by Acronis Backup & Recovery 10 in its log).

Example of varbind values:

1.3.6.1.4.1.24769.100.200.1.0:Information

1.3.6.1.4.1.24769.100.200.2.0:I0064000B

Supported operations

Acronis Backup & Recovery 10 **supports only TRAP operations**. It is not possible to manage Acronis Backup & Recovery 10 using GET- and SET- requests. This means that you need to use an SNMP Trap receiver to receive TRAP-messages.

About the management information base (MIB)

The MIB file **acronis-abr.mib** is located in the Acronis Backup & Recovery 10 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"

2.11 Proprietary Acronis technologies

This section describes the proprietary technologies inherited by Acronis Backup & Recovery 10 from Acronis True Image Echo and Acronis True Image 9.1 product families.

2.11.1 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.

Certain Windows applications, such as Acronis disk management tools, can access the zone.

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.
- Being an internal archive storage, 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 dual destination (p. 67) backup.

Limitations

- The zone cannot be organized on a dynamic disk or a disk using the GPT partitioning style.

Managing the Acronis Secure Zone

Acronis Secure Zone is considered as a personal vault (p. 241). Once created on a managed machine, the zone is always present in the list of **Personal vaults**. Centralized backup plans (p. 232) can use Acronis Secure Zone as well as local plans (p. 237).

If you have used Acronis Secure Zone before, please note a radical change in the zone functionality. The zone does not perform automatic cleanup, that is, deleting old archives, anymore. Use backup schemes with automatic cleanup to back up to the zone, or delete outdated backups manually using the archive management functionality.

With the new Acronis Secure Zone behavior, you obtain the ability to:

- list archives located in the zone and backups contained in each archive
- examine a backup's content

- mount a disk backup to copy files from the backup to a physical disk
- safely delete archives and backups from the archives.

For more information about operations available in Acronis Secure Zone, see the Personal vaults (p. 81) section.

Upgrade from Acronis True Image Echo

When upgrading from Acronis True Image Echo to Acronis Backup & Recovery 10, Acronis Secure Zone will keep the archives created with Echo. The zone will appear in the list of personal vaults and the old archives will be available for recovery.

2.11.2 Acronis Startup Recovery Manager

A modification of the bootable agent (p. 231) can be placed on a system disk and configured to start at boot time when F11 is pressed. This eliminates the need for rescue media or network connection to start the bootable rescue utility. This feature has the trade name "Acronis Startup Recovery Manager".

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. The user can also back up using Acronis Startup Recovery Manager, while on the move.

On machines with the GRUB boot loader installed, the user selects the Acronis Startup Recovery Manager from the boot menu instead of pressing F11.

Activation and deactivation of the Acronis Startup Recovery Manager

The operation that enables using Acronis Startup Recovery Manager is called "activation". To activate Acronis Startup Recovery Manager, select **Actions > Activate Acronis Startup Recovery Manager** from the program menu.

You can activate or deactivate the Acronis Startup Recovery Manager at any time from the **Tools** menu. The deactivation will disable the boot time prompt "Press F11 for Acronis Startup Recovery Manager..." (or removes the corresponding entry from GRUB's boot menu). This means you will need bootable media in case the system fails to boot.

Limitation

Acronis Startup Recovery Manager requires re-activation of third-party loaders after activation.

Upgrade from Acronis True Image Echo

After upgrade from Acronis True Image Echo to Acronis Backup & Recovery 10, Acronis Startup Recovery Manager appears as deactivated regardless of its status before the upgrade. You can activate Acronis Startup Recovery Manager again at any time.

2.11.3 Universal Restore (Acronis Backup & Recovery 10 Universal Restore)

Acronis Backup & Recovery 10 Universal Restore is the Acronis proprietary technology that helps recover and boot up Windows on dissimilar hardware or a virtual machine. The Universal Restore

handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

Acronis Backup & Recovery 10 Universal Restore purpose

A system can be easily recovered from a disk backup (image) onto the same system or to identical hardware. However, if you change a motherboard or use another processor version—a likely possibility in case of hardware failure—the recovered system could be unbootable. An attempt to transfer the system to a new, much more powerful computer will usually produce the same unbootable result because the new hardware is incompatible with the most critical drivers included in the image.

Using Microsoft System Preparation Tool (Sysprep) does not solve this problem, because Sysprep permits installing drivers only for Plug and Play devices (sound cards, network adapters, video cards etc.). As for system Hardware Abstraction Layer (HAL) and mass storage device drivers, they must be identical on the source and the target computers (see Microsoft Knowledge Base, articles 302577 and 216915).

The Universal Restore technology provides an efficient solution for hardware-independent system recovery by replacing the crucial Hardware Abstraction Layer (HAL) and mass storage device drivers.

Universal Restore is applicable for:

1. Instant recovery of a failed system on different hardware.
2. Hardware-independent cloning and deployment of operating systems.
3. Physical-to-physical, physical-to-virtual and virtual-to-physical machine migration.

The Universal Restore principles

1. Automatic HAL and mass storage driver selection.

Universal Restore searches for drivers in the network folders you specify, on removable media and in the default driver storage folders of the system being recovered. Universal Restore analyzes the compatibility level of all found drivers and installs HAL and mass storage drivers that better fit the target hardware. Drivers for network adapters are also searched and passed to the operating system which installs them automatically when first started.

*The Windows default driver storage folder is determined in the registry value **DevicePath**, which can be found in the registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion**. This storage folder is usually **WINDOWS\inf**.*

2. Manual selection of the mass storage device driver.

If the target hardware has a specific mass storage controller (such as a SCSI, RAID, or Fibre Channel adapter) for the hard disk, you can install the appropriate driver manually, bypassing the automatic driver search-and-install procedure.

3. Installing drivers for Plug and Play devices.

Universal Restore relies on the built-in Plug and Play discovery and configuration process to handle hardware differences in devices that are not critical for the system start, such as video, audio and USB. Windows takes control over this process during the logon phase, and if some of the new hardware is not detected, you will have a chance to install drivers for it later manually.

Universal Restore and Microsoft Sysprep

Universal Restore is not a system preparation tool. You can apply it to any Windows image created by Acronis products, including images of systems prepared with Microsoft System Preparation Tool (Sysprep). The following is an example of using both tools on the same system.

Universal Restore does not strip the security identifier (SID) and user profile settings in order to run the system immediately after recovery without re-joining the domain or re-mapping network user profiles. If you are going to change the above settings on a recovered system, you can prepare the system with Sysprep, image it and recover, if need be, using the Universal Restore.

Limitations

Universal Restore is not available:

- when a computer is booted with Acronis Startup Recovery Manager (using F11) or
- the backup image is located in the Acronis Secure Zone or
- when using Acronis Active Restore,

because these features are primarily meant for instant data recovery on the same machine.

Universal Restore is not available when recovering Linux.

Getting Universal Restore

Universal Restore comes free with Acronis Backup & Recovery 10 Advanced Server SBS Edition and Acronis Backup & Recovery 10 Advanced Server Virtual Edition.

Universal Restore for the other product editions is purchased separately, has its own license, and is installed as a separate feature from the setup file. You need to re-create bootable media to make the newly installed add-on operational in the bootable environment.

2.11.4 Acronis Active Restore

Active Restore is the Acronis proprietary technology that brings a system online immediately after the system recovery is started.

Customers familiar with Acronis Recovery for Microsoft Exchange can note that this product uses Active Restore to achieve immediate availability of an Exchange information store after starting the recovery. While based on the same technology, recovery of the Information Store proceeds in quite a different way than the operating system recovery described in this section.

Supported operating systems

Acronis Active Restore is available when recovering Windows starting from Windows 2000.

Limitation

The only supported archive location is a local drive, or more precisely, any device available through the machine's BIOS. This may be Acronis Secure Zone, a USB hard drive, a flash drive or any internal hard drive.

How it works

When configuring a recovery operation, you select disks or volumes to recover from a backup. Acronis Backup & Recovery 10 scans the selected disks or volumes in the backup. If this scan finds a supported operating system, the Acronis Active Restore option becomes available.

If you do not enable the option, the system recovery will proceed in the usual way and the machine will become operational after the recovery is completed.

If you enable the option, the sequence of actions will be set as follows.

Once the system recovery is started, the operating system boots from the backup. The machine becomes operational and ready to provide necessary services. The data required to serve incoming requests is recovered with the highest priority; everything else is recovered in the background.

Because serving requests is performed simultaneously with recovery, the system operation can slow down even if recovery priority in the recovery options is set to **Low**. This way, the system downtime is reduced to a minimum at the cost of a temporary performance downgrade.

Usage scenarios

1. The system uptime is one of the efficiency criteria.

Examples: Client-oriented online services, Web-retailers, polling stations.

2. The system/storage space ratio is heavily biased toward storage.

Some machines are being used as storage facilities, where the operating system claims a small space segment and all other disk space is committed to storage, such as movies, sounds or other multimedia files. Some of these storage volumes can be extremely large as compared to the system and so practically all the recovery time will be dedicated to recovering the files, which might be used much later on, if in any near future at all.

If you opt for Acronis Active Restore, the system will be operational in a short time. Users will be able to open the necessary files from the storage and use them while the rest of the files, which are not immediately necessary, are being recovered in the background.

Examples: movie collection storage, music collection storage, multimedia storage.

How to use

1. Back up the system disk or volume to a location accessible through the system's BIOS. This may be Acronis Secure Zone, a USB hard drive, a flash drive or any internal hard drive.

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

2. Create bootable media.
3. If a system failure occurs, boot the machine using the bootable media. Start the console and connect to the bootable agent.
4. Configure the system recovery: select the system disk or volume and select the **Use Acronis Active Restore** check box.

Acronis Active Restore will choose for the boot-up and subsequent recovery the first operating system found during the backup scan. Do not try to recover more than one operating system using Active Restore if you want the result to be predictable. When recovering a multi-boot system, choose only one system volume and boot volume at a time.

5. Once the system recovery is started, the operating system boots from the backup. The Acronis Active Restore icon appears in the system tray. The machine becomes operational and ready to provide necessary services. The immediate user sees the drive tree and icons and can open files or launch applications even though they were not yet recovered.

The Acronis Active Restore drivers intercept system queries and set the immediate priority for recovery of the files that are necessary to serve the incoming requests. While this on-the-fly recovery proceeds, the continuing recovery process is transferred to the background.

Please do not shut down or reboot the machine until the recovery is completed. If you switch off the machine, all the changes made to the system since the last boot up would be lost. The system will not be recovered, not even partially. The only possible solution in this case will be to restart the recovery process from a bootable media.

6. The background recovery continues until all the selected volumes are recovered, the log entry is made and the Acronis Active Restore icon disappears from the system tray.

3 Options

This section covers Acronis Backup & Recovery 10 options that can be configured using Graphical User Interface. The content of this section is applicable to both stand-alone and advanced editions of Acronis Backup & Recovery 10.

3.1 Console options

The console options define the way information is represented in the Graphical User Interface of Acronis Backup & Recovery 10.

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

3.1.1 Startup page

This option defines whether to show the **Welcome** screen or the **Dashboard** upon connection of the console to a managed machine or to the management server.

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

To make a selection, select or clear the check box for **Show the Dashboard view upon connection of the console to a machine**.

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

3.1.2 Pop-up messages

About tasks that need interaction

This option is effective when the console is connected to a managed machine or to the management server.

The option defines whether to display the pop-up window when one or more tasks 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 tasks in the same place. Until at least one task requires interaction, you can open this window at any time from the managed machine's **Dashboard**. Alternatively, you can review the task execution states in the **Tasks** view and specify your decision on each task in the **Information** pane.

The preset is: **Enabled**.

To make a selection, select or clear the **Pop up the "Tasks Need Interaction" window** check box.

About the task execution results

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

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

3.1.3 Time-based alerts

Last backup

This option is effective when the console is connected to a managed machine (p. 237) or to the management server (p. 238).

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** section of the **Dashboard**. When the console is connected to the management server, this setting will also control the color scheme of the **Last backup** column's value for each machine.

Last connection

This option is effective when the console is connected to the management server or to a registered machine (p. 239).

The option defines whether to alert if no connection was established between a registered machine and the management server for a period of time so indicating that the machine might not be centrally managed (for instance in the case of network connection failure to that machine). You can configure the length of time that is considered critical.

The preset is: alert if the machine's last connection to the management server was more than **5 days** ago.

The alert is displayed in the **Alerts** section of the **Dashboard**. When the console is connected to the management server, this setting will also control the color scheme of the **Last connect** column's value for each machine.

3.1.4 Number of tasks

This option is effective only when the console is connected to the management server.

The option defines how many tasks will be displayed at a time in the **Tasks** view. You can also use filters available in the **Tasks** view to limit the number of displayed tasks.

The preset is: **400**. The adjustment range is: **20 to 500**.

To make a selection, choose the desired value from the **Number of tasks** drop-down menu.

3.1.5 Fonts

This option is effective when the console is connected to a managed machine or to the management server.

The option defines the fonts to be used in the Graphical User Interface of Acronis Backup & Recovery 10. The **Menu** setting affects the drop-down and context menus. The **Application** setting affects the 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 the button to the right.

3.2 Machine options

The machine options define the general behavior of all Acronis Backup & Recovery 10 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.

3.2.1 Event tracing

It is possible to duplicate log events generated by the agent(s), operating on the managed machine, in the Application Event Log of Windows; or send the events to the specified SNMP managers. If you do not modify the event tracing options anywhere except for here, your settings will be effective for each local backup plan and each task created on the machine.

You can override the settings set here, exclusively for the events occurred during backup or during recovery (see Default backup and recovery options (p. 50).) 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.

3.2.1.1 Windows event log

This option is effective only in Windows operating systems.

This option is not available when operating under the bootable media.

This option defines whether the agent(s) operating on the managed machine have to log events in the Application Event Log of Windows (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

You can override the settings set here, exclusively for the events that occur during backup or during recovery, in the Default backup and recovery options (p. 50). In this case, the settings set here will be effective for operations other than backup and recovery, such as archive validation or cleanup.

You can further override the settings set in the default backup and recovery options, when creating a backup plan or a recovery task. The settings you obtain in this case will be plan-specific or task-specific.

The preset is: **Disabled**.

To enable this option, select the **Log events** check box.

Use the **Types of events to log** check box to filter the events to be logged in the Application Event Log of Windows:

- **All events** - all events (information, warnings and errors)
- **Errors and warnings**
- **Errors only.**

To disable this option, clear the **Log events** check box.

3.2.1.2 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 (p. 50). In this case, the settings set here will be effective for operations other than backup and recovery, such as archive validation or cleanup.

You can further override the settings set in the default backup and recovery options, when creating a backup plan or a recovery task. The settings you obtain in this case will be plan-specific or task-specific.

For detailed information about using SNMP with Acronis Backup & Recovery 10, please see "Support for SNMP (p. 38)".

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. 49).

3.2.1.3 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.

3.2.2 Log cleanup rules

This option specifies how to clean up the Acronis Backup & Recovery 10 agent log.

This option defines the maximum size of the agent log folder (in Windows XP/2003 Server, %ALLUSERSPROFILE%\Application Data\Acronis\BackupAndRecovery\MMS\LogEvents).

The preset is: **Maximum log size: 1 GB. On cleanup, keep 95% of the maximum log size.**

When the option is enabled, the program compares the actual log size with the maximum size after every 100 log entries. Once the maximum log size is exceeded, the program deletes the oldest log entries. You can select the amount of log entries to retain. The default 95% setting will keep most of the log. With the minimum 1% setting, the log will be nearly cleared.

This parameter can also be set by using Acronis Administrative Template.

3.2.3 Online backup proxy

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

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

Note: Acronis Backup & Recovery 10 Online supports only HTTP and HTTPS proxy servers.

To set up proxy server settings

1. Select the **Use a proxy server** check box.
2. In **Address**, specify the network name or IP address of the proxy server—for example: **proxy.example.com** or **192.168.0.1**

3. In **Port**, specify the port number of the proxy server—for example: **80**
4. If the proxy server requires authentication, specify the credentials in **User name** and **Password**.
5. To test the proxy server settings, click **Test connection**.

If you do not know the proxy server settings, contact your network administrator or Internet service provider for assistance.

Alternatively, you can try to take these settings from your Web browser's configuration. This is how to find them in three popular browsers.

- **Microsoft Internet Explorer.** On the **Tools** menu, click **Internet Options**. On the **Connections** tab, click **LAN settings**.
- **Mozilla Firefox.** On the **Tools** menu, click **Options** and then click **Advanced**. On the **Network** tab, under **Connection**, click **Settings**.
- **Google Chrome.** In **Options**, click **Under the Hood**. Under **Network**, click **Change proxy settings**.

3.2.4 Customer Experience Program

This option defines whether the machine will participate in the Acronis Customer Experience Program (ACEP).

If you choose **Yes, I want to participate in the ACEP**, information about the hardware configuration, the most and least used features and about any problems will be automatically collected from the machine and sent to Acronis on a regular basis. The end results are intended to provide software improvements and enhanced functionality to better meet the needs of Acronis customers.

Acronis does not collect any personal data. To learn more about the ACEP, read the terms of participation on the Acronis Web site or in the product GUI.

Initially the option is configured during the Acronis Backup & Recovery 10 agent installation. This setting can be changed at any time using the product GUI (**Options > Machine options > Customer Experience Program**). The option can also be configured using the Group Policy infrastructure. A setting defined by a Group Policy cannot be changed using the product GUI unless the Group Policy is disabled on the machine.

3.3 Default backup and recovery options

3.3.1 Default backup options

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

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

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

Availability of the backup options

The set of available backup options depends on:

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

The following table summarizes the availability of the backup options.

	Agent for Windows		Bootable media (Linux-based or PE-based)	
	Disk backup	File backup	Disk backup	File backup
Archive protection (p. 52) (password + encryption)	+	+	+	+
Source files exclusion (p. 53)	+	+	+	+
Pre/Post backup commands (p. 54)	+	+	PE only	PE only
Pre/Post data capture commands (p. 56)	+	+	-	-
Multi-volume snapshot (p. 58)	+	+	-	-
File-level backup snapshot (p. 58)	-	+	-	-
Use VSS (p. 59)	+	+	-	-
Compression level (p. 59)	+	+	+	+
Backup performance:				
Backup priority (p. 60)	+	+	-	-
HDD writing speed (p. 60)	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD
Network connection speed (p. 61)	Dest: network share	Dest: network share	Dest: network share	Dest: network share
Fast incremental/differential backup (p. 64)	+	-	+	-
Backup splitting (p. 64)	+	+	+	+
File-level security (p. 65):				
Preserve files' security settings in archives	-	+	-	-
In archives, store encrypted files in decrypted state	-	+	-	-

Media components	Dest: removable media	Dest: removable media	-	-
Error handling (p. 66):				
Do not show messages and dialogs while processing (silent mode)	+	+	+	+
Re-attempt if an error occurs	+	+	+	+
Ignore bad sectors	+	+	+	+
Dual destination (p. 67)	Dest: local	Dest: local	-	-
Task start conditions (p. 68)	+	+	-	-
Task failure handling (p. 69)	+	+	-	-
Additional settings (p. 69):				
Overwrite data on a tape without prompting for user confirmation	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape
Dismount media after backup has finished	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media
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	-	-	+	+
Use FTP in Active mode	Dest: FTP server	Dest: FTP server	Dest: FTP server	Dest: FTP server
Notifications:				
E-mail (p. 61)	+	+	-	-
Win Pop-up (p. 62)	+	+	-	-
Event tracing:				
Windows events log (p. 63)	+	+	-	-
SNMP (p. 63)	+	+	-	-

3.3.1.1 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.

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.

3.3.1.2 Source files exclusion

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

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

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

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

To specify which files and folders to exclude:

Set up any of the following parameters:

- **Exclude all hidden files and folders**

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

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

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

- **Exclude files matching the following criteria**

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

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

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

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

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

Exclusion examples

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

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

3.3.1.3 Pre/Post commands

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

The option enables you to define the commands to be automatically executed before and after the backup procedure.

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

Pre-backup command	Backup	Post-backup command
--------------------	--------	---------------------

Examples of how you can use the pre/post commands:

- Delete some temporary files from the disk before starting backup
- Configure a third-party antivirus product to be started each time before the backup starts
- Copy an archive to another location after the backup ends.

The program does not support interactive commands, i.e. commands that require user input (for example, "pause").

To specify pre/post commands

1. Enable pre/post commands execution by checking the following options:
 - **Execute before the backup**
 - **Execute after the backup**
2. Do any of the following:
 - Click **Edit** to specify a new command or a batch file
 - Select the existing command or the batch file from the drop-down list
3. Click **OK**.

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.

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. If successful execution of the command is critical for your backup strategy, select the **Fail the task if the command execution fails** check box. In case the command execution fails, the program will remove the resulting TIB file and temporary files if possible, and the task will fail.
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 on the **Dashboard**.
5. Click **Test Command** to check if the command is correct.

3.3.1.4 Pre/Post data capture commands

This option is effective for both Windows and Linux operating systems.

The option enables you to define the commands to be automatically executed before and after data capture (that is, taking the data snapshot). Data capture is performed by Acronis Backup & Recovery 10 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 (p. 59) 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. 54), 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**.

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.

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	N/A	Continue the

	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.	backup after the command is executed despite command execution failure or success.		backup concurrently with the command execution and irrespective of the command execution result.
--	--	--	--	--

3.3.1.5 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.

3.3.1.6 Multi-volume snapshot

This option is effective only for Windows operating systems.

This option applies to disk-level backup. This option also applies to file-level backup when the file-level backup is performed by taking a snapshot. (The File-level backup snapshot (p. 58) option determines whether a snapshot will be taken during file-level backup).

The option determines whether to take snapshots of multiple volumes at the same time or one by one.

The preset is: **Enable.**

When this option is set to **Enable**, snapshots of all volumes being backed up will be created simultaneously. Use this option to create a time-consistent backup of data spanned across multiple volumes, for instance for an Oracle database.

When this option is set to **Disable**, the volumes' snapshots will be taken one after the other. As a result, if the data spans across several volumes, the resulting backup may be not consistent.

3.3.1.7 Volume Shadow Copy Service

This option is effective only for Windows operating systems.

The option defines whether a Volume Shadow Copy Service (VSS) provider—either Acronis VSS or Microsoft VSS—has to notify VSS-aware applications that the backup is about to start. This ensures the consistent state of all data used by the applications, in particular, completion of all database transactions, at the moment of taking the data snapshot by Acronis Backup & Recovery 10. Data consistency, in turn, ensures that the application will be recovered in the correct state and become operational immediately after recovery.

The preset is: **Create snapshots using VSS**

Acronis Backup & Recovery 10 will select the VSS provider automatically based on the operating system running on the machine and whether the machine is a member of an Active Directory domain.

Create snapshots without using VSS

Choose this option if your database is incompatible with VSS. The data snapshot will be taken by Acronis Backup & Recovery 10. Backup process is fastest, but data consistency of the applications whose transactions are not completed at the time of taking a snapshot cannot be guaranteed. You may use Pre/Post data capture commands (p. 56) to indicate which commands should be performed before and after taking the snapshot, to ensure that the data is being backed up in a consistent state. For instance, specify pre-data capture commands that will suspend the database and flush all caches to ensure that all transactions are completed; and specify post-data capture commands that will resume the database operations after the snapshot is taken.

Volume shadow copy writers

Before backing up the data of VSS-aware applications, make sure that the volume shadow copy writers for those applications are turned on, by examining the list of writers that are present in the operating system. To view this list, run the following command:

```
vssadmin list writers
```

Note: In Microsoft Windows Small Business Server 2003, the writer for Microsoft Exchange Server 2003 is turned off by default. For instructions on how to turn it on, see the corresponding Microsoft Help and Support article (<http://support.microsoft.com/kb/838183/en>).

3.3.1.8 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.

3.3.1.9 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.

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.

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.

Network connection speed

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

This option is available when a location on the network (network share, managed vault or an FTP/SFTP server) is selected as the backup destination.

The option defines the amount of network connection bandwidth allocated for transferring the backup data.

By default the speed is set to maximum, i.e. the software uses all the network bandwidth it can get when transferring the backup data. Use this option to reserve a part of the network bandwidth to other network activities.

The preset is: **Maximum**.

To set the network connection speed for backup

Do any of the following:

- Click **Transferring speed stated as a percentage of the estimated maximum speed of the network connection**, and then drag the slider or type a percentage in the box
- Click **Transferring speed stated in kilobytes per second**, and then enter the bandwidth limit for transferring backup data in kilobytes per second.

3.3.1.10 Notifications

Acronis Backup & Recovery 10 provides the ability of notifying users about backup completion through e-mail or the messaging service.

E-mail

This option is effective for Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option enables you to receive e-mail notifications about the backup task's successful completion, failure or need for interaction along with the full log of the task.

The preset is: **Disabled**.

To configure e-mail notification

1. Select the **Send e-mail notifications** check box to activate notifications.
2. In the **E-mail addresses** field, type the e-mail address to which notifications will be sent. You can enter several addresses separated by semicolons.

3. Under **Send notifications**, select the appropriate check boxes as follows:
 - **When backup completes successfully** – to send a notification when the backup task has completed successfully
 - **When backup fails** – to send a notification when the backup task has failedThe **When user interaction is required** check box is always selected.
4. For the e-mail message to include the log entries related to the backup, select the **Add full log to the notification** check box.
5. Click **Additional e-mail parameters**, to configure additional e-mail parameters as follows, then click **OK**:
 - **From** - type the e-mail address of the user from whom the message will be sent. If you leave this field empty, messages will be constructed as if they are from the destination address.
 - **Use encryption** – you can opt for encrypted connection to the mail server. SSL and TLS encryption types are available for selection.
 - Some Internet service providers require authentication on the incoming mail server before being allowed to send something. If this is your case, select the **Log on to incoming mail server** check box to enable a POP server and to set up its settings:
 - **Incoming mail server (POP)** – enter the name of the POP server.
 - **Port** – set the port of the POP server. By default, the port is set to 110.
 - **User name** – enter the user name
 - **Password** – enter the password.
 - Select the **Use the specified outgoing mail server** check box to enable an SMTP server and to set up its settings:
 - **Outgoing mail server (SMTP)** – enter the name of the SMTP server.
 - **Port** – set the port of the SMTP server. By default, the port is set to 25.
 - **User name** – enter the user name.
 - **Password** – enter the password.
6. Click **Send test e-mail message** to check if the settings are correct.

Messenger service (WinPopup)

This option is effective for Windows and Linux operating systems on the sending machine and only for Windows on the receiving machine.

This option is not available when operating under bootable media.

The option enables you to receive WinPopup notifications about the backup task's successful completion, failure or need for interaction.

The preset is: **Disabled**.

Before configuring WinPopup notifications, make sure the Messenger service is started on both the machine executing the task and the machine that will receive messages.

The Messenger service is not started by default in the Microsoft Windows Server 2003 family. Change the service Startup mode to Automatic and start the service.

To configure WinPopup notifications:

1. Select the **Send WinPopup notifications** check box.

2. In the **Machine name** field, enter the name of the machine to which notifications will be sent. Multiple names are not supported.

Under **Send notifications**, select the appropriate check boxes as follows:

- **When backup completes successfully** – to send notification when the backup operation is completed successfully
- **When backup fails** – to send notification when the backup operation is failed

The **When user interaction is required** check box – to send notification during the operation when user interaction is required – is always selected.

Click **Send test WinPopup message** to check if the settings are correct.

3.3.1.11 Event tracing

It is possible to duplicate log events of the backup operations, performed on the managed machine, in the Application Event Log of Windows; or send the events to the specified SNMP managers.

Windows event log

This option is effective only in Windows operating systems.

This option is not available when operating under the bootable media.

This option defines whether the agent(s) operating on the managed machine have to log events of the backup operations in the Application Event Log of Windows (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

The preset is: **Use the setting set in the Machine options.**

To select whether to log the backup operations events in the Application Event Log of Windows:

Choose one of the following:

- **Use the setting set in the Machine options** – to use the setting specified for the machine. For more information refer to Machine options (p. 47).
- **Log the following event types** – to log events of the backup operations in the Application Event Log. Specify the types of events to be logged:
 - **All events** – log all events (information, warnings and errors)
 - **Errors and warnings**
 - **Errors only**
- **Do not log** - to disable logging events of the backup operations in the Application Event Log.

SNMP notifications

This option is effective for both Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option defines whether the agent(s) operating on the managed machine have to send the log events of the backup operations to the specified Simple Network Management Protocol (SNMP) managers. You can choose the types of events to be sent.

For detailed information about using SNMP with Acronis Backup & Recovery 10, please see "Support for SNMP (p. 38)".

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 (p. 47).
- **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.

3.3.1.12 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.

3.3.1.13 Backup splitting

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

The option defines how a backup can be split.

The preset is: **Automatic.**

The following settings are available.

Automatic

With this setting, Acronis Backup & Recovery 10 will act as follows.

- **When backing up to a hard disk:**

A single backup file will be created if the destination disk's file system allows the estimated file size.

The backup will automatically be split into several files if the destination disk's file system does not allow the estimated file size. Such might be the case when the backup is placed on FAT16 and FAT32 file systems that have a 4GB file size limit.

If the destination disk runs out of free space while creating the backup, the task enters the **Need interaction** state. You have the ability to free additional space and retry the operation. If you do so, the resulting backup will be split into the parts created before and after the retry.

- **When backing up to removable media** (CD, DVD or a tape device locally attached to the managed machine):

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

Fixed size

Enter the desired file size or select it from the drop-down list. The backup will then be split into multiple files of the specified size. This comes in handy when creating a backup that you plan to burn to multiple CDs or DVDs later on. You might also want to split the backup destined to an FTP server, since data recovery directly from an FTP server requires the backup to be split into files no more than 2GB in size.

3.3.1.14 File-level security

These options are effective only for file-level backup in Windows operating systems.

In archives, store encrypted files in a decrypted state

This option defines whether to decrypt files before saving them to a backup archive.

The preset is: **Disabled**.

Simply ignore this option if you do not use the encryption. Enable the option if encrypted files are included in the backup and you want them to be accessed by any user after recovery. Otherwise, only the user who encrypted the files/folders will be able to read them. Decryption may also be useful if you are going to recover encrypted files on a different machine.

*File encryption is available in Windows using the NTFS file system with the Encrypting File System (EFS). To access a file or folder encryption setting, select **Properties > General > Advanced Attributes > Encrypt contents to secure data**.*

Preserve file security settings in archives

This option defines whether to back up NTFS permissions for files along with the files.

The preset is: **Enabled**.

When the option is enabled, files and folders are saved in the archive with the original permissions to read, write or execute the files for each user or user group. If you recover a secured file/folder on a machine without the user account specified in the permissions, you may not be able to read or modify this file.

To completely eliminate this kind of problem, disable preserving file security settings in archives. The recovered files and folders will always inherit the permissions from the folder to which they are recovered or from the disk, if recovered to the root.

Alternatively, you can disable recovery (p. 74) of the security settings, even if they are available in the archive. The result will be the same - the files will inherit the permissions from the parent folder.

To access file or folder NTFS permissions, select **Properties > Security**.

3.3.1.15 Media components

This option is effective for both Windows and Linux operating systems, when the backup destination is removable media.

When backing up to removable media, you can make this media work as regular Linux-based bootable media (p. 231) by writing additional components to it. As a result, you will not need a separate rescue disc.

The preset is: **None selected**.

Select the check boxes for the components you want to put on the bootable media:

- **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**, all data will be immediately silently recovered to its initial location.

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.

- **Bootable agent** is a bootable rescue utility (based on Linux kernel) that includes most of the functionality of the Acronis Backup & Recovery 10 agent. Put this component on the media if you want more functionality during recovery. You will be able to configure the recovery operation in the same way as under regular bootable media; use Active Restore or Universal Restore. If the media is being created in Windows, the disk management functionality will also be available.

3.3.1.16 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: 5**. **Interval between attempts: 30 seconds**.

When a recoverable error occurs, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts is performed, depending on which comes first.

For example, if the backup destination on the network becomes unavailable or not reachable, the program will attempt to reach the destination every 30 seconds, but no more than 5 times. The attempts will be stopped as soon as the connection is resumed OR the specified number of attempts is performed, depending on which comes first.

Ignore bad sectors

The preset is: **Disabled**.

When the option is disabled, the program will display a pop-up window each time it comes across a bad sector and ask for a user decision as to whether to continue or stop the backup procedure. In order to back up the valid information on a rapidly dying disk, enable ignoring bad sectors. The rest of the data will be backed up and you will be able to mount the resulting disk backup and extract valid files to another disk.

3.3.1.17 Dual destination

This option is effective for both Windows and Linux operating systems, when the primary backup destination is a *local folder or Acronis Secure Zone* and the secondary destination is *another local folder or network share*. Managed vaults and FTP servers are not supported as secondary destinations.

The preset is: **Disabled**.

When dual destination is enabled, the agent will automatically copy each backup being created locally to the secondary destination such as a network share. Once the backup to the primary destination is completed, the agent compares the updated archive contents to the secondary archive contents, and copies to the secondary destination all backups that are missing there along with the new backup.

This option enables quick machine backup to the internal drive as an intermediate step before saving the ready backup on the network. This comes in handy in cases of slow or busy networks and time-consuming backup procedures. Disconnection during the copy transfer will not affect the backup operation as opposed to backing up directly to the remote location.

Other advantages:

- Replication enhances the archive reliability.
- Roaming users can back up their portable computers to Acronis Secure Zone while on the road. When the portable computer is connected to the corporate network, all changes made to the archive will be transferred to its stationary copy after the first backup operation.

If you select the password-protected Acronis Secure Zone as the primary destination, keep in mind that the archive in the secondary destination will not be protected with a password.

To use Dual destination:

1. Select the check box for **Use dual destination**.
2. Browse to the secondary destination or enter the full path to the destination manually.
3. Click **OK**.

You might have to provide the access credentials for the secondary destination. Enter the credentials on prompt.

3.3.1.18 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. 87) and Conditions (p. 97).

The preset is: **Wait until the conditions are met.**

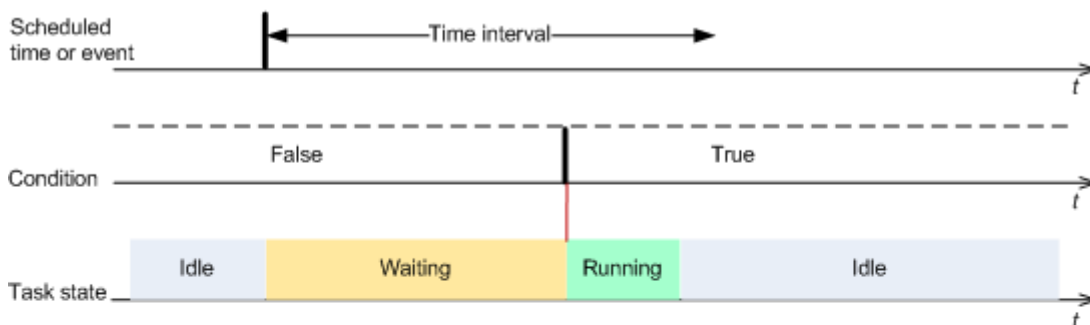
Wait until the conditions are met

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.

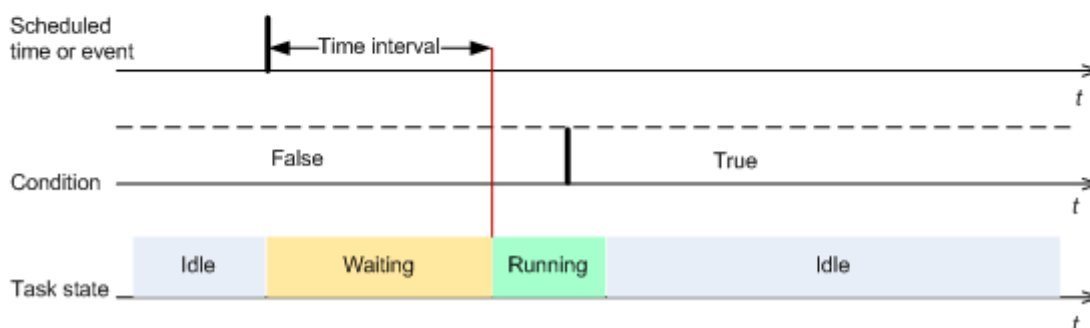
To handle the situation when the conditions are not met for too long and further delaying the backup is becoming risky, you can set the time interval after which the task will run irrespective of the condition. Select the **Run the task anyway after** check box and specify the time interval. The task 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.

3.3.1.19 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.

3.3.1.20 Additional settings

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

Overwrite data on a tape without prompting for user confirmation

This option is effective only when backing up to a tape device.

The preset is: **Disabled**.

When starting backup to a non-empty tape in a locally attached tape device, the program will warn that you are about to lose data on the tape. To disable this warning, select this check box.

Dismount media after backup has finished

This option is effective in Windows and Linux operating systems.

This option is effective when backing up to a removable media (CD, DVD, tape or floppy disk.)

The preset is: **Disabled**.

The destination CD/DVD can be ejected or the tape can be dismounted after the backup is completed.

Ask for the first media while backing up to removable media

This option is effective only when backing up to removable media.

The option defines whether to display the **Insert First Media** prompt when backing up to removable media.

The preset is: **Enabled**.

When the option is enabled, backing up to removable media may be not possible if the user is away, because the program will wait for someone to press OK in the prompt box. Hence, you should disable the prompt when scheduling a backup to removable media. Then, if the removable media is available (for example, a DVD is inserted), the task can run unattended.

Reset archive bit

The option is effective only for file-level backup in Windows operating systems and in bootable media.

The preset is: **Disabled**.

In Windows operating systems, each file has the **File is ready for archiving** attribute, available by selecting **File -> Properties -> General -> Advanced -> Archive and Index attributes**. This attribute, also known as the archive bit, is set by the operating system each time the file is changed and can be reset by backup applications each time they include the file in a backup. The archive bit value is used by various applications such as databases.

When the **Reset archive bit** check box is selected, Acronis Backup & Recovery 10 will reset the archive bits of all files being backed up. Acronis Backup & Recovery 10 itself does not use the archive bit value. When performing incremental or differential backup, it determines whether a file has changed by the file size and the date/time when the file was last saved.

Restart the machine automatically after backup is finished

This option is available only when operating under bootable media.

The preset is: **Disabled**.

When the option is enabled, Acronis Backup & Recovery 10 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.

Deduplicate backup only after transferring it to the vault (do not deduplicate at source)

This option is available only in advanced editions of Acronis Backup & Recovery 10.

This option is effective for Windows and Linux operating systems and bootable media, when the backup destination is a deduplicating vault.

The preset is: **Disabled**.

Enabling this option turns off deduplicating backups at source, meaning that deduplication will be performed by Acronis Backup & Recovery 10 Storage Node after the backup is saved to the vault (this is called deduplication at target).

Turning off deduplication at source may lead to faster backup processes but greater network traffic and heavier load of the storage node. The eventual size of the backup in the vault is independent of whether deduplication at source is turned on.

Deduplication at source and deduplication at target are described in [Deduplication overview](#).

Save software RAID and LVM metadata along with backups

This option is effective only for disk-level backups of machines running Linux.

The preset is: **Enabled**.

When this option is enabled, Acronis Backup & Recovery 10 will save information about the structure of logical volumes (known as LVM volumes) and of Linux Software RAID devices (known as MD devices) to the **/etc/Acronis** directory before creating the backup.

When recovering MD devices and LVM volumes under bootable media, you can use this information to automatically recreate the volume structure. For instructions, see [Recovering MD devices and logical volumes](#) (p. 189).

When using this option, make sure that the volume containing the **/etc/Acronis** directory is among the volumes to back up.

Use FTP in Active mode

The preset is: **Disabled**.

Enable this option if the FTP server supports active mode and you want this mode to be used for file transfers.

3.3.2 Default recovery options

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

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

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

Availability of the recovery options

The set of available recovery options depends on:

- The environment the agent operates in (Windows, bootable media)
- The type of data being recovered (disk, file)
- The operating system being recovered from the disk backup

The following table summarizes the availability of the recovery options.

	Agent for Windows		Bootable media (Linux-based or PE-based)	
	Disk recovery	File recovery (also from a disk backup)	Disk recovery	File recovery (also from a disk backup)
Pre/Post recovery commands (p. 73)	+	+	PE only	PE only
Recovery priority (p. 74)	+	+	-	-
File-level security (p. 74):				
Recover files with their security settings	-	+	-	+
Error handling (p. 77):				
Do not show messages and dialogs while processing (silent mode)	+	+	+	+
Re-attempt if an error occurs	+	+	+	+
Additional settings (p. 78):				
Set current date and time for recovered files	-	+	-	+
Validate backup archive before recovery	+	+	+	+
Check file system after recovery	+	-	+	-
Reboot machine automatically if it is required for recovery	+	+	-	-
Change SID after recovery	Windows recovery	-	Windows recovery	-
Notifications:				
E-mail (p. 75)	+	+	-	-
Win Pop-up (p. 76)	+	+	-	-
Event tracing:				
Windows events log (p. 76)	+	+	-	-
SNMP (p. 77)	+	+	-	-

3.3.2.1 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**.

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.	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

	Fail the task if the command execution failed.			command execution result.
--	--	--	--	---------------------------

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. If successful execution of the command is critical for you, select the **Fail the task if the command execution fails** check box. In case 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 or the errors and warnings displayed on the **Dashboard**.

5. Click **Test command** to check if the command is correct.

A post-recovery command will not be executed if the recovery proceeds with reboot.

3.3.2.2 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.

3.3.2.3 File-level security

This option is effective only for recovery from file-level backup of Windows files.

This option defines whether to recover NTFS permissions for files along with the files.

The preset is: **Recover files with their security settings**.

If the file NTFS permissions were preserved during backup (p. 65), you can choose whether to recover the permissions or let the files inherit the NTFS permissions from the folder to which they are recovered.

3.3.2.4 Notifications

Acronis Backup & Recovery 10 provides the ability of notifying users about recovery completion through e-mail or the messaging service.

E-mail

This option is effective for Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option enables you to receive e-mail notifications about the recovery task's successful completion, failure or need for interaction along with the full log of the task.

The preset is: **Disabled**.

To configure e-mail notification

1. Select the **Send e-mail notifications** check box to activate notifications.
2. In the **E-mail addresses** field, type the e-mail address to which notifications will be sent. You can enter several addresses separated by semicolons.
3. Under **Send notifications**, select the appropriate check boxes as follows:
 - **When backup completes successfully** – to send a notification when the backup task has completed successfully
 - **When backup fails** – to send a notification when the backup task has failedThe **When user interaction is required** check box is always selected.
4. For the e-mail message to include the log entries related to the backup, select the **Add full log to the notification** check box.
5. Click **Additional e-mail parameters**, to configure additional e-mail parameters as follows, then click **OK**:
 - **From** - type the e-mail address of the user from whom the message will be sent. If you leave this field empty, messages will be constructed as if they are from the destination address.
 - **Use encryption** – you can opt for encrypted connection to the mail server. SSL and TLS encryption types are available for selection.
 - Some Internet service providers require authentication on the incoming mail server before being allowed to send something. If this is your case, select the **Log on to incoming mail server** check box to enable a POP server and to set up its settings:
 - **Incoming mail server (POP)** – enter the name of the POP server.
 - **Port** – set the port of the POP server. By default, the port is set to 110.
 - **User name** – enter the user name
 - **Password** – enter the password.
 - Select the **Use the specified outgoing mail server** check box to enable an SMTP server and to set up its settings:
 - **Outgoing mail server (SMTP)** – enter the name of the SMTP server.
 - **Port** – set the port of the SMTP server. By default, the port is set to 25.

- **User name** – enter the user name.
- **Password** – enter the password.

Click **Send test e-mail message** to check if the settings are correct.

Messenger service (WinPopup)

This option is effective for Windows and Linux operating systems.

This option is not available when operating under bootable media.

The option enables you to receive WinPopup notifications about the recovery task's successful completion, failure or need for interaction.

The preset is: **Disabled**.

Before configuring WinPopup notifications, make sure the Messenger service is started on both the machine executing the task and the machine that will receive messages.

The Messenger service is not started by default in the Microsoft Windows Server 2003 family. Change the service Startup mode to Automatic and start the service.

To configure WinPopup notifications:

1. Select the **Send WinPopup notifications** check box.
2. In the **Machine name** field, enter the name of the machine to which notifications will be sent. Multiple names are not supported.
3. Under **Send notifications**, select the appropriate check boxes as follows:
 - **When recovery completes successfully** – to send notification when the recovery task has completed successfully
 - **When recovery fails** – to send notification when the recovery task has failed.
 - **When user interaction is required** – to send notification during the operation when user interaction is required.
4. Click **Send Test WinPopup Message** to check if the settings are correct.

3.3.2.5 Event tracing

It is possible to duplicate log events of the recovery operations, performed on the managed machine, in the Application Event Log of Windows; or send the events to the specified SNMP managers.

Windows event log

This option is effective only in Windows operating systems.

This option is not available when operating under the bootable media.

This option defines whether the agent(s) operating on the managed machine have to log events of the recovery operations in the Application Event Log of Windows (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**). You can filter the events to be logged.

The preset is: **Use the setting set in the Machine options**.

To select whether to log the recovery operations events in the Application Event Log of Windows:

Select one of the following:

- **Use the setting set in the Machine options** – to use the setting specified for the machine. For more information refer to Machine options (p. 47).
- **Log the following event types** – to log events of the recovery operations in the Application Event Log. Specify the types of events to be logged:
 - **All events** – log all events (information, warnings and errors)
 - **Errors and warnings**
 - **Errors only**
- **Do not log** - to disable logging events of the recovery operations in the Application Event Log.

SNMP notifications

This option is effective for both Windows and Linux operating systems.

This option is not available when operating under the bootable media.

The option defines whether the agent(s) operating on the managed machine have to send the log events of the recovery operations to the specified Simple Network Management Protocol (SNMP) managers. You can choose the types of events to be sent.

For detailed information about using SNMP with Acronis Backup & Recovery 10, please see "Support for SNMP (p. 38)".

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 (p. 47).
- **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.

3.3.2.6 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: 5**. **Interval between attempts: 30 seconds**.

When a recoverable error occurs, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts is performed, depending on which comes first.

For example, if the network location becomes unavailable or not reachable, the program will attempt to reach the location every 30 seconds, but no more than 5 times. The attempts will be stopped as soon as the connection is resumed OR the specified number of attempts is performed, depending on which comes first.

3.3.2.7 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 backup 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.

When operating under bootable media, this option is not effective for the NTFS file system.

The preset is **Disabled**.

This option defines whether to check the integrity of the file system after a disk or volume recovery.

Restart machine automatically if it is required for recovery

This option is effective when recovery takes place on a machine running an operating system.

The preset is **Disabled**.

The option defines whether to reboot the machine automatically if it is required for recovery. Such might be the case when a volume locked by the operating system has to be recovered.

Reboot machine after recovery

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.

Change SID after the recovery is finished

This option is not effective when recovery to a virtual machine is performed by Acronis Backup & Recovery 10 Agent for ESX/ESXi or Acronis Backup & Recovery 10 Agent for Hyper-V.

The preset is **Disabled**.

Acronis Backup & Recovery 10 can generate a unique security identifier (SID) for the recovered system. You do not need a new SID when recovering a system over itself or when creating a system replica that will replace the original system. Generate a new SID if the original and the recovered systems will work concurrently in the same workgroup or domain.

Use FTP in Active mode

The preset is: **Disabled**.

Enable this option if the FTP server supports active mode and you want this mode to be used for file transfers.

4 Vaults

A vault is a location for storing backup archives. For ease of use and administration, a vault is associated with the archives' metadata. Referring to this metadata makes for fast and convenient operations with archives and backups stored in the vault.

A vault can be organized on a local or networked drive, detachable media or a tape device attached to the Acronis Backup & Recovery 10 Storage Node.

There are no settings for limiting a vault size or number of backups in a vault. You can limit the size of each archive using cleanup, but the total size of archives stored in the vault is limited by the storage size only.

Why create vaults?

We recommend that you create a vault in each destination where you are going to store backup archives. This will ease your work as follows.

Quick access to the vault

You will not have to remember paths to the folders where the archives are stored. When creating a backup plan or a task that requires selection of an archive or an archive destination place, the list of vaults will be available for quick access without drilling down through the folders tree.

Easy archive management


A vault is available for access from the **Navigation** pane. Having selected the vault, you can browse the archives stored there and perform the following archive management operations:


- get a list of backups included in each archive
- recover data from a backup
- examine backup content
- validate all archives in the vault or individual archives or backups
- mount a volume backup to copy files from the backup to a physical disk
- safely delete archives and backups from the archives.

Creating vaults is highly recommended but is not obligatory. You may choose not to use the shortcuts and always specify the full path to the archive vault. All of the above operations except for archive and backup deletion can be performed without creating vaults.

The operation of creating a vault results in adding the vault name to the **Vaults** section of the **Navigation** pane.

Way of working with the "Vaults" view

 **Vaults** (on the navigation pane) - top element of the vaults tree. Click this item to display groups of centralized and personal vaults.

 **Personal**. This group is available when the console is connected to a managed machine. Expand this group to display a list of personal vaults created on the managed machine.

Click any personal vault in the vaults tree to open the detailed view of this vault (p. 81) and to take actions on the vault (p. 82), archives (p. 84) and backups (p. 85) stored in there.

4.1 Personal vaults

A vault is called personal if it was created using direct connection of the console to a managed machine. Personal vaults are specific for each managed machine. Personal vaults are visible to any user that can log on to the system. A user's right to back up to a personal vault is defined by the user's permission for the folder or device where the vault is located.

A personal vault can be organized on a network share, FTP server, detachable media or removable, Acronis Online Backup Storage, tape device, or on a hard drive local to the machine. Acronis Secure Zone is considered as a personal vault available to all users that can log on the system. Personal vaults are created automatically when backing up any of the above locations.

Personal vaults can be used by local backup plans or local tasks. Centralized backup plans cannot use personal vaults except for Acronis Secure Zone.

Sharing a personal vault

Multiple machines can refer to the same physical location; for example, to the same shared folder. However, each of the machines has its own shortcut in the **Vaults** tree. Users that back up to a shared folder can see and manage each other's archives according to their access permissions for that folder. To ease archive identification, the **Personal vault** view has the **Owner** column that displays the owner of each archive. To find out more about the owner concept see Owners and credentials (p. 24).

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.

4.1.1 Working with the "Personal vault" view


This section briefly describes the main elements of the **Personal vault** view, and suggests the ways to work with them.


Vault toolbar

The toolbar contains operational buttons that let you perform operations with the selected personal vault. See the Actions on personal vaults (p. 82) section for details.

Pie chart with legend

The **pie chart** lets you estimate the vault's load: it shows the proportion of the vault's free space and occupied space.

 - free space: 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 free space of the appropriate volume.

 - occupied space: total size of backup archives and their metadata, if it is located in the vault. Other files that may be put to this folder by a user, are not counted.

The **legend** displays the following information about the vault:

- full path to the vault

- total number of archives and backups stored in the vault
- the ratio of the occupied space to the original data size.

Vault content

The **Vault content** section contains the archives table and toolbar. The archives table displays archives and backups that are stored in the vault. Use the archives toolbar to perform actions on the selected archives and backups. The list of backups is expanded by clicking the "plus" sign to the left of the archive's name. All the archives are grouped by type on the following tabs:

- The **Disk archives** tab lists all the archives that contain disk or volume backups (images).
- The **File archives** tab lists all the archives that contain file backups.

Related sections:

Operations with archives stored in a vault (p. 84)

Operations with backups (p. 85)

Filtering and sorting archives (p. 86)

Bars of the "Actions and tools" pane

- **[Vault Name]** The **Actions** bar is available when clicking the vault in the vaults tree. Duplicates actions of the vault's toolbar.
- **[Archive Name]** The **Actions** bar is available when you select an archive in the archives table. Duplicates actions of the archives toolbar.
- **[Backup Name]** The **Actions** bar is available when you expand the archive and click on any of its backups. Duplicates actions of the archives toolbar.




4.1.2 Actions on personal vaults






To access actions

1. Connect the console to the management server.
2. In the **Navigation** pane, click **Vaults > Personal**.

All the operations described here are performed by clicking the corresponding buttons on the vaults toolbar. These operations can be also accessed from the **[Vault name] actions** item of the main menu.

The following is a guideline for you to perform operations with personal vaults.

To	Do
Create a personal vault	Click  Create . The procedure of creating personal vaults is described in-depth in the Creating a personal vault (p. 83) 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. 174) 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. 160) page, where this vault is already pre-selected as a source. The vault validation checks all the archives stored in the vault.
Delete a vault	Click  Delete . 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.

4.1.2.1 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. In the **Path** field, click **Change...**
In the opened **Personal Vault Path** window, specify a path to the folder that will be used as the vault. A personal vault can be organized on detachable or removable media, on a network share, or on FTP.
4. Click **OK**. As a result, the created vault appears in the **Personal** group of the vaults tree.

4.1.2.2 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 temporarily disable (p. 111) schedules of the given plans.
2. Move the vault folder with all its archives 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 temporarily disable (p. 111) schedules of the given plans.
2. Move the archives 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*.





4.2 Common operations

4.2.1 Operations with archives stored in a vault

To perform any operation with an archive, you have to select it first. If the archive is protected with a password, you will be asked to provide it.

All the operations described below are performed by clicking the corresponding buttons on the toolbar. These operations can be also accessed from the **[Archive name] actions** bar (on the **Actions and tools** pane) and from the **[Archive name] actions** item of the main menu respectively.

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. 160) 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. 167) page will be opened with the pre-selected archive as a source.</p> <p>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 the archive or one of the archives you want to delete.2. Click  Delete. <p>The program duplicates your selection in the Backups deletion (p. 86) 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 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>

4.2.2 Operations with backups

To perform any operation with a backup, you have to select it first. To select a backup, expand the archive, then click the backup. If the archive is protected with a password, you will be asked to provide it.

All the operations described below are performed by clicking the corresponding buttons on the toolbar. These operations can be also accessed from the '**[Backup name]**' **actions** bar (on the **Actions and tools** pane) and 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 page will be opened with the pre-selected backup as a source.
Recover a disk/volume as a virtual machine	Right-click the disk backup, then select Recover as virtual machine . The Recover data page will be opened with the pre-selected backup as a source. Select the location and the type of new virtual machine and then proceed as with regular disk or volume recovery.
Validate a backup	Click  Validate . The Validation (p. 160) 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. 167) 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.
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. 86) 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. 86) window that has check boxes for each archive and each backup. Review the selection and correct if need be, then confirm the deletion.

4.2.3 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.

The filters in this window are from the archives list of the vault view. Thus, if some filters have been applied to the archives list, only the archives and backups corresponding to these filters are displayed here. To see all content, clean all the filter fields.

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.

4.2.4 Filtering and sorting archives

The following is a guideline for you to filter and sort archives in the archives table.

To	Do
Sort backup archives by any column	Click the column's header to sort the archives in ascending order. Click it once again to sort the archives in descending order.
Filter archives by name, owner, or machine.	In the field below the corresponding column's header, type the archive name (the owner name, or the machine name). As a result, you will see the list of the archives, whose names (owner names, or machine names) fully or just partly coincide with the entered value.

Configuring the archives table

By default, the table has seven columns that are displayed, others are hidden. If required, you can hide the displayed columns and show hidden ones.

To show or hide columns

1. Right-click any column header to open the context menu. The menu items that are ticked off correspond to the column headers presented in the table.
2. Click the items you want to be displayed/hidden.

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

Schedule

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

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

Condition

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

Condition: run the task only if
User is idle (a screen saver is running or the machine is locked)
Location's host is available
The task run time is within the specified time interval

All users are logged off
The specified period of time has passed since the last successful backup completed

The scheduler behavior, in case the event occurs but the condition (or any of multiple conditions) is not met is defined by the Task start conditions (p. 68) 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.

5.1 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.

Advanced scheduling settings are available only for machines registered on Acronis Backup & Recovery 10 Management Server. To specify these settings, click **Change** in the **Advanced settings** area.

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

Examples

"Simple" daily schedule

Run the task every day at 6PM.

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

1. Every: **1** day(s).
2. Once at: **06:00:00 PM**.
3. Effective:
From: **not set**. The task will be started on the current day, if it has been created before 6PM. If you have created the task after 6 PM, the task will be started for the first time on the next day at 6 PM.
To: **not set**. The task will be performed for an indefinite number of days.

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

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

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

1. Every: **1** day(s).
2. Every: **3** hours
From: **12:00:00 AM** (midnight) Until: **09:00:00 PM** - thus, the task will be performed 8 times a day with a 3 hour time interval. After the last daily recurrence at 9 PM, the next day comes and the task starts over again from midnight.
3. Effective:
From: **09/15/2009**. If September 15, 2009 is the current date of the task's creation and, say, 01:15 PM is the task's creation time, the task will be started when the nearest time interval comes: at 03:00 PM in our example.
To: **12/15/2009**. On this date the task will be performed for the last time, but the task itself is still available in the **Tasks** view.

Several daily schedules for one task

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

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

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

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

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

First daily schedule

1. Every: **3** day(s).
2. Every: **4** hours.
From: **08:00:00 AM** Until: **12:00:00 PM**.
3. Effective:
From: **09/20/2009**.
To: **not set**.

Second daily schedule

1. Every: **3** day(s).
2. Every: **2** hour(s).
From: **03:00:00 PM** Until: **07:00:00 PM**.
3. Effective:
From: **09/20/2009**.
To: **not set**.

5.2 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.

Advanced scheduling settings are available only for machines registered on Acronis Backup & Recovery 10 Management Server. To specify these settings, click **Change** in the **Advanced settings** area.

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

Examples

"One day in the week" schedule

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

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

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

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

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

This schedule is widely used when creating a custom backup scheme. The "One day in the week"-like schedule is added to the full backups, while the incremental backups are scheduled to be performed on workdays. For more details, see the Full and incremental backups plus cleanup example in the Custom backup scheme (p. 135) 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. 135) section.

Several weekly schedules for one task

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

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

- Monday: twice at 12 PM (noon) and 9 PM
- Tuesday: every 3 hours from 9 AM till 9 PM
- Wednesday: every 3 hours from 9 AM till 9 PM

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

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

First schedule

1. Every: **1** week(s) on: **Mon, Fri**.
2. Every: **9** hours
From: **12:00:00 PM** Until: **09:00:00 PM**.
3. Effective:
From: **not set**.
To: **not set**.

Second schedule

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

Third schedule

1. Every: **1** week(s) on: **Sat, Sun**.
2. Once at: **09:00:00 PM**.
3. Effective:
From: **not set**.
To: **not set**.

5.3 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: <...>	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.
From: <...> Until: <...>	

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

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

Advanced scheduling settings are available only for machines registered on Acronis Backup & Recovery 10 Management Server. To specify these settings, click **Change** in the **Advanced settings** area.

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

Examples

"Last day of every month" schedule

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

The schedule's parameters are set up as follows.

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

This schedule is widely used when creating a custom backup scheme. The "Last day of every month" schedule is added to the full backups, while the differential backups are scheduled to be performed once a week and incremental on workdays. For more details, see the Monthly full, weekly differential, and daily incremental backups plus cleanup example in the Custom backup scheme (p. 135) section.

"Season" schedule

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

The schedule's parameters are set up as follows.

1. Months: **September, October, November**.
2. On: **<all> <workdays>**.
3. Every: **6** hours.
 - From: **12:00:00 AM** Until: **06:00:00 PM**.
4. Effective:
 - From: **08/30/2009**. Actually the task will be started on the first workday of September. By setting up this date we just define that the task must be started in 2009.

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

Several monthly schedules for one task

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

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

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

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

First schedule

1. Months: **December, January, February**.
2. On: **<All> <All workdays>**
3. Once at: **10:00:00 PM**.
4. Effective:
From: **11/01/2009**.
To: **not set**.

Second schedule

1. Months: **March, April, May, September, October, November**.
2. On: **<All> <All workdays>**.
3. Every: **12 hours**
From: **12:00:00 AM** Until: **12:00:00 PM**.
4. Effective:
From: **11/01/2009**.
To: **not set**.

Third schedule

1. Months: **June, July, August**.
2. Days: **1, 15**.
3. Once at: **10:00:00 PM**.
4. Effective:
From: **11/01/2009**.
To: **not set**.

5.4 At Windows Event Log event

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

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

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

Parameters

Log name

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

Event source

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

Event type

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

Event ID

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

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

Examples

"Bad block" emergency backup

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

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

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

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

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

Pre-update backup in Vista

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

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

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

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

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

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

How to view events in Event viewer

To open a log in Event Viewer

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

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

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

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

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

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

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

5.5 Upon an Acronis Drive Monitor alert

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

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

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

Alerts about disk temperature do not start the backup.

Tip: If your backup plan uses the custom backup scheme (p. 135), you can set up this emergency backup simply by adding an extra schedule to the same backup plan. When using a different backup scheme, you will need to create a separate backup plan.

5.6 Conditions

Conditions add more flexibility to the scheduler, enabling to execute backup tasks with respect to certain conditions. Once a specified event occurs (see the "Scheduling (p. 87)" section for the list of available events), the scheduler checks the specified condition and executes the task if the condition is met.

The scheduler behavior in case the event occurs but the condition (or any of multiple conditions) is not met, is defined by the **Task start conditions** (p. 68) backup option. There, you can specify how important the conditions are for the backup strategy:

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

Conditions are available only when the custom backup scheme (p. 135) is used. You can set conditions for full, incremental and differential backup separately.

Adding multiple conditions

Multiple conditions must be met simultaneously to enable task execution.

Example:

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

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

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

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

5.6.1 User is idle

Applies to: Windows

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

Example:

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

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

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

As a result,

- (1) If the user becomes idle before 9PM, the backup task will start at 9PM.
- (2) If the user becomes idle between 9PM and 11PM, the backup task will start immediately after the user becomes idle.
- (3) If the user is still active at 11PM, the backup task starts anyway.

5.6.2 Location's host is available

Applies to: Windows, Linux

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

Example:

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

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

As a result,

- (1) If 9:00 PM comes and the location's host is available, the backup task starts right on time.
- (2) If 9:00 PM comes but the host is unavailable at the moment, the backup task will start on the next workday if the location's host is available.
- (3) If the location's host will never be available on workdays at 9:00 PM, the task never starts.

5.6.3 Fits time interval

Applies to: Windows, Linux

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

Example

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

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

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

As a result,

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

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

What if...

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

For example:

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

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

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

5.6.4 User logged off

Applies to: Windows

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

Example

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

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

As a result,

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

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

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

5.6.5 Time since last backup

Applies to: Windows, Linux

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

Example:

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

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

As a result,

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

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

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

6 Direct management

This section covers operations that can be performed directly on a managed machine by using the direct console-agent connection. The content of this section is applicable to both stand-alone and advanced editions of Acronis Backup & Recovery 10.

6.1 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.

6.1.1 Dashboard





Use the Dashboard to estimate at a glance whether the data is successfully protected on the machine. The dashboard shows the summary of Acronis Backup & Recovery 10 agent's activities and enables you to rapidly identify and resolve any issues.







Alerts

The alerts section draws your attention to issues that have occurred on the machine and offers you ways of fixing or examining them. The most critical issues are displayed on the top. If there are no alerts or warnings at the moment, the system displays "No alerts or warnings".

Types of alerts

The table below illustrates the types of messages you may observe.

	Description	Offer	Comment
	Failed tasks: X	Resolve	Resolve will open the Backup plans and Tasks view with failed tasks, where you can examine the reason of failure.
	Tasks that need interaction: X	Resolve	Each time a task needs human interaction, the Dashboard shows a message to inform you what action has to be performed (for example, insert new CD or Stop/Retry/Ignore on an error).
	Failed to check the license for the current edition. X day(s) remaining until the software stops working. Please make sure you have a valid license on Acronis License Server.	Connect	Acronis Backup & Recovery 10 agent connects to Acronis License Server at the start and then every 1–5 days (the default is 1 day), as specified by the agent configuration parameters. If the license check does not succeed for 1–60 days, as specified by the agent configuration parameters (the default is 30 days), the agent will stop working until there has been a successful last license check.
	Cannot check the license key for the current edition for X days. Either Acronis License Server was unavailable, or the license key data was corrupted. Check connectivity to the server and run Acronis License Server to manage	Connect	Acronis Backup & Recovery 10 is stopped. For the past X days, the agent was unable to check whether its license is available on Acronis License Server. This is probably due to the license server being unavailable. You may also want to ensure that the licenses are present on the license server, or that the license key data was not corrupted.

	licenses. Please make sure you have a valid license on Acronis License Server.		After a successful license check the agent will start working.
	Trial version of product expires in X day(s) Please make sure you have a valid license on Acronis License Server.	Connect	Once the trial version of the product is installed, the program starts the countdown of days remaining until the trial period expires.
	Trial period is over. Start the installer and enter a full license key. Please make sure you have a valid license on Acronis License Server.	Connect	15 day trial period has expired. Enter a full license key.
	Vaults with low free space: X	View vaults	View vaults will take you to the Vaults view where you can examine the vault size, free space, content and take the necessary steps to increase the free space.
	Bootable media was not created	Create now	To be able to recover an operating system when the machine fails to boot, you must: 1. Back up the system volume (and the boot volume, if it is different) 2. Create at least one bootable media (p. 231). Create now will launch the Bootable Media Builder (p. 238).
	No backups have been created for X days	Back up now	The Dashboard warns you that no data was backed up on the machine for a relatively long period of time. Back up now will take you to Create a Backup Plan page where you can instantly configure and run the backup operation. To configure the time interval that is considered as critical, select Options > Console options > Time-based alerts .
	Not connected to management server for X days	View the machines	This type of message can appear on a machine that is registered on a management server. The Dashboard warns you that the connection might be lost or the server might be unavailable and the machine is not centrally managed as a result.

Activities

The calendar lets you explore the history of the Acronis Backup & Recovery 10 agent's activities on the machine. Right-click on any highlighted date and select **View log** to see the list of log entries filtered by date.

On the **View** section (at the right of the calendar), you can select the activities to highlight depending on the presence and severity of the errors.

	How it is determined
Errors	Highlight the date in red if at least one "Error" entry appeared in the log on this date.
Warnings	Highlight the date in yellow if no "Error" entries appeared and at least one "Warning" entry appeared in the log on this date.
Information	Highlight the date in green if only "Information" log entries appeared on this date (normal activity.)

The **Select current date** link focuses selection to the current date.

System view

Shows summarized statistics of backup plans, tasks, and brief information on the last backup. Click the items in this section to obtain the relevant information. This will take you to the **Backup plans and tasks** (p. 103) view with pre-filtered plans or tasks. For instance, if you click **Local** under **Backup plans**, the **Backup plans and tasks** view will be opened with backup plans filtered by the **Local** origin.

6.1.1.1 Tasks need interaction

This window accumulates all the tasks that require user interaction in one place. It enables you to specify your decision, such as to confirm reboot or to retry after freeing-up the disk space, on each of the tasks. Until at least one task requires interaction, you can open this window at any time from the managed machine's **Dashboard** (p. 101).

If you select the check box for the **Do not show this window when tasks require interaction. I will see this information in the tasks' details and dashboard.** parameter, the tasks will be displayed on the **Dashboard** among other alerts and warnings.

Alternatively, you can review the task execution states in the **Backup plans and tasks** (p. 103) view and specify your decision on each task in the **Information** panel (or in the **Task details** (p. 111) window).

6.1.2 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.

A backup plan is a set of rules that specify how the given data will be protected on a given machine. Physically, a backup plan is a bundle of tasks configured for execution on a managed machine. To find out what a backup plan is currently doing on the machine, check the backup plan execution state (p. 104). A backup plan state is a cumulative state of the plan's tasks. The status of a backup plan (p. 104) helps you to estimate whether the data is successfully protected.

A task is a set of sequential actions to be performed on a machine when a certain time comes or certain event occurs. To keep track of a task's current progress, examine its state (p. 105). Check a task status (p. 106) to ascertain the result of a task.

Way of working

- 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. See the Filtering and sorting backup plans and tasks (p. 110) section for details.

- In the backup table, select the backup plan (task).
- Use the toolbar's buttons to take an action on the selected plan (task). See the Actions on backup plans and tasks (p. 107) section for details. You can run, edit, stop and delete the created plans and tasks.
- Use the **Information** panel to review detailed information on the selected plan (task). The panel is collapsed by default. To expand the panel, click the  chevron. The content of the panel is also duplicated in the **Plan details** (p. 113) and **Task details** (p. 111) windows respectively.

6.1.2.1 Understanding states and statuses

Backup plan execution states

A backup plan can be in one of the following execution states: **Idle**; **Waiting**; **Running**; **Stopping**; **Need Interaction**.

Plan states names are the same as task state names because a plan state is a cumulative state of the plan's tasks.

	State	How it is determined	How to handle
1	Need interaction	At least one task needs user interaction. Otherwise, see 2.	Identify the tasks that need interaction (the program will display what action is needed) -> Stop the tasks or enable the tasks to run (change media; provide additional space on the vault; ignore the read error; create the missing Acronis Secure Zone).
2	Running	At least one task is running. 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 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 this way 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	Stopping	At least one task is stopping. Otherwise, see 5.	No action is required.
5	Idle	All the tasks are idle.	No action is required.

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.

	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 in case a local plan has failed ▪ Edit the backup policy on the management server in case a centralized plan has failed When creating a backup plan or policy the administrator can turn on the option to stop executing the backup plan as soon as the backup plan gets the Error status. The backup plan's execution can be resumed using the Restart button.
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 in case none of the tasks has been started yet or some of the tasks are stopped or being stopped. These situations are considered as normal.

Task states

A task can be in one of the following states: **Idle**; **Waiting**; **Running**; **Stopping**; **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 or recovery task cannot run simultaneously on a machine. A backup task and a recovery task also cannot run simultaneously. 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. 68) for details.

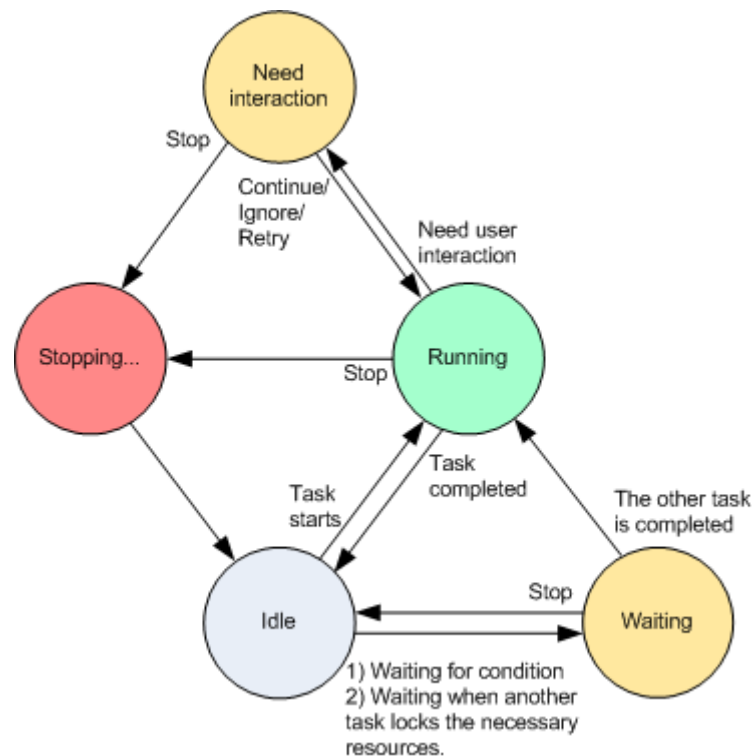
Need interaction

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 **Stopping** (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.)

Stopping

The user can stop a running task or a task that needs interaction. The task changes to the **Stopping** state and then to the **Idle** state. A waiting task can also be stopped. In this case, since the task is not running, "stop" means removing it from the queue .

Task state diagram



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 manuallyEdit the failed task to prevent its future failureEdit the local plan to prevent its future failure in case a local plan has failedEdit the backup policy on the management server in case a





			centralized plan has failed
2	Warning	Last result is "Succeeded with warning"	View the log to read the warnings -> [optionally] Perform actions to prevent the future warnings or failure.
3	OK	Last result is "Succeeded", "-", or "Stopped"	No action is required. The "-" state means that the task has never been started or has been started, but has not finished yet and so its result is not available.



6.1.2.2 Working with backup plans and tasks




Actions on backup plans and tasks

The following is a guideline for you to perform operations with backup plans and tasks.

To	Do
Create a new backup plan, or a task	Click  New , then select one of the following: <ul style="list-style-type: none"> Backup plan (p. 116) Recovery task Validation task (p. 160)
View details of a plan/task	<u>Backup plan</u> Click  View details . In the Plan Details (p. 113) window, review the plan details. <u>Task</u> Click  View details . In the Task Details (p. 111) window, review the task details.
View plan's/task's log	<u>Backup plan</u> Click  View log . You will be taken to the Log (p. 114) view containing the list of the plan-related log entries. <u>Task</u> Click  View log . You will be taken to the Log (p. 114) view containing the list of the task-related log entries.

Run a plan/task	<p><u>Backup plan</u></p> <p>Click  Run.</p> <p>In the Run Backup Plan (p. 111) window, select the task you need to be run.</p> <p>Running the backup plan starts the selected task of that plan immediately in spite of its schedule and conditions.</p> <p><i>Why can't I run the backup plan?</i></p> <ul style="list-style-type: none"> Do not have the appropriate privilege <p>Without the Administrator privileges on the machine, a user cannot run plans owned by other users.</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><u>Backup plan</u></p> <p>Click  Stop.</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>Click  Stop.</p> <p><i>What will happen if I stop the task?</i></p> <p>Generally, stopping the task aborts its operation (backup, recovery, validation, exporting, conversion, migration). The task enters the Stopping state first, then becomes Idle. The task schedule, if created, remains valid. To complete the operation you will have to run the task over again.</p> <ul style="list-style-type: none"> recovery task (from the disk backup): The target volume will be deleted and its space unallocated – you will get the same result if the recovery is unsuccessful. To recover the "lost" volume, you will have to run the task once again. recovery task (from the file backup): The aborted operation may cause changes in the destination folder. Some files may be recovered, but some not, depending on the moment when you stopped the task. To recover all the files, you will have to run the task once again.

Edit a plan/task	<p><u>Backup plan</u></p> <p>Click  Edit.</p> <p>Backup plan editing is performed in the same way as creation (p. 116), 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> <p><i>Why can't I edit the backup plan?</i></p> <ul style="list-style-type: none"> ■ The backup plan is currently running. Editing of the currently running backup plan is impossible. ■ Do not have the appropriate privilege Without the Administrator privileges on the machine, a user cannot edit plans owned by other users. ■ The backup plan has a centralized origin. Direct editing of centralized backup plans is not possible. You need to edit the original backup policy. <p><u>Task</u></p> <p>Click  Edit.</p> <p><i>Why can't I edit the task?</i></p> <ul style="list-style-type: none"> ■ Task belongs to a backup plan Only tasks that do not belong to a backup plan, such as a recovery task, can be modified by direct editing. When you need to modify a task belonging to a local backup plan, edit the backup plan. A task belonging to a centralized backup plan can be modified by editing the centralized policy that spawned the plan. Only the management server administrator can do so. ■ Do not have the appropriate privilege Without the Administrator privileges on the machine, a user cannot modify tasks owned by other users.
------------------	--

Delete a plan/task	<p><u>Backup plan</u></p> <p>Click  Delete.</p> <p><i>What will happen if I delete the backup plan?</i></p> <p>The plan's deletion deletes all its tasks.</p> <p><i>Why can't I delete the backup plan?</i></p> <ul style="list-style-type: none"> ▪ The backup plan is in the "Running" state A backup plan cannot be deleted, if at least one of its tasks is running. ▪ Do not have the appropriate privilege Without the Administrator's privileges on the machine, a user cannot delete plans owned by other users. ▪ The backup plan has a centralized origin. A centralized plan can be deleted by the management server administrator by revoking the backup policy that produced the plan. <p><u>Task</u></p> <p>Click  Delete.</p> <p><i>Why can't I delete the task?</i></p> <ul style="list-style-type: none"> ▪ Task belongs to a backup plan A task belonging to a backup plan cannot be deleted separately from the plan. Edit the plan to remove the task or delete the entire plan. ▪ Do not have the appropriate privilege Without the Administrator privileges on the machine, a user cannot delete tasks owned by other users.
Refresh table	<p>Click  Refresh.</p> <p>The management console will update the list of backup plans and tasks existing on the machine with the most recent information. Though the list is refreshed automatically based on events, the data may not be retrieved immediately from the managed machine, due to some latency. Manual refresh guarantees that the most recent data is displayed.</p>

Filtering and sorting backup plans and tasks

To	Do
Sort backup plans and tasks by: name, state, status, type, origin, etc.	Click the column's header to sort the backup plans and tasks in ascending order. Click it once again to sort the plans and tasks in descending order.
Filter plans/tasks by name or owner.	Type a plan's/task's name or an owner's name in the field below the corresponding header name. As a result you will see the list of tasks, whose names/owners' names fully or just partly coincide with the entered value.
Filter plans and tasks by state, status, type, origin, last result, schedule.	In the field below the corresponding header, select the required value from the list.

Configuring backup plans and the tasks table

By default, the table has six columns that are displayed, others are hidden. If required, you can hide the displayed columns and show hidden ones.

To show or hide columns

1. Right-click any column header to open the context menu. The menu items that are ticked off correspond to the column headers presented in the table.
2. Click the items you want to be displayed/hidden.

Run backup plan

The backup plan is considered as running if at least one of its tasks is running. The **Run backup plan** window lets you run the task of the selected backup plan manually, in spite of its schedule.

To run a task of the selected backup plan


1. Select the task of the backup plan you need to run. To make certain of your selection, check the task information gathered in tabs at the bottom of the window. This information is also duplicated in the **Task details** (p. 111) window.
2. Click **OK**.

Temporarily disabling a backup plan

Temporarily disabling a backup plan is needed when moving archives from one vault to another by means of the third-party file manager.

Applies to backup plans that use custom backup schemes only.

To disable a backup plan

1. Click  **Edit**.
2. Enter the backup scheme scheduling option and disable the schedule for the desired period by changing the **Start date** and/or **End date** parameters.

Task details

The **Task details** window (also duplicated on the **Information** panel) aggregates all information on the selected task.

When a task 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 backup plan.

Types of tasks

The following table summarizes all types of tasks that exist in Acronis Backup & Recovery 10. The actual types of tasks you might observe depend on the product edition and the product component the console is connected to.

Task name	Description
Backup (disk)	Backing up disks and volumes
Backup (file)	Backing up files and folders
Backup (virtual machine)	Backing up an entire virtual machine or its volumes

Recovery (disk)	Disk backup recovery
Recovery (file)	File and folder recovery
Recovery (volume)	Recovery of volumes from a disk backup
Recovery (MBR)	Master boot record recovery
Recovery (disk to existing VM)	Recovery of a disk/volume backup to an existing virtual machine
Recovery (disk to new VM)	Recovery of a disk/volume backup to a new virtual machine
Recovery (existing VM)	Recovery of a virtual machine backup to an existing virtual machine
Recovery (new VM)	Recovery of a virtual machine backup to a new virtual machine
Validation (archive)	Validation of a single archive
Validation (backup)	Validation of backups
Validation (vault)	Validation of all archives stored in a vault
Cleanup	Deleting backups from a backup archive in accordance with retention rules
ASZ creation	Creating Acronis Secure Zone
ASZ management	Resizing, changing password, deleting Acronis Secure Zone
Disk management	Disk management operations
Compacting	Service task performed on a storage node
Indexing	Deduplication task performed by the storage node in the vault after a backup is completed

Depending on the type of task and whether it is running or not, a combination of the following tabs will appear:

Task

The **Task** tab is common for all types of tasks. It provides general information on the selected task.

Archive

The **Archive** tab is available for backup, archive validation and cleanup tasks.

Provides information on the archive: its name, type, size, where it is stored, etc.

Backup

The **Backup** tab is available for recovery, backup validation, and export tasks.

Provides details on the selected backup: when it was created, its type (full, incremental, differential), information on the archive and the vault the backup is stored in.

Settings

The **Settings** tab displays information on scheduling and the options changed against the default values.

Progress

The **Progress** tab is available while the task is running. It is common for all types of tasks. The tab provides information about task progress, elapsed time and other parameters.

Backup plan details

The **Backup plan details** window (also duplicated on the **Information** panel) aggregates in four tabs all the information on the selected backup plan.

The respective message will appear at the top of the tabs, if one of the plan's tasks requires user interaction. It contains a brief description of the problem and action buttons that let you select the appropriate action or stop the plan.

Backup plan

The **Backup plan** tab provides the following general information on the selected plan:

- **Name** - name of the backup plan
- **Origin** - whether the plan was created on the managed machine using direct management (local origin), or appeared on the machine as a result of deploying a backup policy from the management server (centralized origin).
- **Policy** (for backup plans with centralized origin) - name of the backup policy, whose deployment created the backup plan.
- **Account** - the name of the account under which the plan runs
- **Owner** - the name of the user who created or last modified the plan
- **State** - execution state (p. 104) of the backup plan.
- **Status** - status (p. 104) of the backup plan.
- **Schedule** - whether the task is scheduled, or set to start manually.
- **Last backup** - how much time has passed since the last backup.
- **Creation** - backup plan creation date.
- **Comments** - description of the plan (if provided).

Source

The **Source** tab provides the following information on the data selected for backup:

- **Source type** - the type of data (p. 119) selected for backing up.
- **Items to back up** - items selected to back up and their size.

Destination

The **Destination** tab provides the following information:

- **Location** - name of the vault or path to the folder, where the archive is stored.
- **Archive name** - name of the archive.
- **Archive comments** - comments on the archive (if provided).

Settings

The **Settings** tab displays the following information:


- **Backup scheme** - the selected backup scheme and all its settings with schedules.

- **Validation** (if selected) - events before or after which the validation is performed, and validation schedule.
- **Backup options** - backup options changed against the default values.

6.1.3 Log



The Log stores the history of operations performed by Acronis Backup & Recovery 10 on the machine, or actions a user takes on the machine using the program. For instance, when a user edits a task, the respective entry is added to the log. When the program executes a task, it adds multiple entries. With the log, you can examine operations, results of tasks' execution including reasons for failure, if any.

Way of working with log entries

- Use filters to display the desired log entries. You can also hide the unneeded columns and show the hidden ones. See the Filtering and sorting log entries (p. 115) section for details.
- In the log table, select the log entry (or log entries) to take action on it. See the Actions on log entries (p. 114) section for details.
- Use the **Information** panel to review detailed information on the selected log entry. The panel is collapsed by default. To expand the panel, click the  chevron. The content of the panel is also duplicated in the **Log entry details** (p. 116) window.

Opening the Log with pre-filtered log entries

Having selected items in other administration views (**Dashboard**, **Backup plans and tasks**), you can open the **Log** view with pre-filtered log entries for the item in question. Thus, you do not have to configure filters in the log table yourself.






View	Action
Dashboard	In the calendar, right-click on any highlighted date, and then select  View log . The Log view appears with the list of log entries already filtered by the date in question.
Backup plans and tasks	Select a backup plan or a task, and then click  View log . The Log view will display a list of the log entries related to the selected plan or task.

6.1.3.1 Actions on log entries

All the operations described below are performed by clicking the corresponding items on the log **toolbar**. All these operations can also be performed with the context menu (by right-clicking the log entry), or with the **Log actions** bar (on the **Actions and tools** pane).




The following is a guideline for you to perform actions on log entries.

To	Do
Select a single log entry	Click on it.
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 entry. All the 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"> ▪ Click  View Details. The log entry's details will be displayed in a separate window. ▪ Expand the Information panel, by clicking the chevron.
Save the selected log entries to a file	<ol style="list-style-type: none"> 1. Select a single log entry or multiple log entries. 2. Click  Save Selected to File. 3. In the opened window, specify a path and a name for the file.
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.
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. As a result, the log entries of that list will be saved.
Delete all the log entries	<p>Click  Clear Log.</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 entries and when.</p>

6.1.3.2 Filtering and sorting log entries

The following is a guideline for you to filter and sort log entries.

To	Do
Display log entries for a given time period	<ol style="list-style-type: none"> 1. In the From field, select the date starting from which to display the log entries. 2. In the To field, select the date up to which to display the log entries.
Filter log entries by type	<p>Press or release the following toolbar buttons:</p> <p> to filter error messages</p> <p> to filter warning messages</p> <p> to filter information messages</p>
Filter log entries by the original backup plan or managed entity type	<p>Under the Backup plan (or Managed entity type) column header, select the backup plan or the type of managed entity from the list.</p>
Filter log entries by task, managed entity, machine, code, owner	<p>Type the required value (task name, machine name, owner name, etc.) in the field below the respective column header.</p> <p>As a result you will see that the list of log entries fully or just partly coincide with the entered value.</p>
Sort log entries by date and time	<p>Click the column's header to sort the log entries in ascending order. Click it once again to sort the log entries in descending order.</p>

Configuring the log table

By default, the table has seven columns that are displayed, others are hidden. If required, you can hide the shown columns and show the hidden ones.

To show or hide columns

1. Right-click any column header to open the context menu. The menu items that are ticked off correspond to the column headers presented in the table.
2. Click the items you want to be displayed/hidden.

6.1.3.3 Log entry details

Displays detailed information on the log entry you have selected and lets you copy the details to the clipboard.

To copy the details, click the **Copy to clipboard** button.

Log entry data fields

A local log entry contains the following data fields:

- **Type** - type of event (Error; Warning; Information)
- **Date** - date and time of the event occurrence
- **Backup plan** - the backup plan the event relates to (if any)
- **Task** - the task the event relates to (if any)
- **Code** - the program code of the event. Every type of event in the program has its own code. A code is an integer number that may be used by Acronis support service to solve the problem.
- **Module** - 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** - user name of the backup plan owner (only under operating system)
- **Message** - a text description of the event.

The log entry's details that you copy will have the appearance as follows:

```
-----Log Entry Details-----
Type:                               Information
Date and time:                      DD.MM.YYYY HH:MM:SS
Backup plan:                        Backup plan name
Task:                               Task name
Message:                            Description of the operation
Code:                               12(3x45678A)
Module:                             Module name
Owner:                              Owner of the plan
-----
```

Date and time presentation varies depending on your locale settings.

6.2 Creating a backup plan

Before creating your first backup plan (p. 230), please familiarize yourself with the basic concepts (p. 18) used in Acronis Backup & Recovery 10.

To create a backup plan, perform the following steps.

General

Plan name

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

Plan's credentials (p. 119)

[Optional] The backup plan will run on behalf of the user who is creating the plan. You can change the plan account credentials if necessary. To access this option, select the **Advanced view** check box .

Comments

[Optional] Type a description of the backup plan. To access this option, select the **Advanced view** check box.

Label

[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. To access this option, select the **Advanced view** check box.

What to backup

Source type (p. 119)

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

Items to backup (p. 120)

Specify the data items to back up. A list of items to backup depends on the data type, specified previously.

Access credentials (p. 121)

[Optional] Provide credentials for the source data if the plan's account does not have access permissions to the data. To access this option, select the **Advanced view** check box .

Exclusions (p. 122)

[Optional] Set up exclusions for the specific types of files you do not wish to back up. To access this option, select the **Advanced view** check box.

Where to back up

Archive (p. 123)

Specify path to the location, where the backup archive will be stored, and the archive name. It is advisable that the archive name be unique within the location. The default archive name is Archive(N) where N is the sequence number of the archive in the location you have selected.

Name backup files using the archive name, as in Acronis True Image Echo, rather than auto-generated names

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

[Optional] Select this check box if you want to use simplified file naming for the archive's backups.

Access credentials (p. 128)

[Optional] Provide credentials for the location if the plan account does not have access permissions to the location. To access this option, select the **Advanced view** check box.

Archive comments

[Optional] Enter comments on the archive. To access this option, select the **Advanced view** check box.

How to back up

Backup scheme (p. 129)

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. Use well-known optimized backup schemes, such as Grandfather-Father-Son and Tower of Hanoi; create a custom backup scheme, or back up data once.

Archive validation

When to validate (p. 138)

[Optional] Define when and how often to perform validation and whether to validate the entire archive or the latest backup in the archive.

Backup options

Settings

[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. 50) 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 **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 and so you always see only the settings that differ from the default values in this section of the **Create backup plan** page.

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

Convert to VM

Applies to: **Disk/volume** backup, backup of **Entire virtual machines** or **Volumes of a virtual machine**

Not available on machines running Linux

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

When to convert (p. 139)

[Optional] Specify whether to convert every full, every incremental or every differential backup or convert the last created backup on schedule. Specify the conversion schedule if required.

Host (p. 139)

Specify the machine that will perform the conversion. The machine has to have Acronis Backup & Recovery 10 Agent for Windows, Agent for ESX/ESXi or Agent for Hyper-V installed.

Virtualization server (p. 139)

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

Storage (p. 139)

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.

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. 119).

The plan you have created will be accessible for examination and managing in the **Backup plans and tasks** (p. 103) view.

6.2.1 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.

6.2.2 Backup plan's credentials

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

To specify credentials

1. Select one of the following:

- **Run under the current user**

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

- **Use the following credentials**

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

Specify:

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

2. Click **OK**.

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

6.2.3 Source type

Select the type of data you want to be backed up on the managed machine. The list of available data types depends on the agents running on the machine:

Files

Available if the Acronis Backup & Recovery 10 Agent for Windows (or for Linux) is installed.

Select this option to back up specific files and folders.

If you are not concerned about recovery of the operating system along with all the settings and applications, but plan to keep safe only certain data (the current project, for example), choose file backup. This will reduce the archive size, thus saving storage space.

Disks/volumes

Available if the Acronis Backup & Recovery 10 Agent for Windows (or for Linux) is installed.

Select this option to back up disks and/or volumes. To be able to back up disks or volumes, you must have Administrator or Backup operator privileges.

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

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.

6.2.4 Items to back up

The items to backup depend on the source type (p. 119) selected previously.

6.2.4.1 Selecting disks and volumes

To specify disks/volumes to back up

1. Select the check boxes for the disks and/or volumes to back up. You can select a random set of disks and 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.

In Linux, logical volumes and MD devices are shown under **Dynamic and GPT**. For more information about backing up such volumes and devices, see “Backing up LVM volumes and MD devices (Linux)”.

2. [Optional] 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** 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.
3. Click **OK**.

What does a disk or volume backup store?

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

Windows

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

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

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

Linux

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

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

6.2.4.2 Selecting files and folders

To select files and/or folders for backing up

1. Expand the local folders tree items in order to view its nested folders and files.
2. Select an item by checking the corresponding check box in the tree. Selecting a check box for a folder means that all its content (files and folders) will be backed up. That is also the case for new files that will appear there in the future.

A file-based backup is not sufficient for recovery of the operating system. In order to recover your operating system, you have to perform a disk backup.

Use the table in the right part of the window to browse and select the nested items. Selecting the check box beside the **Name** column's header automatically selects all items in the table. Clearing this check box automatically deselects all items.

3. Click **OK**.

6.2.5 Access credentials for source

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

To specify credentials

1. Select one of the following:
 - **Use the plan's credentials**

The program will access the source data using the credentials of the backup plan account specified in the General section.
 - **Use the following credentials**

The program will access the source data using the credentials you specify. Use this option if the plan's account does not have access permissions to the data.

Specify:

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

6.2.6 Exclusions

Set up exclusions for the specific types of files you do not wish to back up. For example, you may not want database, hidden and system files and folders, as well as files with specific extensions, to be stored in the archive.

To specify which files and folders to exclude:

Set up any of the following parameters:

- **Exclude all hidden files and folders**

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

- **Exclude all system files and folders**

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

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

- **Exclude files matching the following criteria**

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

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

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

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

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

Exclusion examples

Criterion	Example	Description
Windows and Linux		
By name	F.log	Excludes all files named "F.log"
	F	Excludes all folders named "F"
By mask (*)	*.log	Excludes all files with the .log extension
	F*	Excludes all files and folders with names starting with "F" (such as folders F, F1 and files F.log, F1.log)
By mask (?)	F????.log	Excludes all .log files with names consisting of four symbols and starting with "F"
Windows		

By file path	C:\Finance\F.log	Excludes the file named "F.log" located in the folder C:\Finance
By folder path	C:\Finance\F\	Excludes the folder C:\Finance\F (be sure to specify the full path starting from the disk letter)
Linux		
By file path	/home/user/Finance/F.log	Excludes the file named "F.log" located in the folder /home/user/Finance
By folder path	/home/user/Finance/	Excludes the folder /home/user/Finance

6.2.7 Archive

Specify where the archive will be stored and the name of the archive.

1. Selecting the destination

Enter the full path to the destination in the **Path** field, or select the desired destination in the folders tree.

- To back up data to Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then, expand the **Online backup storage** group and select the account. Prior to backing up to the online storage, you need to buy a subscription (p. 221) to the online backup service and activate (p. 222) the subscription on the machine(s) you want to back up. Online backup is not available in Linux and under bootable media.

Acronis Backup & Recovery 10 Online might be unavailable in your region. To find more information, click here: <http://www.acronis.eu/my/backup-recovery-online/>

- To back up data to a centralized vault, expand the **Centralized** group and click the vault.
- To back up data to a personal vault, expand the **Personal** group and click the vault.
- To back up data to a local folder on the machine, expand the **Local folders** group and click the required folder.
- To back up 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 back up data to an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

ftp://ftp_server:port _number or sftp://sftp_server:port number

If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

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

- To back up data to a locally attached tape device, expand the **Tape drives** group, then click the required device.

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(1). 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 and click **OK**.

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, but in general you should 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.

6.2.8 Simplified naming of backup files

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

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

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

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

This behavior is useful when you rotate USB hard drives and want each drive to keep a single full backup (p. 126) or all backups created during a week (p. 126). 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. 128) to the archive name.

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

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

Restrictions

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

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

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

6.2.8.1 Usage examples

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

Example 1. Daily backup replacing the old one

Consider the following scenario:

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

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

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

Example 2. Daily full backups with a date stamp

Consider the following scenario:

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

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

Result:

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

See also “The [Date] variable” (p. 128).

Example 3. Hourly backups within a day

Consider the following scenario:

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

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

Result:

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

See also “The [Date] variable” (p. 128).

Example 4. Daily full backups with daily drive swaps

Consider the following scenario:

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

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

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

Example 5. Daily backups with weekly drive swaps

Consider the following scenario:

- You want to perform daily backups of your machine: a full backup each Monday and incremental backups on Tuesday through Sunday.
- You want to back up to the archive **MyMachine** on an external hard disk drive.

- You have two such drives. Either of them has drive letter **D** in the operating system when attached to the machine.
- You want to swap the drives each Monday, so that one drive contains backups of the current week (Monday through Sunday), and the other drive those of the previous week.

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

- When creating the first backup plan, specify **MyMachine** as the archive name and **D:** as the archive location, select the **Name backup files using the archive name...** check box, select **Full** as the backup type, and schedule the backups to run every week on Monday.
- When creating the second backup plan, specify the same settings as in the first backup plan, but select **Incremental** as the backup type and schedule the backups to run every week on Tuesday through Sunday.

Result:

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

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:

- 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.
- When creating the second backup plan, specify the same settings as in the first backup plan, but select **Differential** as the backup type and schedule the backups as follows:
 - **Run the task: Daily**
 - **Every: 1 Hour(s)**
 - **From: 08:00:00 AM**
 - **Until: 05:01:00 PM**

Result:

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

See also "The [Date] variable" (p. 128).

6.2.8.2 The [DATE] variable

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

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

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

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

Examples

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

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

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

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

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

6.2.8.3 Backup splitting and simplified file naming

When a backup is split according to backup splitting (p. 64) 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**.

6.2.9 Access credentials for archive location

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

To specify credentials

1. Select one of the following:
 - **Use the plan's credentials**

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

- **Use the following credentials**

The program will access the source data using the credentials you specify. Use this option if the plan account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

Specify:

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

2. Click **OK**.

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

6.2.10 Backup schemes

Choose one of the available backup schemes:

- **Back up now** – to create a backup task for manual start and run the task immediately after its creation.
- **Back up later** – to create a backup task for manual start OR schedule one-time task execution in the future.
- **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, where you schedule when and how often to back up (sessions) and select the number of backup levels (up to 16). In this scheme, 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.
- **Initial seeding** - to save locally a full backup whose final destination is Acronis Online Backup Storage.

6.2.10.1 Back up now scheme

With the **Back up now** scheme, the backup will be performed immediately, right after you click the **OK** button at the bottom of the page.

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

6.2.10.2 Back up later scheme

With the Back up later scheme, the backup will be performed only once, at the date and time you specify.

Specify the appropriate settings as follows

Backup type	Select the type of backup: full, incremental, or differential. If there is no full backup in the archive, a full backup will be created regardless of your selection.
Date and time	Specify when to start the backup.
The task will be started manually	Select this check box, if you do not need to put the backup task on a schedule and wish to start it manually afterwards.

6.2.10.3 Simple scheme

With the simple backup scheme you just schedule when and how often to back up data and set the retention rule. At the first time a full backup will be created. The next backups will be incremental.

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

Backup	Set up the backup schedule - when and how often to back up the data. To learn more about setting up the schedule, see the Scheduling (p. 87) section.
Retention rule	With the simple scheme, only one retention rule (p. 31) is available. Set the retention period for the backups.

6.2.10.4 Grandfather-Father-Son scheme

At a glance

- Daily incremental, weekly differential, and monthly full backups
- Custom day for weekly and monthly backups
- Custom retention periods for backups of each type

Description

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

	Mo	Tu	We	Th	Fr	Sa	Su
Jan 1—Jan 7	D	D	D	D	W	-	-
Jan 8—Jan 14	D	D	D	D	W	-	-
Jan 15—Jan 21	D	D	D	D	W	-	-
Jan 22—Jan 28	D	D	D	D	M	-	-
Jan 29—Feb 4	D	D	D	D	W	-	-

Feb 5—Feb 11	D	D	D	D	W	-	-
Feb 12—Feb 18	D	D	D	D	W	-	-
Feb 19—Feb 25	D	D	D	D	M	-	-
Feb 26—Mar 4	D	D	D	D	W	-	-

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

- Monthly ("Grandfather") backups are full;
- Weekly ("Father") backups are differential;
- Daily ("Son") backups are incremental.

Parameters

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

Start backup at:	Specifies when to start a backup. The default value is 12:00 PM.
Back up on:	Specifies the days on which to perform a backup. The default value is Workdays.
Weekly/Monthly:	Specifies which of the days selected in the Back up on field you want to reserve for weekly and monthly backups. A monthly backup will be performed every fourth such day. The default value is Friday.
Keep backups:	<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: 7 days (recommended minimum)</p> <p>Weekly: 4 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>

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

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

Examples

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

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

- Back up files every day, including weekends
- Be able to recover files as of any date over the past seven days
- Have access to weekly backups of the past month

- Keep monthly backups indefinitely.

Backup scheme parameters can then be set up as follows.

- Start backup at: **11:00 PM**
- Back up on: **All days**
- Weekly/monthly: **Saturday** (for example)
- Keep backups:
 - Daily: **1 week**
 - Weekly: **1 month**
 - Monthly: **indefinitely**

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

Limited storage

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

Suppose that you need to:

- Perform backups at the end of each working day
- Be able to recover an accidentally deleted or inadvertently modified file if this has been discovered relatively quickly
- Have access to a weekly backup for 10 days after it was created
- Keep monthly backups for half a year.

Backup scheme parameters can then be set up as follows.

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

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

Work schedule

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

- Track changes to the financial statements, spreadsheets, etc. performed on Tuesdays and Thursdays (daily incremental backup).

- Have a weekly summary of file changes since last month (Friday weekly differential backup).
- Have a monthly full backup of your files.

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

The following GFS scheme suits such purposes:

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

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

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

No daily backups

Consider a more exotic GFS scheme:

- Start backup at: **12:00 PM**
- Back up on: **Friday**
- Weekly/monthly: **Friday**
- Keep backups:
 - Daily: **1 week**
 - Weekly: **1 month**
 - Monthly: **indefinitely**

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

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

6.2.10.5 Tower of Hanoi scheme

At a glance

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

Parameters

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

Schedule	Set up a daily (p. 88), weekly (p. 90), or monthly (p. 92) schedule. Setting up schedule parameters allows creating 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.

Example

Schedule parameters are set as follows

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

Number of levels: 4

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.

6.2.10.6 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	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	<p>Specifies on what schedule and under which conditions to perform an incremental backup.</p> <p>If the archive contains no backups at the time of the task run, a full backup is created instead of the incremental backup.</p>
Differential	<p>Specifies on what schedule and under which conditions to perform a differential backup.</p> <p>If the archive contains no full backups at the time of the task run, a full backup is created instead of the differential backup.</p>
Clean up archive	<p>Specifies how to get rid of old backups: either to apply retention rules (p. 31) 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 or centralized vaults.</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 program 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 managed vaults.</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 the rules (only if the retention rules are set)	<p>Specifies when to apply the retention rules (p. 31).</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 the rules.</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. 31).

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. 31).

Resulting tasks

Any custom scheme always produces three backup tasks and—in case the retention rules are specified—a cleanup task. Each task is listed in the list of tasks either as **Scheduled** (if the schedule has been set up) or as **Manual** (if the schedule has not been set up).

You can manually run any backup task or cleanup task at any time, regardless of whether it has a schedule.

In the first of the previous examples, we set up a schedule only for full backups. However, the scheme will still result in three backup tasks, enabling you to manually start a backup of any type:

- Full backup, runs every Friday at 10:00 PM
- Incremental backup, runs manually
- Differential backup, runs manually

You can run any of these backup tasks by selecting it from the list of tasks in the **Backup plans and tasks** section in the left pane.

If you have also specified the retention rules in your backup scheme, the scheme will result in four tasks: three backup tasks and one cleanup task.

6.2.11 Archive validation

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

To set up validation, specify the following parameters

1. **When to validate** – select when to perform the validation. As the validation is a resource-intensive operation, it makes sense to **schedule** the validation to the managed machine's off-peak period. On the other hand, if the validation is a major part of your data protection strategy and you prefer to be immediately informed whether the backed up data is not corrupted and can be successfully recovered, think of starting the validation right after backup creation.
2. **What to validate** – select either to validate the entire archive or the latest backup in the archive. Validation of a file backup imitates recovery of all files from the backup to a dummy destination. Validation of a volume backup calculates a checksum for every data block saved in the backup.

Validation of the archive will validate all the archive's backups and may take a long time and a lot of system resources.

3. **Validation schedule** (appears only if you have selected the on schedule in step 1) - set the schedule of validation. For more information see the Scheduling (p. 87) section.

6.2.12 Setting up regular conversion to a virtual machine

When creating a backup plan (p. 116), you can set up regular conversion of a disk or volume backup to a virtual machine. This section provides information that helps you make the appropriate settings.

6.2.12.1 Setting up a conversion schedule

A disk backup (p. 234) created while executing a backup plan can be converted to a virtual machine immediately or on schedule or you can combine both methods.

The conversion task will be created on the machine being backed up, and will use this machine's date and time.

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

The old virtual machine must be powered off by the time of conversion, otherwise it will not be possible to delete it and 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.

6.2.12.2 Selecting a host that will perform conversion

Specify the machine that will perform the conversion. The machine has to have Acronis Backup & Recovery 10 Agent for Windows, Agent for ESX/ESXi or Agent for Hyper-V installed.

Take into account the following considerations.

Which agent is installed on the host?

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

- **Agent for Windows** is installed on the host
You have a choice of virtual machine types: VMware Workstation, Microsoft Virtual PC, or Parallels Workstation. Files of the new virtual machine will be placed in the folder you select.
- **Agent for ESX/ESXi** is installed on the host
A VMware virtual machine will be created on the ESX/ESXi server.
Virtual machines resulting from backup are not supposed to be backed up and so do not appear on the management server, unless its integration with VMware vCenter Server is enabled. If the integration is enabled, such machines appear as unmanageable. A backup policy cannot be applied to them.
- **Agent for Hyper-V** is installed on the host

You can choose between creating a virtual machine on the Hyper-V server and creating a VMware Workstation, Microsoft Virtual PC or Parallels Workstation machine in the folder you select.

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

What is the host's processing power?

The conversion task will be created on the machine being backed up, and will use this machine's date and time. In fact the task will be executed by the host that you select and so will take that host's CPU resource. If multiple backup plans use the same host, multiple conversion tasks will be queued on that host and it may take considerable time to complete them all.

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 host 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.

Disk space

On VMware ESX/ESXi, new machines are created with pre-allocated disks. 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.

Virtual machines created on a Hyper-V server or workstation type machines (VMware Workstation, Microsoft Virtual PC or Parallels Workstation) use as much disk space as the original data occupies. Since the space is not pre-allocated, the physical disk on which the virtual machine will run is expected to have sufficient free space for the virtual disks to increase in size.

6.3 Recovering data

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. 231) or using Acronis Startup Recovery Manager (p. 40). Then, create a recovery task.

Acronis Universal Restore (p. 40) lets you recover and boot up **Windows on dissimilar hardware** or a virtual machine.

A **Windows system can be brought online in seconds** while it is still being recovered. Using the proprietary Acronis Active Restore (p. 42) technology, Acronis Backup & Recovery 10 will boot the machine into the operating system found in the backup as if the system were on the physical disk. The system becomes operational and ready to provide necessary services. Thus, the system downtime will be minimal.

A **dynamic volume** can be recovered over an existing volume, to unallocated space of a disk group, or to unallocated space of a basic disk. To learn more about recovering dynamic volumes, please turn to the Microsoft LDM (Dynamic volumes) (p. 34) section.

Acronis Backup & Recovery 10 Agent for Windows has the ability to recover a disk (volume) backup to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat KVM. The virtual appliance can then be imported to XenServer. The VMware Workstation machine can be converted to the open virtualization format (OVF) using the VMware OVF tool. With Acronis Backup & Recovery 10 Agent for Hyper-V or Agent for ESX/ESXi, you can create a new virtual machine on the respective virtualization server.

You might need to prepare target disks before recovery. Acronis Backup & Recovery 10 includes a handy disk management utility which enables you to create or delete volumes, change a disk partitioning style, create a disk group and perform other disk management operations on the target hardware, both under the operating system and on bare metal. To find out more about Acronis Disk Director LV, see the Disk management (p. 193) section.

To create a recovery task, perform the following steps

General

Task name

[Optional] Enter a unique name for the recovery task. A conscious name lets you quickly identify the task among the others.

Task credentials (p. 143)

[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, select the **Advanced view** check box .

What to recover

Archive (p. 143)

Select the archive to recover data from.

Data type (p. 144)

Applies to: disk recovery

Choose the type of data you need to recover from the selected disk backup.

Content (p. 144)

Select the backup and content to be recovered.

Access credentials (p. 145)

[Optional] Provide credentials for the archive location if the task account does not have the right to access it. To access this option, select the **Advanced view** check box.

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

Volumes

Acronis Active Restore

[OPTIONAL] The **Acronis Active Restore** check box is available when recovering Windows starting from Windows 2000. Acronis Active Restore brings a system online immediately after the recovery is started. The operating system boots from the backup image and the machine becomes operational and ready to provide necessary services. The data required to

serve incoming requests is recovered with the highest priority; everything else is recovered in the background.

See Acronis Active Restore (p. 42) for details.

Files (p. 151)

You may have to specify credentials for the destination. Skip this step when operating on a machine booted with bootable media.

Access credentials (p. 152)

[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. 153)

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.

[Optional] Acronis Universal Restore

Applies to: Windows OS and system volume recovery

Universal Restore (p. 153)

Use the Acronis Universal Restore when you need to recover and boot up Windows on dissimilar hardware.

Automatic drivers search

Specify where the program should search for HAL, mass storage and network adapter drivers. Acronis Universal Restore will install drivers that better fit the target hardware.

Mass storage drivers to install anyway

[Optional] Specify the mass storage drivers manually if the automatic drivers search has not found the appropriate drivers. To access this option, select the **Advanced view** check box.

Recovery options

Settings

[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. 71) 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 **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 and so you always see only the settings that differ from the default values in the **Settings** section.

Clicking **Reset to default** resets all the settings to default values.

After you complete all the required steps, click **OK** to create the commit creating of the recovery task.

6.3.1 Task credentials

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

To specify credentials

1. Select one of the following:

- **Run under the current user**

The task will run under the credentials with which the user who starts the tasks is logged on. If the task has to run on schedule, you will be asked for the current user's password on completing the task creation.

- **Use the following credentials**

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

Specify:

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

2. Click **OK**.

To learn more about using credentials in Acronis Backup & Recovery 10, see the Owners and credentials (p. 24) section.

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

6.3.2 Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field, or select the desired folder in the folders tree.

- If the archive is stored in Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then expand the **Online backup storage** group and select the account.

Exporting and mounting are not supported for backups stored in Acronis Online Backup Storage.

- If the archive is stored in a centralized vault, expand the **Centralized** group and click the vault.
- If the archive is stored in a personal vault, expand the **Personal** group and click the vault.
- If the archive is stored in a local folder on the machine, expand the **Local folders** group and click the required folder.

If the archive is located on removable media, e.g. DVDs, first insert the last DVD and then insert the discs in order starting from the first one when the program prompts.

- If the archive is stored on a network share, expand the **Network folders** group, then select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.

- If the archive is stored on an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

ftp://ftp_server:port_number or sftp://sftp_server:port number

If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

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

- If the archive is stored on a locally attached tape device, expand the **Tape drives** group, then click the required device.

When operating on a machine booted with bootable media:

- To access a managed vault, type the following string in the **Path** field:
bsp://node_address/vault_name/

- To access an unmanaged centralized vault, type the full path to the vault's folder.

2. In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each vault/folder you select.

While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3. Click **OK**.

6.3.3 Data type

Choose what type of data to recover from the selected disk backup:

- **Disks** - to recover disks
- **Volumes** - to recover volumes
- **Files** - to recover specific files and folders

6.3.4 Content selection

The representation of this window depends on the type of data stored in the archive.

6.3.4.1 Disks/volumes selection

To select a backup and disks/volumes to recover:

1. 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.

Specify the items to recover. By default, all items of the selected backup will be selected. If you do not want to recover certain items, just uncheck them.

To obtain information on a disk/volume, right-click it and then click **Information**.

2. Click **OK**.

Selecting an MBR

You will usually select the disk's MBR if:

- The operating system cannot boot
- The disk is new and does not have an MBR
- Recovering custom or non-Windows boot loaders (such as LILO and GRUB)
- The disk geometry is different to that stored in the backup.

There are probably other times when you may need to recover the MBR, but the above are the most common.

When recovering the MBR of one disk to another Acronis Backup & Recovery 10 recovers Track 0, which does not affect the target disk's partition table and partition layout. Acronis Backup & Recovery 10 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.

6.3.4.2 Files selection

To select a backup and files to recover:

1. Select one of the successive backups by its creation date/time. Thus, you can revert the files/folders to a specific moment in time.
2. Specify the files and folders to recover by selecting the corresponding check boxes in the archives tree.

Selecting a folder automatically selects all its nested folders and files.

Use the table to the right of the archives tree to select the nested items. Selecting the check box for the **Name** column's header automatically selects all items in the table. Clearing this check box automatically deselects all the items.

3. Click **OK**.

6.3.5 Access credentials for location

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 program will access the location using the credentials of the task account specified in the General section.
 - **Use the following credentials**
The program will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.
Specify:
 - **User name.** When entering the name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain).
 - **Password.** The password for the account.
2. Click **OK**.

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

6.3.6 Destination selection

Specify the destination the selected data will be recovered to.

6.3.6.1 Disks

Available disk destinations depend on the agents operating on the machine.

Recover to:

Physical machine

Available when the Acronis Backup & Recovery 10 Agent for Windows or Agent for Linux is installed.

The selected disks will be recovered to the physical disks of the machine the console is connected to. On selecting this, you proceed to the regular disk mapping procedure described below.

New virtual machine (p. 150)

If Acronis Backup & Recovery 10 Agent for Windows is installed.

The selected disks will be recovered to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat KVM. The virtual machine files will be saved to the destination you specify.

If Acronis Backup & Recovery 10 Agent for Hyper-V or Agent for ESX/ESXi is installed.

These agents enable creating a new virtual machine on the virtualization server you specify. The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **Virtual Machine Settings** (p. 150) section. Check the settings and make changes if necessary.

Then you proceed to the regular disk mapping procedure described below.

Existing virtual machine

Available when the Acronis Backup & Recovery 10 Agent for Hyper-V or Agent for ESX/ESXi is installed.

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular disk mapping procedure described below.

Please be aware that the target machine will be powered off automatically before the recovery. If you prefer to power it off manually, modify the VM power management option.

Disk #:

Disk # (MODEL) (p. 148)

Select the destination disk for each of the source disks.

NT signature (p. 147)

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

When the MBR is selected along with the disk backup, you need to retain operating system bootability on the target disk volume. The operating system must have the system volume information (e.g. volume letter) matched with the disk NT signature, which is kept in the MBR disk record. But two disks with the same NT signature cannot work properly under one operating system.

If there are two disks having the same NT signature and comprising of a system volume on a machine, at the startup the operating system runs from the first disk, discovers the same signature on the second one, automatically generates a new unique NT signature and assigns it to the second disk. As a result, all the volumes on the second disk will lose their letters, all paths will be invalid on the disk, and programs won't find their files. The operating system on that disk will be unbootable.

To retain system bootability on the target disk volume, choose one of the following:

- **Select automatically**

A new NT signature will be created only if the existing one differs from the one in the backup. Otherwise, the existing NT signature will be kept.

- **Create new**

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

- **Recover from backup**

The program will replace the NT signature of the target hard disk with one from the disk backup. Recovering the disk signature may be desirable due to the following reasons:

- Acronis Backup & Recovery 10 creates scheduled tasks using the signature of the source hard disk. If you recover the same disk signature, you don't need to re-create or edit the tasks created previously
- Some installed applications use disk signature for licensing and other purposes
- This enables to keep all the Windows Restore Points on the recovered disk
- To recover VSS snapshots used by Windows Vista's "Previous Versions" feature

- **Keep existing**

The program will leave the existing NT signature of the target hard disk as is.

6.3.6.2 Volumes

Available volume destinations depend on the agents operating on the machine.

Recover to:

Physical machine

Available when the Acronis Backup & Recovery 10 Agent for Windows or Agent for Linux is installed.

The selected volumes will be recovered to the physical disks of the machine the console is connected to. On selecting this, you proceed to the regular volume mapping procedure described below.

New virtual machine (p. 150)

If Acronis Backup & Recovery 10 Agent for Windows is installed.

The selected volumes will be recovered to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat KVM. The virtual machine files will be saved to the destination you specify.

If Acronis Backup & Recovery 10 Agent for Hyper-V or Agent for ESX/ESXi is installed.

These agents enable creating a new virtual machine on the virtualization server you specify.

The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **Virtual Machine Settings** (p. 150) section. Check the settings and make changes if necessary.

Then you proceed to the regular volume mapping procedure described below.

Existing virtual machine

Available when the Acronis Backup & Recovery 10 Agent for Hyper-V or Agent for ESX/ESXi is installed.

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular volume mapping procedure described below.

Please be aware that the target machine will be powered off automatically before recovery. If you prefer to power it off manually, modify the VM power management option.

Recover [Disk #] MBR to: [If the Master Boot Record is selected for recovery]

Disk # (p. 148)

Choose the disk to recover the Master Boot Record to.

NT signature: (p. 147)

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 (p. 148)

Sequentially map each of the source volumes to a volume or an unallocated space on the destination disk.

Size:

[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 destination volume:

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...).

Volume properties

Resizing and relocating

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.

Tip: A volume cannot be resized when being recovered from a backup split into multiple DVDs or tapes. To be able to resize the volume, copy all parts of the backup to a single location on a hard disk.

Properties

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

Change the volume file system, if required. By default, the program selects the original volume's file system. Acronis Backup & Recovery 10 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 4GB, so you will not be able to recover a 4GB 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.

Logical drive letter (for Windows only)

Assign a letter to the recovered volume. Select the desired letter from a drop-down list.

- With the default AUTO selection, the first unused letter will be assigned to the volume.
- If you select NO, no letter will be assigned to the recovered volume, hiding it from the OS. You should not assign letters to volumes that are inaccessible to Windows, such as to those other than FAT and NTFS.

6.3.6.3 Virtual machine type / virtualization server selection

The new virtual machine can be created either on a virtualization server (this requires Acronis Backup & Recovery 10 Agent for Hyper-V or Agent for ESX/ESXi to be installed) or in any accessible local or networked folder.

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

1. Choose the **Place on the virtualization server that I select** option.
2. In the left part of the window, select the virtualization server. Use the right part of the window to review details on the selected server.
3. Click **OK** to return to the **Data recovery** page.

To select the type of virtual machine

1. Choose the **Save as files of the VM type that I select to the folder that I specify** option.
2. In the left part of the window, select the virtual machine type. Use the right part of the window to review details on the selected virtual machine type.
3. Click **OK** to return to the **Data recovery** page.

6.3.6.4 Virtual machine settings

The following virtual machine settings can be configured.

Storage

Initial setting: the default storage of the virtualization server if the new machine is created on the virtualization server. Otherwise the current user's documents folder.

This is the place where the new virtual machine will be created. Whether you can change the storage on the virtualization server or not, depends on the virtualization product brand and settings. VMware ESX may have multiple storages. A Microsoft Hyper-V server enables creating a new virtual machine in any local folder.

Memory

Initial setting: if not contained in the backup, the default setting of the virtualization server.

This is the amount of memory allocated to the new virtual machine. The memory adjustment range depends on the host hardware, the host operating system and the virtualization product settings. For example, virtual machines may be allowed to use no more than 30% of memory.

Disks

Initial setting: the number and size of the source machine's disks.

The number of disks is generally equal to that of the source machine, but might be different if the program has to add more disks to accommodate the source machine volumes because of limitations

set by the virtualization product. You can add virtual disks to the machine configuration or, in some cases, delete the proposed disks.

Implementation of Xen machines is based on Microsoft Virtual PC and inherits its limitations: up to 3 IDE disks and 1 processor. SCSI disks are not supported.

Processors

Initial setting: if not contained in the backup or the backed up setting is not supported by the virtualization server, 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.

6.3.6.5 File destination

To specify a destination:

1. 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.
2. Click **OK**.

Recovery exclusions

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

Use the **Add**, **Edit**, **Remove** and **Remove All** buttons to create the list of file masks. Files whose names match any of the masks will be skipped during recovery.

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

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

Exclusion examples

Criterion	Example	Description
Windows and Linux		
By name	F.log	Excludes all files named "F.log"
	F	Excludes all folders named "F"

By mask (*)	*.log F*	Excludes all files with the .log extension Excludes all files and folders with names starting with "F" (such as folders F, F1 and files F.log, F1.log)
By mask (?)	F????.log	Excludes all .log files with names consisting of four symbols and starting with "F"
Windows		
By file path	Finance\F.log	Excludes files named "F.log" from all folders with the name "Finance"
By folder path	Finance\F\ or Finance\F	Excludes folders named "F" from all folders with the name "Finance"
Linux		
By file path	/home/user/Finance/F.log	Excludes the file named "F.log" located in the folder /home/user/Finance

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

Overwriting

Choose what to do if the program finds in the target folder a file with the same name as in the archive:

- **Overwrite existing file** - this will give the file in the backup priority over the file on the hard disk.
- **Overwrite existing file if it is older** - this will give priority to the most recent file modification, whether it be in the backup or on the disk.
- **Do not overwrite existing file** - this will give the file on the hard disk priority over the file in the backup.

If you allow files to be overwritten, you still have an option to prevent overwriting of specific files by excluding (p. 151) them from the recovery operation.

6.3.7 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 General 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**.

6.3.8 When to recover

Select when to start the recovery task:

- **Recover now** - the recovery task will be started immediately after you click the final **OK**.
- **Recover later** - the recovery task will be started at the date and time you specify.

If you do not need to schedule the task and wish to start it manually afterwards, select the **Task will be started manually (do no schedule the task)** check box.

6.3.9 Universal Restore

Use Acronis Backup & Recovery 10 Universal Restore when you need to recover and boot up Windows on dissimilar hardware. Universal Restore handles differences in devices that are critical for the operating system startup, such as storage controllers, motherboard or chipset.

To learn more about the Universal Restore technology, see the Universal Restore (p. 40) section.

Acronis Backup & Recovery 10 Universal Restore is not available when:

- a machine is booted with Acronis Startup Recovery Manager (using F11)
- the backup image is located in Acronis Secure Zone
- you have chosen to use Acronis Active Restore (p. 228)

because these features are primarily meant for instant data recovery on the same machine.

Preparation

Before recovering Windows to dissimilar hardware, make sure that you have the drivers for the new HDD controller and the chipset. These drivers are critical to start the operating system. Use the CD or DVD supplied by the hardware vendor or download the drivers from the vendor's Web site. The driver files should have the *.inf, *.sys or *.oem extensions. If you download the drivers in the *.exe, *.cab or *.zip format, extract them using a third-party application, such as WinRAR (<http://www.rarlab.com/>) or Universal Extractor (<http://legroom.net/software/uniextract>).

The best practice is to store drivers for all the hardware used in your organization in a single repository sorted by device type or by the hardware configurations. You can keep a copy of the repository on a DVD or a flash drive; pick some drivers and add them to the bootable media; create the custom bootable media with the necessary drivers (and the necessary network configuration) for each of your servers. Or you can simply specify the path to the repository every time Universal Restore is used.

Make sure you have access to the device with drivers when working under bootable media. Even if you configure system disk recovery in a Windows environment, the machine will reboot and recovery will proceed in the Linux-based environment. Use WinPE-based media if the device is available in Windows but Linux-based media does not detect it.

Universal Restore settings

Automatic driver search

Specify where the program will search for the Hardware Abstraction Layer (HAL), HDD controller driver and network adapter driver(s):

- If the drivers are on a vendor's disc or other removable media, turn on the **Search removable media**.
- If the drivers are located in a networked folder or on the bootable media, specify the path to the folder in the **Search folder** field.

During recovery, Universal Restore will perform the recursive search in all the sub-folders of the specified folder, find the most suitable HAL and HDD controller drivers of all those available, and install them into the recovered system. Universal Restore also searches for the network adapter driver; the path to the found driver is then transmitted by Universal Restore to the operating system. If the hardware has multiple network interface cards, Universal Restore will try to configure all the cards' drivers. In case Universal Restore cannot find a compatible driver in the specified locations, it will specify the problem device and ask for a disc or a network path to the driver.

Once Windows boots, it will initialize the standard procedure for installing new hardware. The network adapter driver will be installed silently if the driver has the Microsoft Windows signature. Otherwise, Windows will ask for confirmation whether to install the unsigned driver.

After that, you will be able to configure the network connection and specify drivers for the video adapter, USB and other devices.

Mass storage drivers to install anyway

To access this option, select the **Advanced view** check box.

If the target hardware has a specific mass storage controller such as RAID (especially NVIDIA RAID) or a fibre channel adapter, specify the appropriate drivers in the **Drivers** field.

The drivers defined here will have priority. They will be installed, with appropriate warnings, even if the program finds a better driver.

Use this option only if the automatic drivers search does not help to boot the system.

Drivers for a virtual machine

When recovering a system to a new virtual machine, the Universal Restore technology is applied in the background, because the program knows what drivers are required for the supported virtual machines.

When recovering the system to an existing virtual machine that uses SCSI hard drive controller, be sure to specify SCSI drivers for the virtual environment, in the **Mass storage drivers to install anyway** step. Use drivers bundled with your virtual machine software or download the latest drivers versions from the software manufacturer Web site.

6.3.10 How to convert a disk backup to a virtual machine

Rather than converting a TIB file to a virtual disk file, which requires additional operations to bring the virtual disk into use, Acronis Backup & Recovery 10 performs the conversion by recovery of a disk backup to a fully configured and operational new virtual machine. You have the ability to adapt the virtual machine configuration to your needs when configuring the recovery operation.

With **Acronis Backup & Recovery 10 Agent for Windows**, you can recover a disk (volume) backup to a new virtual machine of any of the following types: VMware Workstation, Microsoft Virtual PC, Parallels Workstation, Citrix XenServer Open Virtual Appliance (OVA) or Red Hat KVM.

Files of the new virtual machine will be placed in the folder you select. You can start the machine using the respective virtualization software or prepare the machine files for further usage. The Citrix XenServer Open Virtual Appliance (OVA) can be imported to a XenServer using Citrix XenCenter. The VMware Workstation machine can be converted to the open virtualization format (OVF) using the VMware OVF tool.

With **Acronis Backup & Recovery 10 Agent for Hyper-V** or **Agent for ESX/ESXi**, you can recover a disk (volume) backup to a new virtual machine on the respective virtualization server.

Tip. *Microsoft Virtual PC does not support disks that are larger than 127 GB. Acronis enables you to create a Virtual PC machine with larger disks so that you can attach the disks to a Microsoft Hyper-V virtual machine.*

To convert a disk backup to a virtual machine:

1. Connect the console to a machine where Agent for Windows, Agent for Hyper-V or Agent for ESX/ESXi is installed.
2. Do any of the following:
 - Click **Recover** to open the **Recover data** page. Start creating a recovery task as described in "Recovering data". Select the archive and then select the disk or volume backup you want to convert.
 - Use the **Navigation** pane to navigate to the vault where the archive is stored. Select the archive and then select the disk or volume backup you want to convert. Click **Recover as virtual machine**. The **Recover data** page opens with the pre-selected backup.
3. In **Data type**, select **Disks** or **Volumes** depending on what you need to convert.
4. In **Content**, select the disks to convert or the volumes with the Master Boot Records (MBR) of the corresponding disks.
5. In **Recover to**, select **New virtual machine**.
6. In **VM server**, select the type of the new virtual machine to be created *or* on which virtualization server to create the machine.
7. In **VM name**, enter the name for the new virtual machine.
8. [Optionally] Review the **Virtual machine settings (p. 150)** and make changes if necessary. Here you can change the path to the new virtual machine.

The same type of machines with the same name cannot be created in the same folder. Change either the VM name, or the path if you get an error message caused by identical names.

9. Select the destination disk for each of the source disks or source volumes and MBRs.

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.

10. In **When to recover**, specify when to start the recovery task.
11. [Optionally] 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.
12. Click **OK**. If the recovery task is scheduled for the future, specify the credentials under which the task will run.

You will be taken to the **Backup plans and tasks** view where you can examine the state and progress of the recovery task.

Post-conversion operations

The resulting machine always has SCSI disk interface and basic MBR volumes. If the machine uses a custom boot loader, you might need to configure the loader to point to the new devices and reactivate it. Configuring GRUB is described in "How to reactivate GRUB and change its configuration (p. 157)".

***Tip.** If you want to preserve logical (LVM) volumes on a Linux machine, consider the alternative method of conversion. Create a new virtual machine, boot it using bootable media and perform recovery just like you do on a physical machine. The LVM structure can be automatically recreated (p. 190) during recovery if it has been saved in the backup.*

6.3.11 Bootability troubleshooting

If a system was bootable at the time of backup, you expect that it will boot after recovery. However, the information the operating system stores and uses for booting up may become outdated during recovery, especially if you change volume sizes, locations or destination drives. Acronis Backup & Recovery 10 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 for Windows: Recover the volume once again. When configuring recovery, opt for using Acronis Universal Restore and specify the appropriate HAL and mass storage drivers.
- **Windows was recovered to a dynamic volume that cannot be bootable**
Solution: Recover Windows to a basic, simple or mirrored volume.
- **A system volume was recovered to a disk that does not have an MBR**
When you configure recovery of a system volume to a disk that does not have an MBR, the program prompts whether you want to recover the MBR along with the system volume. Opt for not recovering, only if you do not want the system to be bootable.
Solution: Recover the volume once again along with the MBR of the corresponding disk.
- **The system uses Acronis OS Selector**
Because the Master Boot Record (MBR) can be changed during the system recovery, Acronis OS Selector, which uses the MBR, might become inoperable. If this happens, reactivate Acronis OS Selector as follows.
Solution: Boot the machine from the Acronis Disk Director's bootable media and select in the menu **Tools -> Activate OS Selector**.

- **The system uses GRand Unified Bootloader (GRUB) and was recovered from a normal (not from a raw, that is, sector-by-sector) backup**

One part of the GRUB loader resides either in the first several sectors of the disk or in the first several sectors of the volume. The rest is on the file system of one of the volumes. System bootability can be recovered automatically only when the GRUB resides in the first several sectors of the disk and on the file system to which direct access is possible. In other cases, the user has to manually reactivate the boot loader.

Solution: Reactivate the boot loader. You might also need to fix the configuration file.

- **The system uses Linux Loader (LILO) and was recovered from a normal (not from a raw, that is, sector-by-sector) backup**

LILO contains numerous references to absolute sector numbers and so cannot be repaired automatically except for the case when all data is recovered to the sectors that have the same absolute numbers as on the source disk.

Solution: Reactivate the boot loader. You might also need to fix the loader configuration file for the reason described in the previous item.

- **The system loader points to the wrong volume**

This may happen when system or boot volumes are not recovered to their original location.

Solution:

Modification of the boot.ini or the boot\bcd files fixes this for Windows loaders. Acronis Backup & Recovery 10 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.

6.3.11.1 How to reactivate GRUB and change its configuration

Generally, you should refer to the boot loader manual pages for the appropriate procedure. There is also the corresponding Knowledge Base article on the Acronis Web site.

The following is an example of how to reactivate GRUB in case the system disk (volume) is recovered to identical hardware.

1. Start Linux or boot from the bootable media, and then press CTRL+ALT+F2.
2. Mount the system you are recovering:

```
mkdir /mnt/system/
mount -t ext3 /dev/sda2 /mnt/system/ # root partition
mount -t ext3 /dev/sda1 /mnt/system/boot/ # boot partition
```

3. Mount the **proc** and **dev** file systems to the system you are recovering:

```
mount -t proc none /mnt/system/proc/
mount -o bind /dev/ /mnt/system/dev/
```

4. Save a copy of the GRUB menu file, by running one of the following commands:

```
cp /mnt/system/boot/grub/menu.lst /mnt/system/boot/grub/menu.lst.backup
```

or

```
cp /mnt/system/boot/grub/grub.conf /mnt/system/boot/grub/grub.conf.backup
```

5. Edit the **/mnt/system/boot/grub/menu.lst** file (for Debian, Ubuntu, and SUSE Linux distributions) or the **/mnt/system/boot/grub/grub.conf** file (for Fedora and Red Hat Enterprise Linux distributions)—for example, as follows:

```
vi /mnt/system/boot/grub/menu.lst
```

6. In the **menu.lst** file (respectively **grub.conf**), find the menu item that corresponds to the system you are recovering. This menu items have the following form:

```
title Red Hat Enterprise Linux Server (2.6.24.4)
    root (hd0,0)
    kernel /vmlinuz-2.6.24.4 ro root=/dev/sda2 rhgb quiet
    initrd /initrd-2.6.24.4.img
```

The lines starting with **title**, **root**, **kernel**, and **initrd** respectively determine:

- The title of the menu item.
 - The device on which the Linux kernel is located—typically, this is the boot partition or the root partition, such as **root (hd0,0)** in this example.
 - The path to the kernel on that device and the root partition—in this example, the path is **/vmlinuz-2.6.24.4** and the root partition is **/dev/sda2**. You can specify the root partition by label (such as **root=LABEL=/**), identifier (in the form **root=UUID=some_uuid**), or device name (such as **root=/dev/sda2**).
 - The path to the **initrd** service on that device.
7. Edit the file **/mnt/system/etc/fstab** to correct the names of any devices that have changed as a result of the recovery.
 8. Start the GRUB shell by running one of the following commands:

```
chroot /mnt/system/ /sbin/grub
```

or

```
chroot /mnt/system/ /usr/sbin/grub
```

9. Specify the disk on which GRUB is located—typically, the boot or root partition:

```
root (hd0,0)
```

10. Install GRUB. For example, to install GRUB in the master boot record (MBR) of the first disk, run the following command:

```
setup (hd0)
```

11. Exit the GRUB shell:

```
quit
```

12. Unmount the mounted file systems and then reboot:

```
umount /mnt/system/dev/
umount /mnt/system/proc/
umount /mnt/system/boot/
umount /mnt/system/
reboot
```

13. Reconfigure the bootloader by using tools and documentation from the Linux distribution that you use. For example, in Debian and Ubuntu, you may need to edit some commented lines in the **/boot/grub/menu.lst** file and then run the **update-grub** script; otherwise, the changes might not take effect.

6.3.11.2 About Windows loaders

Windows NT/2000/XP/2003

A part of the loader resides in the partition boot sector, the rest is in the files ntldr, boot.ini, ntddetect.com, ntbootdd.sys. boot.ini is a text file that contains the loader configuration. Example:

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Microsoft Windows XP Professional"
/noexecute=optin /fastdetect
```

Windows Vista/2008

A part of the loader resides in the partition boot sector, the rest is in the files bootmgr, boot\bcd. At starting Windows, boot\bcd is mounted to the registry key HKLM \BCD00000000.

6.3.12 Recovering a vast number of files from a file backup

Applies to: Microsoft Windows Server 2003

When recovering a very large number of files at a time (hundreds of thousands or millions) from a file backup, you might encounter the following problem:

- The recovery process fails, and the message "Error reading the file" appears.
- Not all of the files are recovered.

The most likely cause of the problem is an insufficient amount of memory allocated to the recovery process by the operating system's cache manager. You can either work around this problem or modify the registry to increase the amount of allocated memory, as described below.

To resolve the problem, do either of the following:

- Recover the files as two or more groups. For example, if the problem occurs when recovering 1 million files, try recovering the first 500,000 of them and then the remaining 500,000.
- Modify the registry as follows:

Note: This procedure requires restarting the machine. Use standard precautions when modifying the registry.

1. In Registry Editor, open the following registry subkey:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management
```

2. Add the **PoolUsageMaximum** entry to the subkey:

- Entry type: **DWORD Value**
- Base: **Decimal**
- Value: **40**

3. Add the **PagedPoolSize** entry to the subkey:

- Entry type: **DWORD Value**
- Base: **Hexadecimal**
- Value: **FFFFFFF**

4. Quit Registry Editor, and then restart the machine.

If this does not resolve the problem, or for more details on adding these registry settings, see the corresponding Microsoft Help and Support article.

Tip: In general, if a volume contains many files, consider using a disk-level backup instead of a file-level one. In this case, you will be able to recover the entire volume as well as particular files stored on it.

6.4 Validating vaults, archives and backups

Validation is an operation that checks the possibility of data recovery from a backup.

Validation of a file backup imitates recovery of all files from the backup to a dummy destination. Validation of a disk or volume backup calculates a checksum for every data block saved in the backup. Both procedures are resource-intensive.

Validation of an archive will validate all the archive's backups. A vault (or a location) validation will validate all archives stored in this vault (location).

While successful validation means high probability of successful recovery, it does not check all factors that influence the recovery process. If you back up the operating system, only a test recovery in bootable environment to a spare hard drive can guarantee success of the recovery. At least ensure that the backup can be successfully validated using the bootable media.

Different ways to create a validation task

Using the Validation page is the most general way to create a validation task. Here you can validate immediately or set up a validation schedule for any backup, archive or location 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 the Creating a backup plan (p. 116) section.

You can access the **Validation** page from the **Vaults** (p. 80) view. Right-click the object to validate (archive, backup or vault) and 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.

General

Task name

[Optional] Enter a unique name for the validation task. A conscious name lets you quickly identify the task among the others.

Credentials (p. 161)

[Optional] The validation task will run on behalf of the user who is creating the task. You can change the task credentials if necessary. To access this option, select the **Advanced view** check box.

What to validate

Validate

Choose an object to validate:

Archive (p. 161) - in that case, you need to specify the archive.

Backup (p. 162) - specify the archive first, and then select the desired backup in this archive.

Vault (p. 162) - select a vault (or other location), which archives to validate.

Access Credentials (p. 163)

[Optional] Provide credentials for accessing the source if the task account does not have enough privileges to access it. To access this option, select the check box for **Advanced view**.

When to validate

Validate (p. 163)

Specify when and how often to perform validation.

After you configure all the required settings, click **OK** to create the validation task.

6.4.1 Task credentials

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

To specify credentials

1. Select one of the following:

- **Run under the current user**

The task will run under the credentials with which the user who starts the tasks is logged on. If the task has to run on schedule, you will be asked for the current user's password on completing the task creation.

- **Use the following credentials**

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

Specify:

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

2. Click **OK**.

To learn more about using credentials in Acronis Backup & Recovery 10, see the Owners and credentials (p. 24) section.

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

6.4.2 Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field, or select the desired folder in the folders tree.
 - If the archive is stored in Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then expand the **Online backup storage** group and select the account.

Exporting and mounting are not supported for backups stored in Acronis Online Backup Storage.

- If the archive is stored in a centralized vault, expand the **Centralized** group and click the vault.
- If the archive is stored in a personal vault, expand the **Personal** group and click the vault.

- If the archive is stored in a local folder on the machine, expand the **Local folders** group and click the required folder.

If the archive is located on removable media, e.g. DVDs, first insert the last DVD and then insert the discs in order starting from the first one when the program prompts.

- If the archive is stored on a network share, expand the **Network folders** group, then select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.

- If the archive is stored on an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

ftp://ftp_server:port_number or sftp://sftp_server:port number

If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

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

- If the archive is stored on a locally attached tape device, expand the **Tape drives** group, then click the required device.

When operating on a machine booted with bootable media:

- To access a managed vault, type the following string in the **Path** field:

bsp://node_address/vault_name/

- To access an unmanaged centralized vault, type the full path to the vault's folder.

2. In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each vault/folder you select.

While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3. Click **OK**.

6.4.3 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**.

6.4.4 Location selection

To select a location

Enter the full path to the location in the **Path** field or select the desired location in the **folders tree**.

- To select a centralized vault, expand the **Centralized** group and click the appropriate vault.
- To select a personal vault, expand the **Personal** group and click the appropriate vault.
- To select a local folder (CD/DVD drive, or locally attached tape device), expand the **Local folders** group and click the required folder.
- To select a network share, expand the **Network folders** group, select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.
- To select **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.

Using the archives table

To assist you with choosing the right location, 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.

6.4.5 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 program will access the location using the credentials of the task account specified in the General section.

- **Use the following credentials**

The program will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

Specify:

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

2. Click **OK**.

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

6.4.6 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. 87) section.

6.5 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.

Limitation: Mounting of volume backups stored on Acronis Backup & Recovery 10 Storage Node is not possible.

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 give 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. 165)

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

Backup (p. 166)

Select the backup.

Access credentials (p. 166)

[Optional] Provide credentials for the archive location. To access this option, select the **Advanced view** check box.

Mount settings

Volumes (p. 166)

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.

6.5.1 Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field, or select the desired folder in the folders tree.
 - If the archive is stored in Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then expand the **Online backup storage** group and select the account.

Exporting and mounting are not supported for backups stored in Acronis Online Backup Storage.

- If the archive is stored in a centralized vault, expand the **Centralized** group and click the vault.
- If the archive is stored in a personal vault, expand the **Personal** group and click the vault.
- If the archive is stored in a local folder on the machine, expand the **Local folders** group and click the required folder.

If the archive is located on removable media, e.g. DVDs, first insert the last DVD and then insert the discs in order starting from the first one when the program prompts.

- If the archive is stored on a network share, expand the **Network folders** group, then select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.

- If the archive is stored on an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

ftp://ftp_server:port_number or sftp://sftp_server:port number

If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

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

- If the archive is stored on a locally attached tape device, expand the **Tape drives** group, then click the required device.

When operating on a machine booted with bootable media:

- To access a managed vault, type the following string in the **Path** field:
bsp://node_address/vault_name/
 - To access an unmanaged centralized vault, type the full path to the vault's folder.
2. In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each vault/folder you select.

While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3. Click **OK**.

6.5.2 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**.

6.5.3 Access credentials

To specify credentials

1. Select one of the following:

- **Use the current user credentials**

The program will access the location using the credentials of the current user.

- **Use the following credentials**

The program will access the location using the credentials you specify. Use this option if the current user account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

Specify:

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

2. Click **OK**.

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

6.5.4 Volume selection

Select the volumes to mount and configure the mounting parameters for each of the selected volumes as follows:

1. Select the check box for each volume you need to mount.
2. Click on the selected volume to set its mounting parameters.
 - **Access mode** - choose the mode you want the volume to be mounted in:
 - **Read only** - enables exploring and opening files within the backup without committing any changes.
 - **Read/write** - with this mode, the program assumes that the backup content will be modified, and creates an incremental backup to capture the changes.
 - **Assign letter** (in Windows) - Acronis Backup & Recovery 10 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**.

6.6 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**.

6.7 Exporting archives and backups

The export operation creates a copy of an archive or a self-sufficient part copy of an archive in the location you specify. The original archive remains untouched.

The export operation can be applied to:

- **a single archive** - an exact archive copy will be created
- **a single backup** - an archive consisting of a single full backup will be created. The export of an incremental or a differential backup is performed using consolidation of the preceding backups up to the nearest full backup
- **your choice of backups** belonging to the same archive - the resulting archive will contain only the specified backups. Consolidation is performed as required, so the resulting archive may contain full, incremental and differential backups.

Usage scenarios

Export enables you to separate a specific backup from a chain of incremental backups for fast recovery, writing onto removable or detachable media or other purposes.

Example. When backing up data to a remote location through an unstable or low-bandwidth network connection (such as backing up through WAN using VPN access), you may want to save the initial full backup to a detachable media. Then, send the media to the remote location. There, the backup will

be exported from the media to the target storage. Subsequent incremental backups, which are usually much smaller, can be transferred over the network.

By exporting a managed vault to a detachable media, you obtain a portable unmanaged vault that can be used in the following scenarios:

- Keeping an off-site copy of your vault or of the most important archives.
- Physical transportation of a vault to a distant branch office.
- Recovery without access to the storage node in case of networking problems or failure of the storage node.
- Recovery of the storage node itself.

Export from an HDD-based vault to a tape device can be considered as simple on-demand archive staging.

The resulting archive's name

By default, the exported archive inherits the name of the original archive. Because having multiple archives of the same names in the same location is not advisable, the following actions are disabled with the default archive name:

- Exporting part of an archive to the same location.
- Exporting an archive or part of an archive to a location where an archive of the same name exists.
- Exporting an archive or part of an archive to the same location twice.

In any of the above cases, provide an archive name that is unique to the destination folder or vault. If you need to redo the export using the same archive name, first delete the archive that resulted from the previous export operation.

The resulting archive's options

The exported archive inherits the options of the original archive, including encryption and the password. When exporting a password-protected archive, you are prompted for the password. If the original archive is encrypted, the password is used to encrypt the resulting archive.

Source and destination locations

When the console is connected to a **managed machine**, you can export an archive or part of an archive to and from any location accessible to the agent residing on the machine. These include personal vaults, locally attached tape devices, removable media and, in the advanced product versions, managed and unmanaged centralized vaults.

When the console is connected to a **management server**, two export methods are available:

- Export from a **managed vault**. The export is performed by the storage node that manages the vault. The destination can be a network share or a local folder of the storage node.
- Export from an **unmanaged centralized vault**. The export is performed by the agent installed on the managed machine you specify. The destination can be any location accessible to the agent, including a managed vault.

Tip. When configuring export to a deduplicating managed vault, choose a machine where the deduplication add-on to the agent is installed. Otherwise the export task will fail.

Operations with an export task

An export task starts immediately after you complete its configuration. An export task can be stopped or deleted in the same way as any other task.

Once the export task is completed, you can run it again at any time. Before doing so, delete the archive that resulted from the previous task run if the archive still exists in the destination vault. Otherwise the task will fail. You cannot edit an export task to specify another name for the destination archive (this is a limitation).

Tip. You can implement the staging scenario manually, by regularly running the archive deletion task followed by the export task.

Different ways to create an export task

Using the **Export** page is the most general way to create an export task. Here, you can export any backup, or archive you have permission to access.

You can access the **Export** page from the **Vaults** view. Right-click the object to export (archive or backup) and select **Export** from the context menu. 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.

General

Task name

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

Task credentials (p. 170)

[Optional] The export task will run on behalf of the user who is creating the task. You can change the task credentials if necessary. To access this option, select the **Advanced view** check box.

What to export

Export

Select an object to export:

Archive (p. 143) - in that case, you need to specify the archive only.

Backups (p. 171) - specify the archive first, and then select the desired backup(s) in this archive

Access credentials (p. 171)

[Optional] Provide credentials for accessing the source if the task account does not have enough privileges to access it. To access this option, select the **Advanced view** check box.

Where to export

Archive (p. 172)

Enter 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.

Access credentials (p. 173)

[Optional] Provide credentials for the destination if the task credentials do not have enough privileges to access it. To access this option, select the **Advanced view** check box.

After you have performed all the required steps, click **OK** to start the export task.

6.7.1 Task credentials

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

To specify credentials

1. Select one of the following:

- **Run under the current user**

The task will run under the credentials with which the user who starts the tasks is logged on. If the task has to run on schedule, you will be asked for the current user's password on completing the task creation.

- **Use the following credentials**

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

Specify:

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

2. Click **OK**.

To learn more about using credentials in Acronis Backup & Recovery 10, see the Owners and credentials (p. 24) section.

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

6.7.2 Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field, or select the desired folder in the folders tree.
 - If the archive is stored in Acronis Online Backup Storage, click **Log in** and specify the credentials to log in to the online storage. Then expand the **Online backup storage** group and select the account.

Exporting and mounting are not supported for backups stored in Acronis Online Backup Storage.

- If the archive is stored in a centralized vault, expand the **Centralized** group and click the vault.
- If the archive is stored in a personal vault, expand the **Personal** group and click the vault.
- If the archive is stored in a local folder on the machine, expand the **Local folders** group and click the required folder.

If the archive is located on removable media, e.g. DVDs, first insert the last DVD and then insert the discs in order starting from the first one when the program prompts.

- If the archive is stored on a network share, expand the **Network folders** group, then select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.

- If the archive is stored on an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

ftp://ftp_server:port_number or sftp://sftp_server:port number

If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

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

- If the archive is stored on a locally attached tape device, expand the **Tape drives** group, then click the required device.

When operating on a machine booted with bootable media:

- To access a managed vault, type the following string in the **Path** field:

bsp://node_address/vault_name/

- To access an unmanaged centralized vault, type the full path to the vault's folder.

2. In the table to the right of the tree, select the archive. The table displays the names of the archives contained in each vault/folder you select.

While you are reviewing the location content, archives can be added, deleted or modified by another user or by the program itself according to scheduled operations. Use the **Refresh** button to refresh the list of archives.

3. Click **OK**.

6.7.3 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**.

6.7.4 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 task credentials**

The program will access the location using the credentials of the task account specified in the General section.

- **Use the following credentials**

The program will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

Specify:

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

2. Click **OK**.

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

6.7.5 Location 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 folders tree.

- To export data to a centralized unmanaged vault, expand the **Centralized vaults** group and click the vault.
- To export data to a personal vault, expand the **Personal vaults** group and click the vault.
- To export data to a local folder on the machine, expand the **Local folders** group and click the required folder.
- To export data to a network share, expand the **Network folders** group, select the required networked machine and then click the shared folder. If the network share requires access credentials, the program will ask for them.

Note for Linux users: To specify a Common Internet File System (CIFS) network share which is mounted on a mount point such as /mnt/share, select this mount point instead of the network share itself.

- To export data to an **FTP** or **SFTP** server, type the server name or address in the **Path** field as follows:

ftp://ftp_server:port _number or **sftp://sftp_server:port number**

If the port number is not specified, port 21 is used for FTP and port 22 is used for SFTP.

After entering access credentials, the folders on the server become available. Click the appropriate folder on the server.

You can access the server as an anonymous user if the server enables such access. To do so, click **Use anonymous access** instead of entering credentials.

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

- To export data to a locally attached tape device, expand the **Tape drives** group, then click the required device.

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.

6.7.6 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 task credentials**

The program will access the location using the credentials of the task account specified in the General section.

- **Use the following credentials**

The program will access the location using the credentials you specify. Use this option if the task account does not have access permissions to the location. You might need to provide special credentials for a network share or a storage node vault.

Specify:

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

2. Click **OK**.

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

6.8 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.

Certain Windows applications, such as Acronis disk management tools, can access the zone.

To learn more about the advantages and limitations of the Acronis Secure Zone, see the Acronis Secure Zone (p. 39) topic in the "Proprietary Acronis technologies" section.

6.8.1 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.

Space

Disk (p. 174)

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. 175)

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.

Settings

Password (p. 175)

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

6.8.1.1 Acronis Secure Zone Disk

The Acronis Secure Zone can be located on any fixed hard drive. Acronis Secure Zone is always created at the end of the hard disk. A machine can have only one Acronis Secure Zone. Acronis Secure Zone is created using unallocated space, if available, or at the expense of the volumes' free space.

The Acronis Secure Zone cannot be organized on a dynamic disk or a disk using the GPT partitioning style.

To allocate space for Acronis Secure Zone

1. Choose a hard disk (if several) on which to create the zone. The unallocated space is 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. 175) window.
3. Click **OK**.

6.8.1.2 Acronis Secure Zone Size

Enter the Acronis Secure Zone size or drag the slider to select any size between the minimum and the maximum ones. The minimum size is approximately 50MB, depending on the geometry of the hard disk. The maximum size is equal to the disk's unallocated space plus the total free space on all the volumes you have selected in the previous step.

If you have to take space from the boot or the system volume, please bear the following in mind:

- Moving or resizing of the volume from which the system is currently booted will require a reboot.
- Taking all free space from a system volume may cause the operating system to work unstably and even fail to start. Do not set the maximum zone size if the boot or the system volume is selected.

6.8.1.3 Password for Acronis Secure Zone

Setting up a password protects the Acronis Secure Zone from unauthorized access. The program will ask for the password at any operation relating to the zone and the archives located there, such as data backup and recovery, validating archives, resizing and deleting the zone.

To set up a password

1. Choose **Use password**.
2. In the **Enter the password** field, type a new password.
3. In the **Confirm the password** field, re-type the password.
4. Click **OK**.

To disable password

1. Choose **Do not use**.
2. Click **OK**.

6.8.1.4 Result confirmation

The **Result confirmation** window displays the expected partition layout according to the settings you have chosen. Click **OK**, if you are satisfied with the layout and the Acronis Secure Zone creation will start.

How the settings you make will be processed

This helps you to understand how creating the Acronis Secure Zone will transform a disk containing multiple volumes.

- Acronis Secure Zone is always created at the end of the hard disk. When calculating the final layout of the volumes, the program will first use unallocated space at the end.
- If there is no or not enough unallocated space at the end of the disk, but there is unallocated space between volumes, the volumes will be moved to add more unallocated space to the end.
- When all unallocated space is collected but it is still not enough, the program will take free space from the volumes you select, proportionally reducing the volumes' size. Resizing of locked volumes requires a reboot.
- However, there should be free space on a volume, so that the operating system and applications can operate; for example, for creating temporary files. The program will not decrease a volume where free space is or becomes less than 25% of the total volume size. Only when all volumes on the disk have 25% or less free space, will the program continue decreasing the volumes proportionally.

As is apparent from the above, setting the maximum possible zone size is not advisable. You will end up with no free space on any volume which might cause the operating system or applications to work unstably and even fail to start.

6.8.2 Managing Acronis Secure Zone

Acronis Secure Zone is considered as a personal vault (p. 241). Once created on a managed machine, the zone is always present in the list of **Personal vaults**. Centralized backup plans can use Acronis Secure Zone as well as local plans.

If you have used the Acronis Secure Zone before, please note a radical change in the zone functionality. The zone does not perform automatic cleanup, that is, deleting old archives, anymore. Use backup schemes with automatic cleanup to back up to the zone, or delete outdated archives manually using the vault management functionality.

With the new Acronis Secure Zone behavior, you obtain the ability to:

- list archives located in the zone and backups included in each archive
- examine backup content
- mount a volume backup to copy files from the backup to a physical disk
- safely delete archives and backups from the archives.

To learn more about operations with vaults, see the Vaults (p. 80) section.

6.8.2.1 Increasing Acronis Secure Zone

To increase Acronis Secure Zone

1. On the **Manage Acronis Secure Zone** page, click **Increase**.
2. Select volumes from which free space will be used to increase the Acronis Secure Zone.
3. Specify the new size of the zone by:
 - dragging the slider and selecting any size between the current and maximum values. The maximum size is equal to the disk's unallocated space plus the total free space of all selected partitions;
 - typing an exact value in the Acronis Secure Zone Size field.

When increasing the size of the zone, the program will act as follows:

- first, it will use the unallocated space. Volumes will be moved, if necessary, but not resized. Moving of locked volumes requires a reboot.
- If there is not enough unallocated space, the program will take free space from the selected volumes, proportionally reducing the volumes' size. Resizing of locked partitions requires a reboot.

Reducing a system volume to the minimum size might prevent the machine's operating system from booting.

4. Click **OK**.

6.8.2.2 Decreasing Acronis Secure Zone

To decrease Acronis Secure Zone

1. On the **Manage Acronis Secure Zone** page, click **Decrease**.
2. Select volumes that will receive free space after the zone is decreased.

3. Specify the new size of the zone by:
 - dragging the slider and selecting any size between the current and minimum values. The minimum size is approximately 50MB, depending on the geometry of the hard disk;
 - typing an exact value in the **Acronis Secure Zone Size** field.
4. Click **OK**.

6.8.2.3 Deleting Acronis Secure Zone

To delete Acronis Secure Zone:

1. In the **Acronis Secure Zone Actions** bar (on the **Actions and tools** pane), select **Delete**.
2. In the **Delete Acronis Secure Zone** window, select volumes to which you want to add the space freed from the zone and then click **OK**.

If you select several volumes, the space will be distributed proportionally to each partition. If you do not select any volume, the freed space becomes unallocated.

After you click **OK**, Acronis Backup & Recovery 10 will start deleting the zone.

6.9 Acronis Startup Recovery Manager

Acronis Startup Recovery Manager is a modification of the bootable agent (p. 231), 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.

Activate

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). If the system fails to boot, you will be able to start the bootable rescue utility, by pressing F11 or by selecting it from the menu, respectively.

The system disk (or, the /boot partition in Linux) should have at least 70 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 ASRM. 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).

See the Bootable media (p. 178) section for details.

6.10 Bootable media

Bootable media

Bootable media is physical media (CD, DVD, USB drive or other media supported by a machine BIOS as a boot device) that boots on any PC-compatible machine and enables you to run Acronis Backup & Recovery 10 Agent either in a Linux-based environment or Windows Preinstallation Environment (WinPE), without the help of an operating system. Bootable media is most often used to:

- recover an operating system that cannot start
- access and back up the data that has survived in a corrupted system
- deploy an operating system on bare metal
- create basic or dynamic volumes on bare metal
- back up sector-by-sector a disk with an unsupported file system
- back up offline any data that cannot be backed up online because of restricted access, being permanently locked by the running applications or for any other reason.

A machine can be booted into the above environments either with physical media, or using the network boot from Acronis PXE Server, Windows Deployment Services (WDS) or Remote Installation Services (RIS). These servers with uploaded bootable components can be thought of as a kind of bootable media too. You can create bootable media or configure the PXE server or WDS/RIS using the same wizard.

Linux-based bootable media

Linux-based media contains Acronis Backup & Recovery 10 Bootable Agent based on Linux kernel. The agent can boot and perform operations on any PC-compatible hardware, including bare metal and machines with corrupted or non-supported file systems. The operations can be configured and controlled either locally or remotely using the management console.

PE-based bootable media

PE-based bootable media contains a minimal Windows system called Windows Preinstallation Environment (WinPE) and Acronis Plug-in for WinPE, that is, a modification of Acronis Backup & Recovery 10 Agent that can run in the preinstallation environment.

WinPE proved to be the most convenient bootable solution in large environments with heterogeneous hardware.

Advantages:

- Using Acronis Backup & Recovery 10 in Windows Preinstallation Environment provides more functionality than using Linux-based bootable media. Having booted PC-compatible hardware into WinPE, you can use not only Acronis Backup & Recovery 10 Agent, but also PE commands and scripts and other plug-ins you've added to the PE.
- PE-based bootable media helps overcome some Linux-related bootable media issues such as support for certain RAID controllers or certain levels of RAID arrays only. Media based on PE 2.x, that is, Windows Vista or Windows Server 2008 kernel, allows for dynamic loading of the necessary device drivers.

6.10.1 How to create bootable media

To enable creating physical media, the machine must have a CD/DVD recording drive or allow a flash drive to be attached. To enable PXE or WDS/RIS configuration, the machine must have a network connection. Bootable Media Builder can also create an ISO image of a bootable disk to burn it later on a blank disk.

Linux-based bootable media

Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component.

Select the way volumes and network resources will be handled—called the media style:

- A media with Linux-style volume handling displays the volumes as, for example, hda1 and sdb2. It tries to reconstruct MD devices and logical (LVM) volumes before starting a recovery.
- A media with Windows-style volume handling displays the volumes as, for example, C: and D:. It provides access to dynamic (LDM) volumes.

The wizard will guide you through the necessary operations. Please refer to Linux-based bootable media (p. 180) for details.

PE-based bootable media

Acronis Plug-in for WinPE can be added to WinPE distributions based on any of the following kernels:

- Windows XP Professional with Service Pack 2 (PE 1.5)
- Windows Server 2003 with Service Pack 1 (PE 1.6)
- Windows Vista (PE 2.0)
- Windows Vista SP1 and Windows Server 2008 (PE 2.1)
- Windows 7 (PE 3.0)

If you already have media with PE1.x distribution, unpack the media ISO to a local folder and start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component. The wizard will guide you through the necessary operations. Please refer to Adding the Acronis Plug-in to WinPE 1.x (p. 184) for details.

To be able to create or modify PE 2.x or 3.0 images, install Bootable Media Builder on a machine where Windows Automated Installation Kit (AIK) is installed. The further operations are described in the Adding the Acronis Plug-in to WinPE 2.x or 3.0 (p. 185) section.

If you do not have a machine with WAIK, prepare as follows:

1. Download and install Windows Automated Installation Kit (WAIK).

Automated Installation Kit (AIK) for Windows Vista (PE 2.0):

<http://www.microsoft.com/Downloads/details.aspx?familyid=C7D4BC6D-15F3-4284-9123-679830D629F2&displaylang=en>

Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008 (PE 2.1):

<http://www.microsoft.com/downloads/details.aspx?FamilyID=94bb6e34-d890-4932-81a5-5b50c657de08&DisplayLang=en>

Automated Installation Kit (AIK) for Windows 7 (PE 3.0):

<http://www.microsoft.com/downloads/details.aspx?familyid=696DD665-9F76-4177-A811-39C26D3B3B34&displaylang=en>

You can find system requirements for installation by following the above links.

2. [optional] Burn the WAIK to DVD or copy to a flash drive.
3. Install the Microsoft .NET Framework v.2.0 from this kit (NETFXx86 or NETFXx64, depending on your hardware.)
4. Install Microsoft Core XML (MSXML) 5.0 or 6.0 Parser from this kit.
5. Install Windows AIK from this kit.
6. Install Bootable Media Builder on the same machine.

It is recommended that you familiarize yourself with the help documentation supplied with Windows AIK. To access the documentation, select **Microsoft Windows AIK -> Documentation** from the start menu.

Using Bart PE

You can create a Bart PE image with Acronis Plug-in using the Bart PE Builder. Please refer to Building Bart PE with Acronis Plug-in from Windows distribution (p. 186) for details.

6.10.1.1 Linux-based bootable media

When using the media builder, you have to specify:

1. [optional] The parameters of the Linux kernel. Separate multiple parameters with spaces.
For example, to be able to select a display mode for the bootable agent each time the media starts, type: **vga=ask**
For a list of parameters, see Kernel parameters (p. 181).
2. The Acronis bootable components to be placed on the media.
 - Universal Restore can be enabled if Acronis Backup & Recovery 10 Universal Restore is installed on the machine where the media is created.
3. [optional] The timeout interval for the boot menu plus the component that will automatically start on timeout.
 - If not configured, the Acronis loader waits for someone to select whether to boot the operating system (if present) or the Acronis component.
 - If you set, say, **10 sec.** for the bootable agent, the agent will launch 10 seconds after the menu is displayed. This enables unattended onsite operation when booting from a PXE server or WDS/RIS.
4. [optional] Remote logon settings:
 - user name and password to be entered on the console side at connection to the agent. If you leave these fields empty, the connection will be enabled on typing any symbols in the prompt window.
5. [optional] Network settings (p. 182):
 - TCP/IP settings to be assigned to the machine network adapters.
6. [optional] Network port (p. 183):
 - the TCP port that the bootable agent listens for incoming connection.
7. The type of media to create. You can:
 - create CD, DVD or other bootable media such as removable USB flash drives if the hardware BIOS allows for boot from such media
 - build an ISO image of a bootable disc to burn it later on a blank disc
 - upload the selected components to Acronis PXE Server

- upload the selected components to a WDS/RIS.
8. [optional] Windows system drivers to be used by Acronis Universal Restore (p. 183). This window appears only if the Acronis Universal Restore add-on is installed and a media other than PXE or WDS/RIS is selected.
 9. Path to the media ISO file or the name or IP and credentials for PXE or WDS/RIS.

Kernel parameters

This window lets you specify one or more parameters of the Linux kernel. They will be automatically applied when the bootable media starts.

These parameters are typically used when experiencing problems while working with the bootable media. Normally, you can leave this field empty.

You also can specify any of these parameters by pressing F11 while in the boot menu.

Parameters

When specifying multiple parameters, separate them with spaces.

acpi=off

Disables Advanced Configuration and Power Interface (ACPI). You may want to use this parameter when experiencing problems with a particular hardware configuration.

noapic

Disables Advanced Programmable Interrupt Controller (APIC). You may want to use this parameter when experiencing problems with a particular hardware configuration.

vga=ask

Prompts for the video mode to be used by the bootable media's graphical user interface. Without the **vga** parameter, the video mode is detected automatically.

vga=mode_number

Specifies the video mode to be used by the bootable media's graphical user interface. The mode number is given by *mode_number* in the hexadecimal format—for example: **vga=0x318**

Screen resolution and the number of colors corresponding to a mode number may be different on different machines. We recommend using the **vga=ask** parameter first to choose a value for *mode_number*.

quiet

Disables displaying of startup messages when the Linux kernel is loading, and starts the management console after the kernel is loaded.

This parameter is implicitly specified when creating the bootable media, but you can remove this parameter while in the boot menu.

Without this parameter, all startup messages will be displayed, followed by a command prompt. To start the management console from the command prompt, run the command: **/bin/product**

nousb

Disables loading of the USB (Universal Serial Bus) subsystem.

nousb2

Disables USB 2.0 support. USB 1.1 devices still work with this parameter. This parameter allows you to use some USB drives in the USB 1.1 mode if they do not work in the USB 2.0 mode.

nodma

Disables direct memory access (DMA) for all IDE hard disk drives. Prevents the kernel from freezing on some hardware.

nofw

Disables the FireWire (IEEE1394) interface support.

nopcmcia

Disables detection of PCMCIA hardware.

nomouse

Disables mouse support.

module_name=off

Disables the module whose name is given by *module_name*. For example, to disable the use of the SATA module, specify: **sata_sis=off**

pci=bios

Forces the use of PCI BIOS instead of accessing the hardware device directly. You may want to use this parameter if the machine has a non-standard PCI host bridge.

pci=nobios

Disables the use of PCI BIOS; only direct hardware access methods will be allowed. You may want to use this parameter when the bootable media fails to start, which may be caused by the BIOS.

pci=biosirq

Uses PCI BIOS calls to get the interrupt routing table. You may want to use this parameter if the kernel is unable to allocate interrupt requests (IRQs) or discover secondary PCI buses on the motherboard.

These calls might not work properly on some machines. But this may be the only way to get the interrupt routing table.

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.

Network port

While creating bootable media, you have an option to pre-configure the network port that the bootable agent listens for incoming connection. The choice is available between:

- the default port
- the currently used port
- the new port (enter the port number).

If the port has not been pre-configured, the agent uses the default port number (9876.) This port is also used as default by the Acronis Backup & Recovery 10 Management Console.

Drivers for Universal Restore

While creating bootable media, you have an option to add Windows drivers to the media. The drivers will be used by Universal Restore when recovering Windows on a machine with a dissimilar processor, different motherboard or different mass storage device than in the backed up system.

You will be able to configure the Universal Restore:

- to search the media for the drivers that best fit the target hardware

- to get the mass-storage drivers that you explicitly specify from the media. This is necessary when the target hardware has a specific mass storage controller (such as a SCSI, RAID, or Fiber Channel adapter) for the hard disk.

For more information please refer to Universal Restore (p. 153).

The drivers will be placed in the visible Drivers folder on the bootable media. The drivers are not loaded into the target machine RAM, therefore, the media must stay inserted or connected throughout the Universal Restore operation.

Adding drivers to bootable media is available on the condition that:

1. The Acronis Backup & Recovery 10 Universal Restore add-on is installed on the machine where the bootable media is created AND
2. You are creating a removable media or its ISO or detachable media, such as a flash drive. Drivers cannot be uploaded on a PXE server or WDS/RIS.

The drivers can be added to the list only in groups, by adding the INF files or folders containing such files. Selecting individual drivers from the INF files is not possible, but the media builder shows the file content for your information.

To add drivers:

1. Click **Add** and browse to the INF file or a folder that contains INF files.
2. Select the INF file or the folder.
3. Click **OK**.

The drivers can be removed from the list only in groups, by removing INF files.

To remove drivers:

1. Select the INF file.
2. Click **Remove**.

6.10.1.2 Adding the Acronis Plug-in to WinPE 1.x

Acronis Plug-in for WinPE can be added to:

- Windows PE 2004 (1.5) (Windows XP Professional with Service Pack 2)
- Windows PE 2005 (1.6) (Windows Server 2003 with Service Pack 1).

To add Acronis Plug-in to WinPE 1.x:

1. Unpack all files of your WinPE 1.x ISO to a separate folder on the hard disk.
2. Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component.
3. Select **Bootable media type: Windows PE**.
 - Select **Use WinPE files located in the folder I specify**
4. Specify path to the folder where the WinPE files are located.
5. Specify network settings (p. 182) for the machine network adapters or choose DHCP auto configuration.
6. Specify the full path to the resulting ISO file including the file name.
7. Check your settings in the summary screen and click **Proceed**.
8. Burn the .ISO to CD or DVD using a third-party tool or copy to a flash drive.

Once a machine boots into the WinPE, Acronis Backup & Recovery 10 starts automatically.

6.10.1.3 Adding the Acronis Plug-in to WinPE 2.x or 3.0

Bootable Media Builder provides three methods of integrating Acronis Backup & Recovery 10 with WinPE 2.x or 3.0:

- Adding the Acronis Plug-in to the existing PE ISO. This comes in handy when you have to add the plug-in to the previously configured PE ISO that is already in use.
- Creating the PE ISO with the plug-in from scratch.
- Adding the Acronis Plug-in to a WIM file for any future purpose (manual ISO building, adding other tools to the image and so on).

To be able to perform any of the above operations, install Bootable Media Builder on a machine where Windows Automated Installation Kit (WAIK) is installed. If you do not have such machine, prepare as described in How to create bootable media (p. 179).

Bootable Media Builder supports only x86 WinPE 2.x or 3.0. These WinPE distributions can also work on x64 hardware.

A PE image based on Win PE 2.0 requires at least 256MB RAM to work. The recommended memory size for PE 2.0 is 512MB. A PE image based on Win PE 3.0 requires at least 512MB RAM to work.

Adding Acronis Plug-in to WinPE 2.x or 3.0 ISO

To add Acronis Plug-in to WinPE 2.x or 3.0 ISO:

1. When adding the plug-in to the existing Win PE ISO, unpack all files of your Win PE ISO to a separate folder on the hard disk.
2. Start the Bootable Media Builder either from the management console, by selecting **Tools > Create Bootable Media** or, as a separate component.
3. Select **Bootable media type: Windows PE**.

When creating a new PE ISO:

- Select **Create Windows PE 2.x or 3.0 automatically**
- The software runs the appropriate script and proceeds to the next window.

When adding the plug-in to the existing PE ISO:

- Select **Use WinPE files located in the folder I specify**
- Specify path to the folder where the WinPE files are located.

4. Specify network settings (p. 182) for the machine network adapters or choose DHCP auto configuration.
5. [optional] Specify Windows drivers to be added to Windows PE. Once you boot a machine into Windows PE, the drivers can help you access the device where the backup archive is located. Click **Add** and specify the path to the necessary *.inf file for a corresponding SCSI, RAID, SATA controller, network adapter, tape drive or other device. You will have to repeat this procedure for each driver you want to be included in the resulting WinPE boot media.
6. Choose whether you want to create ISO or WIM image or upload the media on Acronis PXE Server.
7. Specify the full path to the resulting image file including the file name, or specify the PXE server and provide the user name and password to access it.
8. Check your settings in the summary screen and click **Proceed**.
9. Burn the .ISO to CD or DVD using a third-party tool or copy to a flash drive.

Once a machine boots into WinPE, Acronis Backup & Recovery 10 starts automatically.

To create a PE image (ISO file) from the resulting WIM file:

- replace the default boot.wim file in your Windows PE folder with the newly created WIM file. For the above example, type:

```
copy c:\AcronisMedia.wim c:\winpe_x86\ISO\sources\boot.wim
```

- use the **Oscdimg** tool. For the above example, type:

```
oscdimg -n -bc:\winpe_x86\etfsboot.com c:\winpe_x86\ISO  
c:\winpe_x86\winpe_x86.iso
```

For more information on customizing Windows PE, see the Windows Preinstallation Environment User's Guide (Winpe.chm).

6.10.1.4 Building Bart PE with Acronis Plug-in from Windows distribution

1. Get the Bart PE builder.
2. Install Bootable Media Builder from the Acronis Backup & Recovery 10 setup file.
3. Change the current folder to the folder where the Acronis Plug-in for WinPE is installed—by default: C:\Program Files\Acronis\Bootable Components\WinPE.
If the plug-in is installed in a folder other than the default folder, change the path accordingly (check the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\Bootable Components\Settings\WinPE for the plug-in location).
4. Unpack the WinPE.zip file to the current folder.
5. Run the following command:

```
export_license.bat
```
6. Copy the contents of the current folder—by default: C:\Program Files\Acronis\Bootable Components\WinPE—to the %BartPE folder%\plugins\Acronis.
7. Insert your Windows distribution CD if you do not have a copy of Windows installation files on the HDD.
8. Start the Bart PE builder.
9. Specify the path to the Windows installation files or Windows distribution CD.
10. Click **Plugins** and check whether the Acronis Backup & Recovery 10 plug-in is enabled. Enable if disabled.
11. Specify the output folder and the full path to the resulting ISO file including the file name or the media to create.
12. Build the Bart PE.
13. Burn the ISO to CD or DVD (if this has not been done yet) or copy to a flash drive.

Once the machine boots into the Bart PE and you configure the network connection, select **Go -> System -> Storage -> Acronis Backup & Recovery 10** to start.

6.10.2 Connecting to a machine booted from media

Once a machine boots from bootable media, the machine terminal displays a startup window with the IP address(es) obtained from DHCP or set according to the pre-configured values.

Remote connection

To connect to the machine remotely, select **Connect -> Manage a remote machine** in the console menu and specify one of the machine's IP addresses. Provide the user name and password if these have been configured when creating the bootable media.

Local connection

Acronis Backup & Recovery 10 Management Console is always present on the bootable media. Anyone who has physical access to the machine terminal can run the console and connect. Just click **Run management console** in the bootable agent startup window.

6.10.3 Working under bootable media

Operations on a machine booted with bootable media are very similar to backup and recovery under the operating system. The difference is as follows:

1. Disk letters seen under Windows-style bootable media might differ from the way Windows identifies drives. For example, the D: drive under the rescue utility might correspond to the E: drive in Windows.

Be careful! To be on the safe side, it is advisable to assign unique names to the volumes.

2. The Linux-style bootable media shows local disks and volumes as unmounted (sda1, sda2...).
3. 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.
4. You can switch the bootable media between the Windows style and the Linux style by selecting **Tools > Change volume representation**.
5. There is no **Navigation** tree in the media GUI. Use the **Navigation** menu item to navigate between views.
6. Tasks cannot be scheduled; in fact, tasks are not created at all. If you need to repeat the operation, configure it from scratch.
7. The log lifetime is limited to the current session. You can save the entire log or the filtered log entries to a file.
8. Centralized vaults are not displayed in the folder tree of the **Archive** window.

To access a managed vault, type the following string in the **Path** field:

bsp://node_address/vault_name/

To access an unmanaged centralized vault, type the full path to the vault's folder.

After entering access credentials, you will see a list of archives located in the vault.

6.10.3.1 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. 180) section for details).

6.10.3.2 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**.

6.10.4 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

- `acronis`
- `asamba`
- `lash`
- `restoreraids`
- `trueimagecmd`
- `trueimagemnt`

Linux commands and utilities

busybox	ifconfig	rm
cat	init	rmmod
cdrecord	insmod	route
chmod	iscsiadm	scp
chown	kill	scsi_id
chroot	kpartx	sed
cp	ln	sg_map26
dd	ls	sh
df	lspci	sleep
dmesg	lvm	ssh
dmraid	mdadm	sshd
e2fsck	mkdir	strace
e2label	mke2fs	swapoff
echo	mknod	swapon
egrep	mkswap	sysinfo
fdisk	more	tar
fsck	mount	tune2fs
fxload	mtx	udev
gawk	mv	udevinfo
gpm	pccardctl	udevstart
grep	ping	umount
growisofs	pktsetup	uuidgen
grub	poweroff	vconfig
gunzip	ps	vi
halt	raidautorun	zcat
hexdump	readcd	
hotplug	reboot	

6.10.5 Recovering MD devices and logical volumes

To recover MD devices, known as Linux Software RAID, and/or devices created by Logical Volume Manager (LVM), known as logical volumes, you need to create the corresponding volume structure before starting the recovery.

You can create the volume structure in either of the following ways:

- Automatically in Linux-based bootable media by using the management console or a script—see *Creating the volume structure automatically* (p. 190).
- Manually by using the **mdadm** and **lvm** utilities—see *Creating the volume structure manually* (p. 191).

6.10.5.1 Creating the volume structure automatically

Let's assume that you saved the volume structure to the `/etc/Acronis` directory and that the volume with this directory is included in the archive.

To recreate the volume structure in Linux-based bootable media, use either of the methods described below.

Caution: *As a result of the following procedures, the current volume structure on the machine will be replaced with the one stored in the archive. 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, each of its own size. If you replaced any of these disks between backup and recovery—or if you are recovering the volumes to a different machine—make sure that the new disk configuration includes enough disks whose sizes are at least those of the original disks.

To create the volume structure by using the management console

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 & Recovery 10 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.

To create the volume structure by using a script

1. Boot the machine from a Linux-based bootable media.
2. Click **Acronis Bootable Agent**. Then, click **Run management console**.
3. On the toolbar, click **Actions**, and then click **Start shell**. Alternatively, you can press CTRL+ALT+F2.
4. Run the **restoreraids.sh** script, specifying the full file name of the archive—for example:

```
/bin/restoreraids.sh  
smb://server/backups/linux_machine_2010_01_02_12_00_00_123D.tib
```
5. Return to the management console by pressing CTRL+ALT+F1, or by running the command:
/bin/product
6. Click **Recover**, then specify the path to the archive and any other required parameters, and then click **OK**.

If Acronis Backup & Recovery 10 could not create the volume structure (or if it is not present in the archive), create the structure manually.

6.10.5.2 Creating the volume structure manually

The following are 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 recover MD devices and logical volumes

1. Boot the machine from a Linux-based bootable media.
2. Click **Acronis Bootable Agent**. 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 **trueimagecmd** utility. Also, you can use the **trueimagemnt** 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: **lvm pvcreate**, **lvm vgcreate**, etc.

6. If you previously mounted the backup by using the **trueimagemnt** 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, or by running the command:
/bin/product
(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 does not work when connected to Acronis Backup & Recovery 10 Bootable Agent remotely, because the command shell is not available in this case.

Example

Suppose that you previously performed a disk 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 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      0qoQ4l-Vk7W-yDG3-uF1l-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. 188). For detailed descriptions of the **trueimagecmd** and **trueimagemnt** utilities, see the Acronis Backup & Recovery 10 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 **--list** command to list the volumes which are stored in the backup. For example:

```
trueimagecmd --list --filename:smb://server/backups/linux_machine.tib
```

The output will contain lines similar to the following:

Num	Idx	Partition	Flags	Start	Size	Type

Disk 1:						
		Table		0		Table
Disk 2:						
		Table		0		Table
...						
Dynamic & GPT Volumes:						
DYN1	4	my_volgroup-my_logvol		12533760		Ext2

You will need the volume's index, given in the **Idx** column, in the next step.

2. Use the **--mount** command, specifying the volume's index in the **-i** parameter. For example:

```
trueimagemnt --mount /mnt --filename smb://server/backups/linux_machine.tib -i 4
```

This command mounts the logical volume DYN1, whose index in the backup is 4, on the mount point /mnt.

To unmount a backup volume

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

```
trueimagemnt --unmount /mnt
```

6.11 Disk management

Acronis Disk Director Lite is a tool for preparing a machine disk/volume configuration for recovering the volume images saved by the Acronis Backup & Recovery 10 software.

Sometimes after the volume has been backed up and its image placed into a safe storage, the machine disk configuration might change due to a HDD replacement or hardware loss. In such case with the help of Acronis Disk Director Lite, the user has the possibility to recreate the necessary disk

configuration so that the volume image can be recovered exactly “as it was” or with any alteration of the disk or volume structure the user might consider necessary.

All operations on disks and volumes involve a certain risk of data damage. Operations on system, bootable or data volumes must be carried out very carefully to avoid potential problems with the booting process or hard disk data storage.

Operations with hard disks and volumes take a certain amount of time, and any power loss, unintentional turning off of the machine or accidental pressing of the Reset button during the procedure could result in volume damage and data loss.

All operations on volumes of dynamic disks in Windows XP and Windows 2000 require Acronis Managed Machine Service to be run under an account with administrator's rights.

Please take all necessary precautions (p. 194) to avoid possible data loss.

6.11.1 Basic precautions

To avoid any possible disk and volume structure damage or data loss, please take all necessary precautions and follow these simple rules:

1. Back up the disk on which volumes will be created or managed. Having your most important data backed up to another hard disk, network share or removable media will allow you to work on disk volumes being reassured that your data is safe.
2. Test your disk to make sure it is fully functional and does not contain bad sectors or file system errors.
3. Do not perform any disk/volume operations while running other software that has low-level disk access. Close these programs before running Acronis Disk Director Lite.

With these simple precautions, you will protect yourself against accidental data loss.

6.11.2 Running Acronis Disk Director Lite

You can run Acronis Disk Director Lite under Windows or start it from a bootable media.

Running Acronis Disk Director Lite under Windows

If you run Acronis Backup & Recovery 10 Management Console, and connect it to a managed machine, the **Disk management** view will be available in the **Navigation** tree of the console, with which you can start Acronis Disk Director Lite.

Running Acronis Disk Director Lite from a bootable media

You can run Acronis Disk Director Lite on a bare metal, on a machine that cannot boot or on a non-Windows machine. To do so, boot the machine from a bootable media (p. 231) created with the Acronis Bootable Media Builder; run the management console and then click **Disk Management**.

6.11.3 Choosing the operating system for disk management

On a machine with two or more operating systems, representation of disks and volumes depends on which operating system is currently running.

A volume may have a different letter in different Windows operating systems. For example, volume E: might appear as D: or L: when you boot another Windows operating system installed on the same

machine. (It is also possible that this volume will have the same letter E: under any Windows OS installed on the machine.)

A dynamic disk created in one Windows operating system is considered as a **Foreign Disk** in another Windows operating system or might be unsupported by this operating system.

When you need to perform a disk management operation on such machine, it is necessary to specify for which operating system the disk layout will be displayed and the disk management operation will be performed.

The name of the currently selected operating system is shown on the console toolbar after “**The current disk layout is for:**”. Click the OS name to select another operating system in the **Operating System Selection** window. Under bootable media, this window appears after clicking **Disk management**. The disk layout will be displayed according to the operating system you select.

6.11.4 "Disk management" view

Acronis Disk Director Lite is controlled through the **Disk management** view of the console.

The top part of the view contains a disks and volumes table enabling data sorting and columns customization and toolbar. The table presents the numbers of the disks, as well as assigned letter, label, type, capacity, free space size, used space size, file system, and status for each volume. The toolbar comprises of icons to launch the **Undo**, **Redo** and **Commit** actions intended for pending operations (p. 208).

The graphic panel at the bottom of the view also graphically depicts all the disks and their volumes as rectangles with basic data on them (label, letter, size, status, type and file system).

Both parts of the view also depict all unallocated disk space that can be used in volume creation.

Starting the operations

Any operation can be launched:

- From the volume or disk context menu (both in the table and the graphic panel)
- From the **Disk management** menu of the console
- From the **Operations** bar on the **Actions and Tools** pane

*Note that the list of available operations in the context menu, the **Disk management** menu, and the **Operations** bar depends on the selected volume or disk type. The same is true for unallocated space as well.*

Displaying operation results

The results of any disk or volume operation, you have just planned, are immediately displayed in the **Disk management** view of the console. For example, if you create a volume, it will be immediately shown in the table, as well as in graphical form at the bottom of the view. Any volume changes, including changing the volume letter or label, are also immediately displayed in the view.

6.11.5 Disk operations

Acronis Disk Director Lite includes the following operations that can be performed on disks:

- Disk Initialization (p. 196) - initializes the new hardware added to the system
- Basic disk cloning (p. 196) - transfers complete data from the source basic MBR disk to the target

- Disk conversion: MBR to GPT (p. 198) - converts an MBR partition table to GPT
- Disk conversion: GPT to MBR (p. 199) - converts a GPT partition table to MBR
- Disk conversion: Basic to Dynamic (p. 199) - converts a basic disk to dynamic
- Disk conversion: Dynamic to Basic (p. 200) - converts a dynamic disk to basic

The full version of Acronis Disk Director will provide more tools and utilities for working with disks.

Acronis Disk Director Lite must obtain exclusive access to the target disk. This means no other disk management utilities (like Windows Disk Management utility) can access it at that time. If you receive a message stating that the disk cannot be blocked, close the disk management applications that use this disk and start again. If you cannot determine which applications use the disk, close them all.

6.11.5.1 Disk initialization

If you add any new disk to your machine, Acronis Disk Director Lite will notice the configuration change and scan the added disk to include it to the disk and volume list. If the disk is still not initialized or, possibly, has a file structure unknown to the machine system, that means that no programs can be installed on it and you will not be able to store any files there.

Acronis Disk Director Lite will detect that the disk is unusable by the system and needs to be initialized. The **Disk management** view will show the newly detected hardware as a gray block with a grayed icon, thus indicating that the disk is unusable by the system.

If you need to initialize a disk:

1. Select a disk to initialize.
2. Right-click on the selected volume, and then click **Initialize** in the context menu. You will be forwarded to the **Disk Initialization** window, that will provide the basic hardware details such as the disk's number, capacity and state to aid you in the choice of your possible action.
3. In the window, you will be able to set the disk partitioning scheme (MBR or GPT) and the disk type (basic or dynamic). The new disk state will be graphically represented in the **Disk Management** view of the console immediately.
4. By clicking **OK**, you'll add a pending operation of the disk initialization.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

After the initialization, all the disk space remains unallocated and so still impossible to be used for program installation or file storage. To be able to use it, proceed normally to the **Create volume** operation.

If you decide to change the disk settings it can be done later using the standard Acronis Disk Director Lite disk tools.

6.11.5.2 Basic disk cloning

Sometimes it is necessary to transfer all the disk data onto a new disk. It can be a case of expanding the system volume, starting a new system layout or disk evacuation due to a hardware fault. In any case, the reason for the **Clone basic disk** operation can be summed up as the necessity to transfer all the source disk data to a target disk exactly as it is.

Acronis Disk Director Lite allows the operation to be carried out to basic MBR disks only.

To plan the **Clone basic disk** operation:

1. Select a disk you want to clone.
2. Select a disk as target for the cloning operation.
3. Select a cloning method and specify advanced options.

The new volume structure will be graphically represented in the **Disk management** view immediately.

*It is advisable that you deactivate Acronis Startup Recovery Manager (p. 228) (ASRM), if it is active, before cloning a system disk. Otherwise the cloned operating system might not boot. You can activate the ASRM again after the cloning is completed. If deactivation is not possible, choose the **As is** method to clone the disk.*

Selecting source and target disks

The program displays a list of partitioned disks and asks the user to select the source disk, from which data will be transferred to another disk.

The next step is selection of a disk as target for the cloning operation. The program enables the user to select a disk if its size will be sufficient to hold all the data from the source disk without any loss.

If there is some data on the disk that was chosen as the target, the user will receive a warning: “**The selected target disk is not empty. The data on its volumes will be overwritten.**”, meaning that all the data currently located on the chosen target disk will be lost irrevocably.

Cloning method and advanced options

The **Clone basic disk** operation usually means that the information from the source disk is transferred to the target “**As is**”. So, if the destination disk is the same size and even if it is larger, it is possible to transfer all the information there exactly as it is stored at the source.

But with the wide range of available hardware it is normal that the target disk would differ in size from the source. If the destination is larger, then it would be advisable to resize the source disk volumes to avoid leaving unallocated space on the target disk by selecting the **Proportionally resize volumes** option. The option to **Clone basic disk** “as is” remains, but the default method of cloning will be carried out with proportional enlargement of all the **source** disk volumes so that no unallocated space remains on the **target** disk .

If the destination is smaller, then the **As is** option of cloning will be unavailable and proportional resizing of the **source** disk volumes will be mandatory. The program analyzes the **target** disk to establish whether its size will be sufficient to hold all the data from the **source** disk without any loss. If such transfer with proportional resizing of the **source** disk volumes is possible, but without any data loss , then the user will be allowed to proceed. If due to the size limitations safe transfer of all the **source** disk data to the **target** disk is impossible even with the proportional resizing of the volumes, then the **Clone basic disk** operation will be impossible and the user will not be able to continue.

If you are about to clone a disk comprising of a **system volume**, pay attention to the **Advanced options**.

By clicking **Finish**, you'll add the pending operation of the disk cloning.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

Using advanced options

When cloning a disk comprising of a **system volume**, you need to retain an operating system bootability on the target disk volume. It means that the operating system must have the system volume information (e.g. volume letter) matched with the disk NT signature, which is kept in the MBR disk record. But two disks with the same NT signature cannot work properly under one operating system.

If there are two disks having the same NT signature and comprising of a system volume on a machine, at the startup the operating system runs from the first disk, discovers the same signature on the second one, automatically generates a new unique NT signature and assigns it to the second disk. As a result, all the volumes on the second disk will lose their letters, all paths will be invalid on the disk, and programs won't find their files. The operating system on that disk will be unbootable.

You have the following two alternatives to retain system bootability on the target disk volume:

1. Copy NT signature – to provide the target disk with the source disk NT signature matched with the Registry keys also copied on the target disk.
2. Leave NT signature – to keep the old target disk signature and update the operating system according to the signature.

If you need to copy the NT signature:

1. Select the **Copy NT signature** check box. You receive the warning: “If there is an operating system on the hard disk, uninstall either the source or the target hard disk drive from your machine prior to starting the machine again. Otherwise, the OS will start from the first of the two, and the OS on the second disk will become unbootable.” The **Turn off the machine after the cloning operation** check box is selected and disabled automatically.
2. Click **Finish** to add the pending operation.
3. Click **Commit** on the toolbar and then click **Proceed** in the **Pending Operations** window.
4. Wait until the operation is finished.
5. Wait until the machine is turned off.
6. Disconnect either the source or the target hard disk drive from the machine.
7. Start up the machine.

If you need to leave an NT signature:

1. Click to clear the **Copy NT signature** check box, if necessary.
2. Click to clear the **Turn off the machine after the cloning operation** check box, if necessary.
3. Click **Finish** to add the pending operation.
4. Click **Commit** on the toolbar and then click **Proceed** in the **Pending Operations** window.
5. Wait until the operation is finished.

6.11.5.3 Disk conversion: MBR to GPT

You would want to convert an MBR basic disk to a GPT basic disk in the following cases:

- If you need more than 4 primary volumes on one disk.
- If you need additional disk reliability against any possible data damage.

If you need to convert a basic MBR disk to basic GPT:

1. Select a basic MBR disk to convert to GPT.
2. Right-click on the selected volume, and then click **Convert to GPT** in the context menu.

You will receive a warning window, stating that you are about to convert MBR into GPT.

3. By clicking **OK**, you'll add a pending operation of MBR to GPT disk conversion.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

Please note: A GPT-partitioned disk reserves the space in the end of the partitioned area necessary for the backup area, which stores copies of the GPT header and the partition table. If the disk is full and the volume size cannot be automatically decreased, the conversion operation of the MBR disk to GPT will fail.

The operation is irreversible. If you have a primary volume, belonging to an MBR disk, and convert the disk first to GPT and then back to MBR, the volume will be logical and will not be able to be used as a system volume.

If you plan to install an OS that does not support GPT disks, the reverse conversion of the disk to MBR is also possible through the same menu items the name of the operation will be listed as **Convert to MBR**.

Dynamic disk conversion: MBR to GPT

Acronis Disk Director Lite does not support direct MBR to GPT conversion for dynamic disks. However you can perform the following conversions to reach the goal using the program:

1. MBR disk conversion: dynamic to basic (p. 200) using the **Convert to basic** operation.
2. Basic disk conversion: MBR to GPT using the **Convert to GPT** operation.
3. GPT disk conversion: basic to dynamic (p. 199) using the **Convert to dynamic** operation.

6.11.5.4 Disk conversion: GPT to MBR

If you plan to install an OS that does not support GPT disks, conversion of the GPT disk to MBR is possible the name of the operation will be listed as **Convert to MBR**.

If you need to convert a GPT disk to MBR:

1. Select a GPT disk to convert to MBR.
2. Right-click on the selected volume, and then click **Convert to MBR** in the context menu.

You will receive a warning window, stating that you are about to convert GPT into MBR.

You will be explained the changes that will happen to the system after the chosen disk is converted from GPT to MBR. E.g. if such conversion will stop a disk from being accessed by the system, the operating system will stop loading after such conversion or some volumes on the selected GPT disk will not be accessible with MBR (e.g. volumes located more than 2 TB from the beginning of the disk) you will be warned here about such damage.

Please note, a volume, belonging to a GPT disk to convert, will be a logical one after the operation and is irreversible.

3. By clicking **OK**, you'll add a pending operation of GPT to MBR disk conversion.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

6.11.5.5 Disk conversion: basic to dynamic

You would want to convert a basic disk to dynamic in the following cases:

- If you plan to use the disk as part of a dynamic disk group.
- If you want to achieve additional disk reliability for data storage.

If you need to convert a basic disk to dynamic:

1. Select the basic disk to convert to dynamic.
2. Right-click on the selected volume, and then click **Convert to dynamic** in the context menu. You will receive a final warning about the basic disk being converted to dynamic.
3. If you click **OK** in this warning window, the conversion will be performed immediately and if necessary, your machine will be restarted.

Please note: A dynamic disk occupies the last megabyte of the physical disk to store the database, including the four-level description (Volume-Component-Partition-Disk) for each dynamic volume. If during the conversion to dynamic it turns out that the basic disk is full and the size of its volumes cannot be decreased automatically, the basic disk to dynamic conversion operation will fail.

Should you decide to revert your dynamic disks back to basic ones, e.g. if you want to start using an OS on your machine that does not support dynamic disks, you can convert your disks using the same menu items, though the operation now will be named **Convert to basic**.

System disk conversion

Acronis Disk Director Lite does not require an operating system reboot after basic to dynamic conversion of the disk, if:

1. There is a single Windows 2008/Vista operating system installed on the disk.
2. The machine runs this operating system.

Basic to dynamic conversion of the disk, comprising of system volumes, takes a certain amount of time, and any power loss, unintentional turning off of the machine or accidental pressing of the Reset button during the procedure could result in bootability loss.

In contrast to Windows Disk Manager the program ensures bootability of an **offline operating system** on the disk after the operation.

6.11.5.6 Disk conversion: dynamic to basic

You would want to convert dynamic disks back to basic ones, e.g. if you want to start using an OS on your machine that does not support dynamic disks.

If you need to convert a dynamic disk to basic:

1. Select the dynamic disk to convert to basic.
2. Right-click on the selected volume, and then click **Convert to basic** in the context menu. You will receive a final warning about the dynamic disk being converted to basic.

You will be advised about the changes that will happen to the system if the chosen disk is converted from dynamic into basic. E.g. if such a conversion will stop the disk from being accessed by the system, the operating system will stop loading after such conversion, or if the disk you want to convert to basic contains any volumes of the types that are only supported by dynamic disks (all volume types except Simple volumes), then you will be warned here about the possible damage to the data involved in the conversion.

Please note, the operation is unavailable for a dynamic disk containing Spanned, Striped, or RAID-5 volumes.

3. If you click **OK** in this warning window, the conversion will be performed immediately.

After the conversion the last 8Mb of disk space is reserved for the future conversion of the disk from basic to dynamic.

In some cases the possible unallocated space and the proposed maximum volume size might differ (e.g. when the size of one mirror establishes the size of the other mirror, or the last 8Mb of disk space are reserved for the future conversion of the disk from basic to dynamic).

System disk conversion

Acronis Disk Director Lite does not require an operating system reboot after dynamic to basic conversion of the disk, if:

1. There is a single Windows 2008/Vista operating system installed on the disk.
2. The machine runs this operating system.

Dynamic to basic conversion of the disk, comprising of system volumes, takes a certain amount of time, and any power loss, unintentional turning off of the machine or accidental pressing of the Reset button during the procedure could result in bootability loss.

In contrast to Windows Disk Manager the program ensures:

- safe conversion of a dynamic disk to basic when it contains volumes **with data** for simple and mirrored volumes
- in multiboot systems, bootability of a system that was **offline** during the operation

6.11.5.7 Changing disk status

Changing disk status is effective for Windows Vista SP1, Windows Server 2008, Windows 7 operating systems and applies to the current disk layout (p. 194).

One of the following disk statuses always appears in the graphical view of the disk next to the disk's name:

- **Online**

The online status means that a disk is accessible in the read-write mode. This is the normal disk status. If you need a disk to be accessible in the read-only mode, select the disk and then change its status to offline by selecting **Change disk status to offline** from the **Operations** menu.

- **Offline**

The offline status means that a disk is accessible in the read-only mode. To bring the selected offline disk back to online, select **Change disk status to online** from the **Operations** menu. If the disk has the offline status and the disk's name is **Missing**, this means that the disk cannot be located or identified by the operating system. It may be corrupted, disconnected, or powered off. For information on how to bring a disk that is offline and missing back online, please refer to the following Microsoft knowledge base article: <http://technet.microsoft.com/en-us/library/cc732026.aspx>.

6.11.5.8 Importing foreign disks

On a machine with two or more operating systems, representation of disks and volumes depends on which operating system is currently running.

Normally, all dynamic disks created within the same machine and operating system are members of the same disk group. When moved to another machine, or added to another operating system on the same machine, a disk group is considered as **foreign**. Foreign group disks cannot be used until they are imported into the existing disk group. A foreign group is imported as is (will have the original name) if no disk group exists on the machine.

To access data on foreign disks, you have to add these disks to your machine's system configuration using the **Import foreign disks** operation.

All dynamic disks of the foreign disk group are imported at the same time, you cannot import just one dynamic disk.

To import foreign disks

1. Right-click one of the foreign disks, and then click **Import foreign disks**.

The appearing window lists all foreign dynamic disks that were added to the machine, and displays information about volumes that will be imported. Volume statuses lets you detect whether you are importing all the required disks of the disk group. When importing all the required disks, all their volumes have the **Healthy** status. Statuses other than **Healthy** indicate that not all of the disks were imported.

For more information on volume statuses please refer to the following Microsoft article:

<http://technet.microsoft.com/en-us/library/cc771775.aspx>

2. Click **OK** to add the pending foreign disks importing operation.

The results of the pending operation are immediately displayed as if the operation had been performed.

To perform the pending operation you will have to commit it. Exiting the program without committing the pending operations will effectively cancel them.

6.11.6 Volume operations

Acronis Disk Director Lite includes the following operations that can be performed on volumes:

- Create Volume (p. 202) - Creates a new volume with the help of the Create Volume Wizard.
- Delete Volume (p. 206) - Deletes the selected volume.
- Set Active (p. 206) - Sets the selected volume Active so that the machine will be able to boot with the OS installed there.
- Change Letter (p. 207) - Changes the selected volume letter
- Change Label (p. 207) - Changes the selected volume label
- Format Volume (p. 208) - Formats a volume giving it the necessary file system

The full version of Acronis Disk Director will provide more tools and utilities for working with volumes.

Acronis Disk Director Lite must obtain exclusive access to the target volume. This means no other disk management utilities (like Windows Disk Management utility) can access it at that time. If you receive a message stating that the volume cannot be blocked, close the disk management applications that use this volume and start again. If you can not determine which applications use the volume, close them all.

6.11.6.1 Creating a volume

You might need a new volume to:

- Recover a previously saved backup copy in the “exactly as was” configuration;
- Store collections of similar files separately — for example, an MP3 collection or video files on a separate volume;
- Store backups (images) of other volumes/disks on a special volume;
- Install a new operating system (or swap file) on a new volume;

- Add new hardware to a machine.

In Acronis Disk Director Lite the tool for creating volumes is the **Create volume Wizard**.

Types of dynamic volumes

Simple Volume

A volume created from free space on a single physical disk. It can consist of one region on the disk or several regions, virtually united by the Logical Disk Manager (LDM). It provides no additional reliability, no speed improvement, nor extra size.

Spanned Volume

A volume created from free disk space virtually linked together by the LDM from several physical disks. Up to 32 disks can be included into one volume, thus overcoming the hardware size limitations, but if at least one disk fails, all data will be lost, and no part of a spanned volume may be removed without destroying the entire volume. So, a spanned volume provides no additional reliability, nor a better I/O rate.

Striped Volume

A volume, also sometimes called RAID 0, consisting of equal sized stripes of data, written across each disk in the volume; it means that to create a striped volume, a user will need two or more dynamic disks. The disks in a striped volume don't have to be identical, but there must be unused space available on each disk that you want to include in the volume and the size of the volume will depend on the size of the smallest space. Access to the data on a striped volume is usually faster than access to the same data on a single physical disk, because the I/O is spread across more than one disk.

Striped volumes are created for improved performance, not for their better reliability - they do not contain redundant information.

Mirrored Volume

A fault-tolerant volume, also sometimes called RAID 1, whose data is duplicated on two identical physical disks. All of the data on one disk is copied to another disk to provide data redundancy. Almost any volume can be mirrored, including the system and boot volumes, and if one of the disks fails, the data can still be accessed from the remaining disks. Unfortunately, the hardware limitations on size and performance are even more severe with the use of mirrored volumes.

Mirrored-Striped Volume

A fault-tolerant volume, also sometimes called RAID 1+0, combining the advantage of the high I/O speed of the striped layout and redundancy of the mirror type. The evident disadvantage remains inherent with the mirror architecture - a low disk-to-volume size ratio.

RAID-5

A fault-tolerant volume whose data is striped across an array of three or more disks. The disks do not need to be identical, but there must be equally sized blocks of unallocated space available on each disk in the volume. Parity (a calculated value that can be used to reconstruct data in case of failure) is also striped across the disk array. And it is always stored on a different disk than the data itself. If a physical disk fails, the portion of the RAID-5 volume that was on that failed disk can be re-created from the remaining data and the parity. A RAID-5 volume provides reliability and is able to overcome the physical disk size limitations with a higher than mirrored disk-to-volume size ratio.

Create volume wizard

The **Create volume** wizard lets you create any type of volume (including system and active), select a file system, label, assign a letter, and also provides other disk management functions.

Its pages will enable you to enter operation parameters, proceeding step-by-step further on and return to any previous step if necessary to change any previously selected options. To help you with your choices, each parameter is supplemented with detailed instructions.

If you want to create a volume:

Run the **Create volume** wizard by selecting **Create volume** on the **Wizards** bar, or right-click any unallocated space and select **Create volume** in the appearing context menu.

Select the type of volume being created

At the first step you have to specify the type of volume you want to create. The following types of volume are available:

- Basic
- Simple/Spanned
- Striped
- Mirrored
- RAID-5

You will obtain a brief description of every type of volume for better understanding of the advantages and limitations of each possible volume architecture.

*If the current operating system, installed on this machine, does not support the selected type of volume, you will receive the appropriate warning. In this case the **Next** button will be disabled and you will have to select another type of volume to proceed with the new volume creation.*

After you click the **Next** button, you will proceed forward to the next wizard page: Select destination disks (p. 204).

Select destination disks

The next wizard page will prompt you to choose the disks, whose space will be used for the volume creation.

To create a basic volume:

- Select a destination disk and specify the unallocated space to create the basic volume on.

To create a Simple/Spanned volume:

- Select one or more destination disks to create the volume on.

To create a Mirrored volume:

- Select two destination disks to create the volume on.

To create a Striped volume:

- Select two or more destination disks to create the volume on.

To create a RAID-5 volume:

- Select three destination disks to create the volume on.

After you choose the disks, the wizard will calculate the maximum size of the resulting volume, depending on the size of the unallocated space on the disks you chose and the requirements of the volume type you have previously decided upon.

If you are creating a **dynamic** volume and select one or several **basic** disks, as its destination, you will receive a warning that the selected disk will be converted to dynamic automatically.

If need be, you will be prompted to add the necessary number of disks to your selection, according to the chosen type of the future volume.

If you click the **Back** button, you will be returned to the previous page: Select the type of volume being created (p. 204).

If you click the **Next** button, you will proceed to the next page: Set the volume size (p. 205).

Set the volume size

On the third wizard page, you will be able to define the size of the future volume, according to the previously made selections. In order to choose the necessary size between the minimum and the maximum values, use the slider or enter the necessary values into the special windows between the minimum and the maximum values or click on the special handle, and hold and drag the borders of the disk's picture with the cursor.

The maximum value normally includes the most possible unallocated space. But in some cases the possible unallocated space and the proposed maximum volume size might differ (e.g. when the size of one mirror establishes the size of the other mirror, or the last 8Mb of the disk space is reserved for the future conversion of the disk from basic to dynamic).

For basic volumes if some unallocated space is left on the disk, you also will be able to choose the position of the new volume on the disk.

If you click the **Back** button, you will be returned to the previous page: Select destination disks (p. 204).

If you click the **Next** button, you will proceed to the next page: Set the volume options (p. 205).

Set the volume options

On the next wizard page you can assign the volume **Letter** (by default - the first free letter of the alphabet) and, optionally, a **Label** (by default – none). Here you will also specify the **File system** and the **Cluster size**.

The wizard will prompt you to choose one of the Windows file systems: FAT16 (disabled, if the volume size has been set at more than 2 GB), FAT32 (disabled, if the volume size has been set at more than 2 TB), NTFS or to leave the volume **Unformatted**.

In setting the cluster size you can choose between any number in the preset amount for each file system. Note, the program suggests the cluster size best suited to the volume with the chosen file system.

If you are creating a basic volume, which can be made into a system volume, this page will be different, giving you the opportunity to select the volume **Type — Primary (Active Primary)** or **Logical**.

Typically **Primary** is selected to install an operating system to a volume. Select the **Active** (default) value if you want to install an operating system on this volume to boot at machine startup. If the

Primary button is not selected, the **Active** option will be inactive. If the volume is intended for data storage, select **Logical**.

*A Basic disk can contain up to four primary volumes. If they already exist, the disk will have to be converted into dynamic, otherwise **Active** and **Primary** options will be disabled and you will only be able to select the **Logical** volume type. The warning message will advise you that an OS installed on this volume will not be bootable.*

*If you use characters when setting a new volume label that are unsupported by the currently installed operation system, you will get the appropriate warning and the **Next** button will be disabled. You will have to change the label to proceed with the creation of the new volume.*

If you click the **Back** button, you will be returned to the previous page: Set the volume size (p. 205).

If you click the **Finish** button, you will complete the operation planning.

To perform the planned operation click **Commit** in the toolbar, and then click **Proceed** in the **Pending Operations** window.

If you set a 64K cluster size for FAT16/FAT32 or on 8KB-64KB cluster size for NTFS, Windows can mount the volume, but some programs (e.g. Setup programs) might calculate its disk space incorrectly.

6.11.6.2 Delete volume

This version of Acronis Disk Director Lite has reduced functionality because it is mainly a tool for preparing bare-metal systems for recovering previously saved volume images. The features of resizing the existing volumes and creating the new volumes, using free space from the existing ones, exist on the full version of the software, so with this version deleting an existing volume sometimes might be the only way to free the necessary disk space without changing the existing disk configuration.

After a volume is deleted, its space is added to unallocated disk space. It can be used for creation of a new volume or to change another volume's type.

If you need to delete a volume:

1. Select a hard disk and a volume to be deleted.
2. Select **Delete volume** or a similar item in the **Operations** sidebar list, or click the **Delete the selected volume** icon on the toolbar.

If the volume contains any data, you will receive the warning, that all the information on this volume will be lost irrevocably.

3. By clicking **OK** in the **Delete volume** window, you'll add the pending operation of volume deletion.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

6.11.6.3 Set active volume

If you have several primary volumes, you must specify one to be the boot volume. For this, you can set a volume to become active. A disk can have only one active volume, so if you set a volume as active, the volume, which was active before, will be automatically unset.

If you need to set a volume active:

1. Select a primary volume on a basic MBR disk to set as active.
2. Right-click on the selected volume, and then click **Mark as active** in the context menu.

If there is no other active volume in the system, the pending operation of setting active volume will be added.

Please note, that due to setting the new active volume, the former active volume letter might be changed and some of the installed programs might stop running.

3. If another active volume is present in the system, you will receive the warning that the previous active volume will have to be set passive first. By clicking **OK** in the **Warning** window, you'll add the pending operation of setting active volume.

Please note: even if you have the Operating System on the new active volume, in some cases the machine will not be able to boot from it. You will have to confirm your decision to set the new volume as active.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the **Disk management** view immediately.

6.11.6.4 Change volume letter

Windows operating systems assign letters (C:, D:, etc) to hard disk volumes at startup. These letters are used by applications and operating systems to locate files and folders in the volumes.

Connecting an additional disk, as well as creating or deleting a volume on existing disks, might change your system configuration. As a result, some applications might stop working normally or user files might not be automatically found and opened. To prevent this, you can manually change the letters that are automatically assigned to the volumes by the operating system.

If you need to change a letter assigned to a volume by the operating system:

1. Select a volume to change a letter.
2. Right-click on the selected volume, and then click **Change letter** in the context menu.
3. Select a new letter in the **Change Letter** window.
4. By clicking **OK** in the **Change Letter** window, you'll add a pending operation to volume letter assignment.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the **Disk management** view immediately.

6.11.6.5 Change volume label

The volume label is an optional attribute. It is a name assigned to a volume for easier recognition. For example, one volume could be called SYSTEM — a volume with an operating system, or PROGRAM — an application volume, DATA — a data volume, etc., but it does not imply that only the type of data stated with the label could be stored on such a volume.

In Windows, volume labels are shown in the Explorer disk and folder tree: LABEL1(C:), LABEL2(D:), LABEL3(E:), etc. LABEL1, LABEL2 and LABEL3 are volume labels. A volume label is shown in all application dialog boxes for opening and saving files.

If you need to change a volume label:

1. Right-click on the selected volume, and then click **Change label**.

2. Enter a new label in the **Change label** window text field.
3. By clicking **OK** in the **Change label** window, you'll add the pending operation of changing the volume label .

*If when setting a new volume label you use characters that are unsupported by the currently installed operating system, you will get the appropriate warning and the **OK** button will be disabled. You will have to use only supported characters to proceed with changing the volume label.*

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new label will be graphically represented in the **Disk Management** view of the console immediately.

6.11.6.6 Format volume

You might want to format a volume if you want to change its file system:

- to save additional space which is being lost due to the cluster size on the FAT16 or FAT32 file systems
- as a quick and more or less reliable way of destroying data, residing in this volume

If you want to format a volume:

1. Select a volume to format.
2. Right-click on the selected volume, and then click **Format** in the context menu.

You will be forwarded to the **Format Volume** window, where you will be able to set the new file system options. You can choose one of the Windows file systems: FAT16 (disabled, if the Volume Size is more than 2 GB), FAT32 (disabled, if the Volume Size is more than 2 TB) or NTFS.

In the text window you will be able to enter the volume label, if necessary: by default this window is empty.

In setting the cluster size you can choose between any number in the preset amount for each file system. Note, the program suggests the cluster size best suited to the volume with the chosen file system.

3. If you click **OK** to proceed with the **Format Volume** operation, you'll add a pending operation of formatting a volume.

(To finish the added operation you will have to commit (p. 208) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the **Disk management** view.

If you set a 64K cluster size for FAT16/FAT32 or an 8KB-64KB cluster size for NTFS, Windows can mount the volume, but some programs (e.g. Setup programs) might calculate its disk space incorrectly.

6.11.7 Pending operations

All operations, which were prepared by the user in manual mode or with the aid of a wizard, are considered pending until the user issues the specific command for the changes to be made permanent. Until then, Acronis Disk Director Lite will only demonstrate the new volume structure that will result from the operations that have been planned to be performed on disks and volumes. This approach enables you to control all planned operations, double-check the intended changes, and, if necessary, cancel operations before they are executed.

To prevent you from performing any unintentional change on your disk, the program will first display the list of all pending operations.

The **Disk management** view contains the toolbar with icons to launch the **Undo**, **Redo** and **Commit** actions intended for pending operations. These actions might also be launched from the **Disk management** menu of the console.

All planned operations are added to the pending operation list.

The **Undo** action lets you undo the latest operation in the list. While the list is not empty, this action is available.

The **Redo** action lets you reinstate the last pending operation that was undone.

The **Commit** action forwards you to the **Pending Operations** window, where you will be able to view the pending operation list. Clicking **Proceed** will launch their execution. You will not be able to undo any actions or operations after you choose the **Proceed** operation. You can also cancel the commitment by clicking **Cancel**. Then no changes will be done to the pending operation list.

Quitting Acronis Disk Director Lite without committing the pending operations effectively cancels them, so if you try to exit **Disk management** without committing the pending operations, you will receive the appropriate warning.

6.12 Collecting system information

The system information collection tool gathers information about the machine to which the management console is connected, and saves it to a file. You may want to provide this file when contacting Acronis technical support.

This option is available under bootable media and for machines where Agent for Windows, Agent for Linux or Acronis Backup & Recovery 10 Management Server is installed.

To collect system information

1. In the management console, select from the top menu **Help > Collect system information from 'machine name'**.
2. Specify where to save the file with system information.

7 Online backup

This section provides details about using the Acronis Backup & Recovery 10 Online service. This service enables you to do online backups to Acronis Online Backup Storage.

Acronis Backup & Recovery 10 Online might be unavailable in your region. To find more information, click here: <http://www.acronis.eu/my/backup-recovery-online/>

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

Creating a backup plan (p. 116)

Creating a backup policy

Recovering data

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

Host-based backups of virtual machines are possible with Acronis Backup & Recovery 10 Virtual Edition. You can back up all virtual machines managed by Agent for ESX/ESXi or Agent for Hyper-V with a single subscription for virtual machines.

7.1 Introduction to Acronis Backup & Recovery 10 Online

This section contains a brief overview of Acronis Backup & Recovery 10 Online and answers questions that may arise during evaluation and usage of this product.

7.1.1 What is Acronis Backup & Recovery 10 Online?

Acronis Backup & Recovery 10 Online is a service that enables you to back up data to Acronis Online Backup Storage. To use this service, you need to buy a subscription that determines the amount of storage space reserved for your backups (storage quota) and how long the online service will be available to you.

A separate subscription is required for each physical machine you are going to back up. The server or workstation subscription that you buy will depend on the Windows operating system that the machine is running. Other operating systems are not supported by Acronis Backup & Recovery 10 Online.

Example: A 250 GB/ 1 year workstation subscription means that you can back up data from a machine running a non-server Windows operating system for a period of 1 year. The backups can occupy no more than 250 GB.

7.1.2 What data can I back up and recover?

You can back up any files, volumes, or the entire physical machine as often as you wish. Unlike most online backup solutions, Acronis Backup & Recovery 10 Online enables bare metal recovery directly from the online storage. Files can be recovered from disk-level backups as well as from file-level backups.

7.1.3 How long will my backups be kept in the online storage?

Your backups remain in the online storage until you delete them or until the subscription expires. Recovering data from the online storage is possible for 30 days following the subscription expiration date.

For effective use of the storage space, you have the option to set up the "**Delete backups older than**" retention rule.

Example

You might want to use the following backup strategy for a file server.

Back up the critical files twice a day on a schedule. Set the retention rule "**Delete backups older than**" 7 days. This means that after every backup the software will check for backups older than 7 days and delete them automatically.

Run backup of the server's system volume manually as required. For example, after the operating system updates. Manually delete the backups that you do not need.

7.1.4 How to secure my data?

Backups can be encrypted using the Advanced Encryption Standard (AES) cryptographic algorithm and the password you set. This guarantees that your data is not accessed by anyone else.

7.1.5 Backup and recovery FAQ

This section answers questions related to backup and recovery processes.

7.1.5.1 What backup methods are available?

Full and incremental backup methods are available through the following backup schemes.

Back up now (immediate start) or **Back up later** (postponed start). You can select either the full or incremental backup method. On the first task run a full backup is created. If you run the backup task again, the selected backup method will be applied. This way, you will obtain either a new full backup, or an incremental one. Since full backups take more storage space, Acronis recommends that you choose incremental backup if you want to run backups manually.

Simple (start on schedule). The first backup is full, the later backups are incremental. With this backup scheme, you can set up a retention rule to automatically delete old backups.

An additional backup scheme that is available only for online storage is **Initial seeding**. This is the **Back up now** scheme using local destination and full backup method. To use this scheme, you need a license for the Initial Seeding (p. 213) service.

7.1.5.2 Is the online storage available under Acronis bootable media?

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

7.1.5.3 Can I use Acronis Universal Restore when recovering a system from the online storage?

Yes. Acronis Universal Restore is always available when recovering a system from the online storage. Using Acronis Universal Restore when recovering from other types of storage will require a separate license.

7.1.5.4 What if a network connection is lost during online backup or recovery?

The software will try to reach the online storage every 30 seconds. After five unsuccessful attempts the backup or recovery task will fail.

You can change the number of attempts and the interval between the attempts in the **Error handling > Re-attempt, if an error occurs** option. Every backup plan or recovery task includes this option.

7.1.5.5 What happens if I run out of space?

When a machine's backups are about to exceed the storage space allowed by its subscription, you receive an e-mail notification with an alert. In addition, you can see this alert on the account management web page near the machine. This means you have to free some space for future backups. You may also want to set or edit the retention rule (p. 211) so that an overflow does not occur. Once the occupied space reaches the limit, the backups will cease to run.

7.1.5.6 What is the cleanup task for?

Any backup plan where the retention rule is set contains a cleanup task in addition to a backup task. The cleanup task checks the archive created by the backup plan for backups that have outlived their lifetime. If such backups are found, the task makes the online storage delete them. Since the deletion is performed on the online storage side, it does not take your machine's CPU resource.

The cleanup task runs after every online backup, even if the backup has failed. The last successful backup is always kept though. For more information about the retention rule please refer to "How long will my backups be kept in the online storage? (p. 211)"

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.

7.1.5.7 How to make a recovered machine recognize its subscription?

When you recover a physical machine from a backup, a new machine identifier is created. Therefore, the machine is not able to back up to the subscription it used before recovery.

To continue backing up the machine to the same subscription, reassign (p. 223) 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.

7.1.6 Initial Seeding FAQ

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

7.1.6.1 What is Initial Seeding?

Initial Seeding is an extra service that lets you save an initial full backup locally and then send it to Acronis on a hard disk drive.

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

The hard disk drive is sent back to you but it is not possible to recover from it. However, recovery from a locally attached device is possible with the Large scale recovery (p. 217) option.

7.1.6.2 Why would I want to use Initial Seeding?

This service helps you save time and network traffic during the initial full backup. It is useful when backing up very large volumes of data or entire machines to the online storage.

7.1.6.3 Is Initial Seeding a paid service?

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

7.1.6.4 What types of hard drive can I use for Initial Seeding?

Acronis accepts hard disk drives of the following interface types: IDE, ATA, SATA, USB connected drives. SCSI drives are not accepted.

You can back up directly to the device or back up to a local or network folder and then copy the backup to the device. Make sure that the device has only one volume and that the file system on that volume is NTFS or FAT32.

7.1.6.5 Can I send more than one backup under a single Initial Seeding license?

Yes, if the backups are taken on the same machine and sent on the same disk.

For example, you may want to upload a volume backup and a number of file backups. Do as many "initial seeding" backups as you wish and then send them to Acronis on the same disk. The Initial Seeding license gets used when uploading to the online storage starts. You will receive an e-mail notification informing you about this. Afterwards, creating a new "initial seeding" backup on the same machine will require a new Initial Seeding license.

7.1.6.6 Can I send backups taken from a number of machines on a single hard drive?

Yes. However, the number of required licenses is still one per machine.

7.1.6.7 How to buy an Initial Seeding license?

You can buy an Initial Seeding license from an Acronis partner or in Acronis online store. Follow the link <http://www.acronis.eu/my/backup-recovery-online/#buy> to locate a partner or to buy online.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Click **Enter new registration code** on the same web page and register the license. The license becomes available on the **Initial Seeding / Recovery** tab.

A license purchased in Acronis online store becomes available immediately after the payment is processed.

7.1.6.8 How to perform initial seeding?

1. Decide on the media (p. 213) that you will send.
2. 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.
3. Start Acronis Backup & Recovery 10, click **Back up** and create a backup plan on this machine:
 - In **Where to back up**, specify **Online Backup Storage**.
 - In **What to back up**, select disks, volumes or files/folders you want to back up.
 - 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**.
 - [Optional, but strongly recommended] Enable backup validation in **Archive validation**.The backup starts immediately after you click the final **OK**.
4. [Optional] If you want to add more backups to the media, repeat step 3 selecting the different data in **What to back up**. (Do not edit the backup plan, create a new one!)
5. [Optional] If you want to add backups from another machine, attach the media to that machine and perform the same steps.
6. Package (p. 214) the media along with a prepaid return shipping label and send it to Acronis by physical mail. The address is available on your account management Web page > **Initial Seeding / Recovery** tab > **Orders in process > Initial Seeding orders > Datacenter address**.
7. On the same web page, mark the order as "shipped" and track (p. 216) the order status.
8. Once you observe that the backup has been uploaded on the online storage, you can edit the backup plan to do incremental backups:
 - In **Backup scheme**, select **Back up now**, **Back up later** for manual backups or **Simple** for scheduled backups.
 - For **Simple**, specify the schedule and (optionally) the retention rule.
 - Click **Save**.

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

7.1.6.9 How to package a hard drive for shipment?

It is very important that your hard drive be packaged carefully. Careful packaging will protect your drive from any damage during shipment.

Hard drive types

Acronis accepts hard disk drives of the following interface types: IDE, ATA, SATA, USB connected drives.

SCSI drives are not accepted.

Packaging

If possible, use the original packaging. Otherwise, packaging materials can be obtained at any shipping outlet or stationary store.

The following are instructions about how to package your hard disk drive.

Step 1

Delicately remove your hard disk drive from the machine.



Step 2

Place the hard drive into an anti-static bag to protect the drive from electrostatic discharge. If you do not have an anti-static bag, simply wrap the hard drive into aluminum foil.




Step 3

Use a sturdy box that is at least twice the size of the drive. Pack the drive with a bubble wrap around all 6 sides so it can fit tight into the box and cannot be moved within.

DO NOT use Styrofoam *peanuts* for packing as they do not provide enough protection.

DO NOT send your media in *jiffy* bags



Step 4

Choose the transport company that you will use for shipping. On this company's web site, prepare and print two prepaid shipping labels:

1. **Shipping label for sending** your hard drive. This label is placed on the top of the box. You should send your package to one of the Acronis data centers. The data center address can be obtained on the **Initial seeding/Recovery** tab of your account management page by clicking **Show data center address**.

We recommend that you use the overnight shipping, if you want to start doing incremental backups as soon as possible. The data is generally available on 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.



7.1.6.10 How do I track an Initial Seeding order status?

On the Acronis Web site, the **Initial Seeding / Recovery** tab shows you the status of all your orders. In addition, you will receive e-mail notifications about the most important events.

- **Available** – The license is available for using on any machine.

- **An order was created** – The first backup is about to start and 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.
- **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. 214) 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 on the same machine will require a new Initial Seeding license.
- **The data upload has started** – The process of uploading data to Acronis Online Backup Storage has started.
- **The data upload has been completed** – The initial full backup has been successfully uploaded to the online storage. You can do incremental online backups now.
- **The order has been completed. The media has been returned (or: Returning the media was not requested)** – Your media has been shipped back (the carrier and the tracking number are specified). If a prepaid shipping label was not provided with the media, the media will be discarded.
- [Occasional] **The order is on hold** – Your order was placed on hold due to technical difficulties processing the order. Acronis is working on resolving these issues.
- [Occasional] **The order has been cancelled** – The order had been cancelled before the media was shipped, so returning the media is not required.
- [Occasional] **The order has been cancelled. The media has been returned (or: Returning the media was not requested)** – The order was cancelled while the media was in the datacenter. The media has been shipped back (the carrier and the tracking number are specified). If a prepaid shipping label was not provided with the media, the media will be discarded.

7.1.7 Large Scale Recovery FAQ

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

7.1.7.1 What is Large Scale Recovery?

Large Scale Recovery is an extra service that enables you to obtain a copy of the backups you have in the online storage. You can then recover data from this copy.

Once you order Large Scale Recovery for a particular machine, Acronis sends you a USB hard disk drive with all of the backups made from this machine. You can recover data directly from the disk or copy the backups to a local or network folder.

7.1.7.2 Why would I use Large Scale Recovery?

In the event of a disaster or the need to recover large volumes of data or the entire machines quickly, this service helps you save time and network traffic. Recovering hundreds of gigabytes over the Internet may take days. This process will deliver a faster recovery.

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

No, these services are independent.

7.1.7.4 Is Large Scale Recovery a paid service?

Yes, you need to buy one Large Scale Recovery license per machine. The license enables you to get a disk with all of the currently available backups of this machine. To obtain backups that will be created in the future, you will need a new Large Scale Recovery license.

7.1.7.5 Can I perform large scale recovery on a different machine?

Yes. You can recover the data an unlimited number of times on any machine you wish. Acronis Universal Restore is included to help you recover an operating system to dissimilar hardware.

7.1.7.6 Can I obtain backups taken from a number of machines on a single hard drive?

No. A separate hard drive is required for each machine.

7.1.7.7 How to buy a Large Scale Recovery license?

You can buy a Large Scale Recovery license from an Acronis partner or in Acronis online store. Follow the link <http://www.acronis.eu/my/backup-recovery-online/#buy> to locate a partner or to buy online.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Click **Enter new registration code** on the same web page and register the license. The license becomes available on the **Initial Seeding / Recovery** tab.

A license purchased in Acronis online store becomes available immediately after the payment is processed.

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

On the Acronis Web site, the **Initial Seeding / Recovery** tab shows you the status of all your orders. In addition, you will receive e-mail notifications about most important events.

- **Available** – The license can be used for any machine.
- **An order was created** – This status is set upon completion of the Large Scale Recovery order form. This means that the license cannot be used for any other machine. From this point on, you

can cancel the order if something goes wrong. This will return the license to the pool of available licenses.

- **The order is being processed** - Order processing in the datacenter started.
- **Writing data** – Your backups are being written onto the media. From this point on, you cannot cancel the order.
- **Writing data has been completed** – Your backups have been successfully written to the media.
- **Ready to ship the media** – Your order has been processed and the media will be shipped shortly.
- **The order has been completed. The media has been shipped** – The media has been shipped to you (the carrier and the tracking number are specified).
- [Occasional] **The order is on hold** – Your order was placed on hold due to technical difficulties processing the order. Acronis is working on resolving these issues.
- [Occasional] **The order has been cancelled** – The order has been cancelled.
- [Occasional] **Address is undeliverable** – Acronis cannot send the disk. On the same Web page, click **Change my delivery address** and specify the correct address for the order.
- [Occasional] **Address has been updated** – This status is set after you have updated the delivery address on Acronis web site.

7.1.7.9 How to perform large scale recovery?

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

7.1.8 Subscription lifecycle FAQ

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

7.1.8.1 How to access my account management Web page?

To access this Web page from the Acronis Web site:

1. Select **User Login**.
2. Log in to your account (create one if you are not registered yet).
3. Navigate to **Online backup > For Business**.

To access this Web page from Acronis Backup & Recovery 10:

1. Click **Activate online backup subscriptions** in the **Actions** menu.
2. Click **Go to account management Web page**.
3. Log in to your account (create one if you are not registered yet).

7.1.8.2 Where do I find the subscriptions that I purchased?

If you purchased your subscriptions from an Acronis partner, you should have received an e-mail confirming the registration codes for each subscription. Create an account on the Acronis web site, if you do not have one already, and log in to it. Navigate to **Online backup > For Business**. This is your *account management Web page*. Click **Enter new registration code** and enter the registration codes. The subscriptions will appear in the list of available subscriptions under the **Manage Subscriptions** tab.

If you purchased your subscriptions online using the Acronis Web site, they are available immediately on your account management Web page. The newly obtained subscriptions are listed at the **Manage Subscriptions** tab.

7.1.8.3 When does my subscription begin?

Your subscription begins when you choose to have it begin, not at the time of purchase.

The time count will begin as soon as the subscription is activated. The first activation occurs when you assign a subscription to a certain machine. To do so, you need to have Acronis software installed.

7.1.8.4 What happens when my subscription expires?

A month before the subscription expiration date you receive an e-mail notification with an alert. In addition, you can see this alert on the account management web page near the machine. This means you need to renew (p. 220) the subscription to continue backing up the machine.

If you do not renew the subscription, you will be able to back up data to the online storage for five days following the expiration date. You will be able to recover data from the online storage for 30 days following the expiration date.

7.1.8.5 How do I renew a subscription?

Buy another subscription and specify it as the next subscription of the same machine. The new subscription will be activated as soon as the current subscription expires.

An expired subscription can be renewed within five days after expiration. In such cases, the new subscription will be activated immediately.

Renewing a single subscription

To renew a subscription

1. Go to the account management Web page.
2. Make sure that you have an available subscription with the *same* storage quota.
3. Select the machine that you want to renew the subscription for, and then click **Renew**.

The subscription appears in the **Next subscription** column for the selected machine.

Renewing a number of activated subscriptions at once

This operation is possible if the appropriate number of new subscriptions are identical to the currently used subscriptions.

Make sure the new subscriptions are available on your account management Web page. Then click **Renew all**. The confirmation window will summarize which subscriptions will be renewed. If identical subscriptions are not found for some of the machines, you have the option to cancel automatic renewal and renew each subscription individually.

What does "Auto-renew" mean?

Auto-renewal means that when the current subscription expires, the next subscription will be automatically selected from the available subscriptions. The next subscription must be identical to the current subscription.

If an identical subscription is not found, auto-renewal will not take place and your backups may fail. No subscriptions will be bought automatically. Only those subscriptions available at the time of auto-renewal can be used. You can select auto-renewal for each individual subscription or set up bulk auto-renewal of all of the activated subscriptions you have.

7.1.8.6 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.

7.1.8.7 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. 223) it to a different machine in Acronis Backup & Recovery 10 GUI.

7.1.8.8 Can I cancel my subscription?

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

7.2 Where do I start?

On the Acronis Web site, log in to your account (create one if you are not registered yet) and navigate to **Online backup > For Business**. This is your *account management Web page*. Here you can get a trial subscription, locate an Acronis partner or buy subscriptions online. The newly obtained subscriptions are listed as available subscriptions in the **Manage Subscriptions** tab.

If you purchased your subscriptions from an Acronis partner, register them manually using the **Enter new registration code** link. The registration codes come with the purchase confirmation e-mail.

Next, install Acronis software (if not yet installed) and assign (p. 222) each subscription to a machine. The subscriptions become activated. After that, you can start backing up to Acronis Online Backup Storage.

7.3 Choosing a subscription

Normally, you choose a subscription based on the operating system your machine is running.

Server operating systems supported by Acronis Backup & Recovery 10 Online:

- Windows 2000 Server/Advanced Server
- Windows Server 2003/2003 R2 - the Standard, Enterprise, Small Business Server editions (x86, x64)
- Windows Server 2008 - the Standard, Enterprise, Small Business Server, Foundation editions (x86, x64)

- Windows Server 2008 R2 - the Standard, Enterprise, Small Business Server, Datacenter, Foundation editions
- Windows MultiPoint Server 2010

Workstation operating systems supported by Acronis Backup & Recovery 10 Online:

- Windows 2000 Professional SP4
- Windows XP Professional SP2+ (x86, x64)
- Windows Vista - all editions except for Vista Home Basic and Vista Home Premium (x86, x64)
- Windows 7 - all editions except for the Starter and Home editions (x86, x64)

Virtualization products supported by Acronis Backup & Recovery 10 Online (host-based backup of virtual machines):

- VMware ESX Infrastructure 3.5 Update 2+
- VMware ESX(i) 4.0 and 4.1

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

- Windows Server 2008/2008 R2 (x64) with Hyper-V
- Microsoft Hyper-V Server 2008/2008 R2

If your backups are likely to exceed the storage quota for this type of subscription, you may want to use a subscription with larger storage quota. For example, you can use a server subscription or a subscription for virtual machines on a workstation. Or you can use a subscription for virtual machines on a server which is not a virtualization server.

The inverse usage is not possible. You cannot back up a server using a workstation subscription. If you try to back up ESX(i) or Hyper-V virtual machines from the host that uses a server subscription, the backup will fail. With a server subscription, you can only back up the Windows host. With a subscription for virtual machines, you can back up both the Windows host and its virtual machines.

Trial subscriptions

You can get one free workstation subscription, server subscription, or subscription for virtual machines per account. The storage quota of the trial subscription is equal to that of the standard subscription. The subscription period is limited to 2 months.

Obtaining a trial subscription is possible until you buy a paid subscription. You can use a trial subscription along with paid ones. The same expiration rules apply to trial and paid subscriptions.

To continue using the service after the trial subscription expires, buy the same type subscription and renew the trial subscription specifying the paid subscription. Your backed up data will be kept online. Regular backups of your machines will continue uninterrupted. The service will not need to perform a new full backup.

7.4 Activating online backup subscriptions

To be able to back up a machine to the online storage, you need to purchase and activate a subscription to the Acronis Backup & Recovery 10 Online service. You can purchase subscriptions on the Acronis Web site or from an Acronis reseller.

Before activating a subscription, please take into account the following considerations:

- As soon as a subscription is activated, its subscription period starts. To avoid losing subscription time, activate the subscription only when you are ready to back up the machine.
- If a machine already has a subscription, the new subscription will replace the old one. You can reassign the old subscription to a different machine—see “Reassigning an activated subscription” later in this section.

7.4.1 Activating subscriptions

To activate a subscription

1. Start Acronis Backup & Recovery 10.
2. In the **Actions** pane, click **Activate online backup subscription**.
3. Specify the credentials to log in to the online storage.
4. From **Available subscriptions**, select the subscription that you want to activate for the machine.
5. Click **Activate now**.

7.4.2 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 & Recovery 10 on a machine and want to resume its online backups.
- You recovered a machine to bare metal (or to a state when it did not yet have an activated subscription) and want to resume its online backups.

Reassigning a subscription does not restart its subscription period.


To assign an activated subscription to a machine

1. On the machine to which you want to assign an activated subscription, go to the subscription activation window.
2. From **Activated subscriptions**, select the activated subscription that you want to reassign to the machine.
3. Click **Activate now**.

Example

The diagram below shows what happens if you reassign a subscription to a different machine. Let's assume Machine 1 has four backups in Subscription 1. Machine 2 has two backups in Subscription 2. At that point, you reassign Subscription 1 to Machine 2. Machine 2 does its third backup to Subscription 1.

Depending on your settings, this backup will be either full or incremental. But its size is not likely to be less than a full backup size. Therefore, it is not practical to reassign a subscription to a machine whose first backup was done as an initial seeding. You will need to either redo the initial seeding (which requires a new license) or to transfer the sizeable backup over the Internet.



All earlier created backups remain intact. You can delete them manually if necessary. Keep in mind though, backups can be deleted from a subscription only by the machine to which the subscription is assigned. In our example, you have the following options.

Before reassigning

Delete backups from Subscription 1 using Machine 1 (if it is available and turned on). Delete backups from Subscription 2 using Machine 2.

After reassigning

Delete backups from Subscription 1 using Machine 2. You cannot delete backups from Subscription 2, unless you assign this subscription to any machine.

7.5 Configuring proxy settings

If the machine connects to the Internet through a proxy server, configure Acronis Backup & Recovery 10 to use the proxy server.

1. Start Acronis Backup & Recovery 10.
2. On the toolbar, on the **Options** menu, click **Machine options**.
3. Click **Online backup proxy**.
4. Enter the proxy server settings. For detailed information about the settings please refer to the context-sensitive help.

7.6 Limitations of the online storage

Unlike other types of storage available in Acronis Backup & Recovery 10, the online storage has the following limitations.

Operations

The following operations are not possible.

Backup operations:

- Backing up from bootable media
- Backing up under Linux
- Creating differential backups
- Using the **Grandfather-Father-Son (GFS)**, **Tower of Hanoi**, and **Custom** backup schemes
- Simplified naming of backup files
- Converting a backup to a virtual machine

Recovery operation:

- Recovering a backup as a virtual machine

Operations with backups:

- Exporting a backup
- Mounting a backup

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

- Exporting an archive

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

Backup and recovery options

Some backup and recovery options are not supported by online backups. For example:

- **Backup splitting** (p. 64)
- **Dual destination** (p. 67)

By using the **Backup performance > Network connection speed** option, you can vary the transferring speed as kilobytes per second, but not as a percentage.

Command-line mode

Acronis Backup & Recovery 10 command-line utilities do not support online backup.

7.7 Terminology reference

The following is the list of terms related to the Acronis Backup & Recovery 10 Online service.

Activate a subscription

Allow the machine to use the online storage according to the subscription. Subscription period starts counting down when the subscription is activated.

Activated subscription

A subscription that is currently being used by a machine.

Assign a subscription to a machine

Reserve a subscription for a particular machine. Subscription period does not start counting down until the subscription is activated.

Assigned subscription

A subscription that has been assigned to a machine.

Available subscription

A subscription that is not assigned to any machine.

Extra service

A service that you can use in addition to online backup subscriptions.

Initial Seeding

An extra service that enables you to save an initial full backup locally and then send it to Acronis on a hard disk drive. Acronis uploads the backup to the online storage. After that, you can add incremental backups to this full backup, either manually or on a schedule.

Large Scale Recovery

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

License

Not to be confused with Acronis Backup & Recovery 10 product license.

Permission for a machine to use an extra service of Acronis Backup & Recovery 10 Online.

You can buy Initial Seeding licenses and/or Large Scale Recovery licenses.

Reassign a subscription

Assign a subscription that is already activated, to a different machine.

Registration code

A character string for registering a subscription or license that was bought from an Acronis partner.

When you purchase such subscriptions or licenses, you receive a confirmation e-mail containing the registration codes for each of them. You then enter the registration codes on the account management Web page, and these subscriptions and licenses become available for use.

Renew a subscription

Assign a subscription that has the same storage quota as the current, activated subscription.

This subscription will become activated as soon as the current subscription expires.

Storage quota

The amount of storage space that a machine can use according to the subscription.

Subscription

Permission for a machine to use a specific amount of space in the online storage for a specific period of time.

Subscription period

The period during which the subscription remains activated. You can back up and recover the machine during this period. Recovery is possible for extra 30 days after this period ends.

Unassign a subscription

Make an assigned subscription available again.

You can unassign a subscription as long as it is not activated.

8 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. 234) and the machine becomes operational and ready to provide necessary services. The data required to serve incoming requests is recovered with the highest priority; everything else is recovered in the background. Limitations:

- the backup must be located on the local drive (any device available through the BIOS except for network boot)
- does not work with Linux images.

Acronis Plug-in for WinPE

A modification of Acronis Backup & Recovery 10 Agent for Windows that can run in the preinstallation environment. The plug-in can be added to a WinPE (p. 241) image using Bootable Media Builder. The resulting bootable media (p. 231) can be used to boot any PC-compatible machine and perform, with certain limitations, most of the direct management (p. 233) 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. 233).

Acronis Secure Zone

A secure volume for storing backup archives (p. 229) within a managed machine (p. 237).

Advantages:

- enables recovery of a disk to the same disk where the disk's backup resides
- offers a cost-effective and handy method for protecting data from software malfunction, virus attack, operator error
- eliminates the need for a separate media or network connection to back up or recover the data. This is especially useful for mobile users
- can serve as the primary location for dual destination backup.

Limitations: Acronis Secure Zone cannot be organized on a dynamic disk (p. 235) or a disk using the GPT partitioning style.

Acronis Secure Zone is considered as a personal vault (p. 238).

Acronis Startup Recovery Manager (ASRM)

A modification of the bootable agent (p. 231), 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.

Agent (Acronis Backup & Recovery 10 Agent)

An application that performs data backup and recovery and enables other management operations on the machine (p. 237), such as task management and operations with hard disks.

The type of data that can be backed up depends on the agent type. Acronis Backup & Recovery 10 includes the agents for backing up disks and files and the agents for backing up virtual machines residing on virtualization servers.

Agent-side cleanup

Cleanup (p. 232) performed by an agent (p. 229) according to the backup plan (p. 230) that produces the archive (p. 229). Agent-side cleanup is performed in unmanaged vaults (p. 241).

Agent-side validation

Validation (p. 241) performed by an agent (p. 229) according to the backup plan (p. 230) that produces the archive (p. 229). Agent-side validation is performed in unmanaged vaults (p. 241).

Archive

See Backup archive (p. 229).

B

Backup

A backup is the result of a single backup operation (p. 229). Physically, it is a file or a tape record that contains a copy of the backed up data as of a specific date and time. Backup files created by Acronis Backup & Recovery 10 have a TIB extension. The TIB files which are the result of a backup export (p. 236) or consolidation (p. 233) are also called backups.

Backup archive (Archive)

A set of backups (p. 229) created and managed by a backup plan (p. 230). An archive can contain multiple full backups (p. 236) as well as incremental (p. 237) and differential backups (p. 233). Backups belonging to the same archive are always stored in the same location. Multiple backup plans can back up the same source to the same archive, but the mainstream scenario is "one plan – one archive".

Backups in an archive are entirely managed by the backup plan. Manual operations with archives (validation (p. 241), viewing contents, mounting and deleting backups) should be performed using Acronis Backup & Recovery 10. Do not modify your archives using non-Acronis tools such as Windows Explorer or third-party file managers.

Backup operation

An operation that creates a copy of the data that exists on a machine's (p. 237) 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. 229), 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. 230).

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
- where to store the backup archive (p. 229) (the backup archive name and location)
- the backup scheme (p. 231), that includes the backup schedule and [optionally] the retention rules
- [optionally] the archive validation rules (p. 241)
- the backup options (p. 229).

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 to \\server\backups\ (**this is the backup archive name and location**)
- perform full backup monthly on the last day of the month at 10:00AM and 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. 240) configured for execution on a managed machine (p. 237).

A backup plan can be created directly on the machine (local plan) or appears on the machine as a result of a backup policy (p. 230) deployment (centralized plan (p. 232)).

Backup policy (Policy)

A backup plan template created by the management server (p. 238) administrator and stored on the management server. A backup policy contains the same rules as a backup plan, but might not explicitly specify what data items to back up. Instead, selection rules (p. 239), such as environment variables, can be used. Because of this flexible selection, a backup policy can be centrally applied to multiple machines. If a data item is specified explicitly (e.g. /dev/sda or C:\Windows), the policy will back up this item on each machine where this exact path is found.

By applying a policy to a group of machines, the administrator deploys multiple backup plans with a single action.

The workflow when using policies is as follows.

1. The administrator creates a backup policy.
2. The administrator applies the policy to a group of machines or a single machine (p. 237).
3. The management server deploys the policy to the machines.

4. On each machine, the agent (p. 229) installed on the machine finds data items using the selection rules. For example, if the selection rule is [All volumes], the entire machine will be backed up.
5. On each machine, the agent installed on the machine creates a backup plan (p. 230) using other rules specified by the policy. Such backup plan is called a centralized plan (p. 232).
6. On each machine, the agent installed on the machine creates a set of centralized tasks (p. 232) that will carry out the plan.

Backup scheme

A part of the backup plan (p. 230) that includes the backup schedule and [optionally] the retention rules and the cleanup (p. 232) schedule. For example, perform a full backup (p. 236) monthly on the last day of the month at 10:00AM and an incremental backup (p. 237) on Sundays at 10:00PM. Delete backups that are older than 3 months. Check for such backups every time the backup operation is completed.

Acronis Backup & Recovery 10 provides the ability to use well-known optimized backup schemes such as GFS and Tower of Hanoi, to create a custom backup scheme or to back up data once.

Bootable agent

A bootable rescue utility that includes most of the functionality of the Acronis Backup & Recovery 10 Agent (p. 229). Bootable agent is based on Linux kernel. A machine (p. 237) can be booted into a bootable agent using either bootable media (p. 231) or Acronis PXE Server. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 233).

Bootable media

A physical media (CD, DVD, USB flash drive or other media supported by a machine (p. 237) BIOS as a boot device) that contains the bootable agent (p. 231) or Windows Preinstallation Environment (WinPE) (p. 241) with the Acronis Plug-in for WinPE (p. 228). A machine can also be booted into the above environments using the network boot from Acronis PXE Server or Microsoft Remote Installation Service (RIS). 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. 236) 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 that always exists on a management server (p. 238).

A management server has two built-in groups that contain all machines of each type: All physical machines (p. 238), All virtual machines (p. 241).

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 physical machine from the built-in group except for removing the machine from the management server. Virtual machines are removed as a result of their host server removal.

A backup policy (p. 230) can be applied to a built-in group.

C

Centralized backup plan

A backup plan (p. 230) that appears on the managed machine (p. 237) as a result of deploying a backup policy (p. 230) from the management server (p. 238). Such plan can be modified only by editing the backup policy.

Centralized management

Management of the Acronis Backup & Recovery 10 infrastructure through a central management unit known as Acronis Backup & Recovery 10 Management Server (p. 238). The centralized management operations include:

- creating, applying and managing backup policies (p. 230)
- creating and managing static (p. 239) and dynamic groups (p. 235) of machines (p. 237)
- managing the tasks (p. 240) existing on the machines
- creating and managing centralized vaults (p. 232) for storing archives
- managing storage nodes (p. 239)
- monitoring activities of the Acronis Backup & Recovery 10 components, viewing the centralized log and more.

Centralized task

A task (p. 240) belonging to a centralized backup plan (p. 232). Such task appears on the managed machine (p. 237) as a result of deploying a backup policy (p. 230) from the management server (p. 238) and can be modified only by editing the backup policy.

Centralized vault

A networked location allotted by the management server (p. 238) administrator to serve as storage for the backup archives (p. 229). A centralized vault can be managed by a storage node (p. 239) or be unmanaged. The total number and size of archives stored in a centralized vault are limited by the storage size only.

As soon as the management server administrator creates a centralized vault, the vault name and path to the vault are distributed to all machines registered (p. 239) on the server. The shortcut to the vault appears on the machines in the Centralized vaults list. Any backup plan (p. 230) 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. 229) from a backup archive (p. 229) in order to get rid of outdated backups or prevent the archive from exceeding the desired size.

Cleanup consists in applying to an archive the retention rules set by the backup plan (p. 230) that produces the archive. This operation checks if the archive has exceeded its maximum size and/or for expired backups. This may or may not result in deleting backups depending on whether the retention rules are violated or not.

For more information please refer to Retention rules (p. 31).

Console (Acronis Backup & Recovery 10 Management Console)

A tool for remote or local access to Acronis agents (p. 229) and Acronis Backup & Recovery 10 Management Server (p. 238).

Having connected the console to the management server, the administrator sets up and manages backup policies (p. 230) and accesses other management server functionality, that is, performs centralized management (p. 232). Using the direct console-agent connection, the administrator performs direct management (p. 233).

Consolidation

Combining two or more subsequent backups (p. 229) belonging to the same archive (p. 229) into a single backup.

Consolidation might be needed when deleting backups, either manually or during cleanup (p. 232). For example, the retention rules require to delete a full backup (p. 236) that has expired but retain the next incremental (p. 237) 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

Deduplicating vault

A managed vault (p. 237) in which deduplication (p. 233) is enabled.

Deduplication

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

Acronis Backup & Recovery 10 can apply the deduplication technology to backup archives (p. 229) stored on storage nodes (p. 239). This minimizes storage space taken by the archives, backup traffic and network usage during backup.

Differential backup

A differential backup stores changes to the data against the latest full backup (p. 236). You need access to the corresponding full backup to recover the data from a differential backup.

Direct management

Any management operation that is performed on a managed machine (p. 237) using the direct console (p. 233)-agent (p. 229) connection (as opposed to centralized management (p. 232) when the operations are configured on the management server (p. 238) and propagated by the server to the managed machines).

The direct management operations include:

- creating and managing local backup plans (p. 237)
- creating and managing local tasks (p. 237), such as recovery tasks
- creating and managing personal vaults (p. 238) and archives stored there
- viewing the state, progress and properties of the centralized tasks (p. 232) 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. 231). Some of the direct management operations can also be performed via the management server GUI. This presumes, however, either an explicit or an implicit direct connection to the selected machine.

Disk backup (Image)

A backup (p. 229) that contains a sector-based copy of a disk or a volume in a packaged form. Normally, only sectors that contain data are copied. Acronis Backup & Recovery 10 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. 235) 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. 237) 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:

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. 237) which is populated automatically by the management server (p. 238) according to membership criteria specified by the administrator. Acronis Backup & Recovery 10 offers the following membership criteria:

- Operating system
- Active Directory organization unit
- IP address range.

A machine remains in a dynamic group as long as the machine meets the group's criteria. The machine is removed from the group automatically as soon as

- the machine's properties change so that the machine does not meet the criteria anymore OR

- the administrator changes the criteria so that the machine does not meet them anymore.

There is no way to remove a machine from a dynamic group manually except for deleting the machine from the management server.

Dynamic volume

Any volume located on dynamic disks (p. 235), or more precisely, on a disk group (p. 234). 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. 229) 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. 229), each backup belonging to the archive is encrypted by the agent (p. 229) before saving the backup to its destination.

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 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.

Encrypted vault

A managed vault (p. 237) to which anything written is encrypted and anything read is decrypted transparently by the storage node (p. 239), using a vault-specific encryption key stored on the node. In case the storage medium is stolen or accessed by an unauthorized person, the malefactor will not be able to decrypt the vault contents without access to the storage node. Encrypted archives (p. 236) will be encrypted over the encryption performed by the agent (p. 229).

Export

An operation that creates a copy of an archive (p. 229) 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. 229) or to your choice of backups belonging to the same archive. An entire vault (p. 241) can be exported by using the command line interface.

F

Full backup

A self-sufficient backup (p. 229) 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. 231) aimed to maintain the optimal balance between a backup archive (p. 229) size and the number of recovery points (p. 238) available from the archive. GFS enables recovering with daily resolution for the last several days, weekly resolution for the last several weeks and monthly resolution for any time in the past.

For more information please refer to GFS backup scheme (p. 25).

I

Image

The same as Disk backup (p. 234).

Incremental backup

A backup (p. 229) that stores changes to the data against the latest backup. You need access to other backups from the same archive (p. 229) to restore data from an incremental backup.

L

Local backup plan

A backup plan (p. 230) created on a managed machine (p. 237) using direct management (p. 233).

Local task

A task (p. 240) belonging to a local backup plan (p. 237) or a task that does not belong to any plan, such as a recovery task. A local task belonging to a backup plan can be modified by editing the plan only; other local tasks can be modified directly.

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. 237), either physical or virtual, where at least one Acronis Backup & Recovery 10 Agent (p. 229) is installed.

Managed vault

A centralized vault (p. 232) managed by a storage node (p. 239). Archives (p. 229) 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 storage node-side cleanup (p. 239) and storage node-side validation (p. 240) for each archive stored in the managed vault. An administrator can specify additional operations that the storage node will perform (deduplication (p. 233), encryption).

Any managed vault is self-contained, that is, contains all metadata the storage node needs to manage the vault. In case the storage node is lost or its database is corrupted, the new storage node retrieves the metadata and re-creates the database. When the vault is attached to another storage node, the same procedure takes place.

Management server (Acronis Backup & Recovery 10 Management Server)

A central server that drives data protection within the enterprise network. Acronis Backup & Recovery 10 Management Server provides the administrator with:

- a single entry point to the Acronis Backup & Recovery 10 infrastructure
- an easy way to protect data on numerous machines (p. 237) using backup policies (p. 230) and grouping
- enterprise-wide monitoring functionality
- the ability to create centralized vaults (p. 232) for storing enterprise backup archives (p. 229)
- the ability to manage storage nodes (p. 239).

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. 231).

P

Personal vault

A local or networked vault (p. 241) created using direct management (p. 233). Once a personal vault is created, a shortcut to it appears under the **Personal vaults** item of the **Navigation** pane. Multiple machines can use the same physical location; for example, a network share; as a personal vault.

Physical machine

On Acronis Backup & Recovery 10 Management Server, a physical machine is the same as a registered machine (p. 239). A virtual machine is considered physical if an Acronis Backup & Recovery 10 agent is installed on the machine and the machine is registered on the management server.

Plan

See Backup plan (p. 230).

Policy

See Backup policy (p. 230).

R

Recovery point

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

Registered machine

A machine (p. 237) managed by a management server (p. 238). A machine can be registered on only one management server at a time. A machine becomes registered as a result of the registration (p. 239) procedure.

Registration

A procedure that adds a managed machine (p. 237) to a management server (p. 238).

Registration sets up a trust relationship between the agent (p. 229) 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).

S

Selection rule

A part of the backup policy (p. 230). Enables the management server (p. 238) administrator to select the data to back up within a machine.

Static group

A group of machines which a management server (p. 238) administrator populates by manually adding machines to the group. A machine remains in a static group until the administrator removes it from the group or from the management server.

Storage node (Acronis Backup & Recovery 10 Storage Node)

A server aimed to optimize usage of various resources required for protection of enterprise data. This goal is achieved by organizing managed vaults (p. 237). Storage node enables the administrator to:

- relieve managed machines (p. 237) of unnecessary CPU load by using the storage node-side cleanup (p. 239) and storage node-side validation (p. 240)
- drastically reduce backup traffic and storage space taken by the archives (p. 229) by using deduplication (p. 233)
- prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted vaults (p. 236).

Storage node-side cleanup

Cleanup (p. 232) performed by a storage node (p. 239) according to the backup plans (p. 230) that produce the archives (p. 229) stored in a managed vault (p. 237). Being an alternative to the agent-

side cleanup (p. 229), the cleanup on the storage node side relieves the production servers of unnecessary CPU load.

Since the cleanup schedule exists on the machine (p. 237) the agent (p. 229) resides on, and therefore uses the machine's time and events, the agent has to initiate the storage node-side cleanup every time the scheduled time or event comes. To do so, the agent must be online.

The following table summarizes the cleanup types used in Acronis Backup & Recovery 10.

	Cleanup	
	Agent-side	Storage node-side
Applied to:	Archive	Archive
Initiated by:	Agent	Agent
Performed by:	Agent	Storage node
Schedule set by:	Backup plan	Backup plan
Retention rules set by:	Backup plan	Backup plan

Storage node-side validation

Validation (p. 241) performed by a storage node (p. 239) according to the backup plans (p. 230) that produce the archives (p. 229) stored in a managed location (p. 237). Being an alternative to the agent-side validation (p. 229), the validation on the storage node side relieves the production servers of unnecessary CPU load.

T

Task

In Acronis Backup & Recovery 10, a task is a set of sequential actions to be performed on a managed machine (p. 237) when a certain time comes or a certain event occurs. The actions are described in an xml script file. The start condition (schedule) exists in the protected registry keys.

Tower of Hanoi

A popular backup scheme (p. 231) aimed to maintain the optimal balance between a backup archive (p. 229) size and the number of recovery points (p. 238) available from the archive. Unlike the GFS (p. 236) 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. 29)".

U

Universal Restore (Acronis Backup & Recovery 10 Universal Restore)

The Acronis proprietary technology that helps boot up Windows on dissimilar hardware or a virtual machine. The Universal Restore handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

The Universal Restore is not available:

- when the machine is booted with Acronis Startup Recovery Manager (p. 228) (using F11) or
- the image being recovered is located in Acronis Secure Zone (p. 228) or
- when using Acronis Active Restore (p. 228),

because these features are primarily meant for instant data recovery on the same machine.

Universal Restore is not available when recovering Linux.

Unmanaged vault

Any vault (p. 241) that is not a managed vault (p. 237).

V

Validation

An operation that checks the possibility of data recovery from a backup (p. 229).

Validation of a file backup imitates recovery of all files from the backup to a dummy destination. The previous product versions considered a file backup valid when the metadata contained in its header was consistent. The current method is time-consuming but much more reliable. Validation of a volume backup calculates a checksum for every data block saved in the backup. This procedure is also 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.

Validation rules

A part of the backup plan (p. 230). Rules that define when and how often to perform validation (p. 241) and whether to validate the entire archive (p. 229) or the latest backup in the archive.

Vault

A place for storing backup archives (p. 229). 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. 232), but the total size of archives stored in the vault is limited by the storage size only.

Virtual machine

On Acronis Backup & Recovery 10 Management Server, a machine (p. 237) is considered virtual if it can be backed up from the virtualization host without installing an agent (p. 229) on the machine. A virtual machine appears on the management server after registration of the virtualization server that hosts the machine, provided that Acronis Backup & Recovery 10 agent for virtual machines is installed on that server.

W

WinPE (Windows Preinstallation Environment)

A minimal Windows system based on any of the following kernels:

- Windows XP Professional with Service Pack 2 (PE 1.5)
- Windows Server 2003 with Service Pack 1 (PE 1.6)
- Windows Vista (PE 2.0)
- Windows Vista SP1 and Windows Server 2008 (PE 2.1).

WinPE is commonly used by OEMs and corporations for deployment, test, diagnostic and system repair purposes. A machine can be booted into WinPE via PXE, CD-ROM, USB flash drive or hard disk. The Acronis Plug-in for WinPE (p. 228) enables running the Acronis Backup & Recovery 10 Agent (p. 229) in the preinstallation environment.