

Acronis[®] Backup & Recovery 10[®] Advanced Server Virtual Edition

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	Int	roducing A	Acronis® Backup & Recovery® 10	8
	1.1	Acronis B	Backup & Recovery 10 overview	8
	1.2	Getting s	started	9
	1.2	1 Using	the management console	11
	1.3	Acronis B	Backup & Recovery 10 components	18
	1.3	_	t for Windows	
	1.3 1.3	_	t for Linuxt for ESX/ESXi	
	1.3	_	onents for centralized management	
	1.3	5 Mana	gement Console	23
	1.3		able Media Builder	
	1.3 1.4		nis Wake-on-LAN Proxyed file systems	
			•	
	1.5		ed operating systems	
	1.6	•	equirements	
	1.7	Technica	l Support	28
2	Un	derstandiı	ng Acronis Backup & Recovery 10	29
	2.1	Basic con	ncepts	29
	2.2	User priv	vileges on a managed machine	33
	2.3	Owners a	and credentials	34
	2.4	Full, incre	emental and differential backups	35
	2.5	GFS back	sup scheme	37
	2.6	Tower of	Hanoi backup scheme	40
	2.7	Retention	n rules	43
	2.8	Backing u	up dynamic volumes (Windows)	45
	2.9	Backing u	up LVM volumes and MD devices (Linux)	47
	2.9		ng up logical volumes	
	2.9		ng up MD devices	
	2.9 2.9		g the volume structure information ting logical volumes and MD devices in command line	
	2.10		up hardware RAID arrays (Linux)	
	2.11	_	up virtual machines	
	2.1	•	to install Hyper-V Integration Services	
	2.1	L.2 How t	to install VMware Tools	53
	2.12		port	
	2.1	•	compatibility table	
	2.13	_	a single tape drivefor SNMP	
	2.14	-	ary Acronis technologies nis Secure Zone	
	2.1		nis Secure Zoneis Secure ZoneManager	
	2.1		ersal Restore (Acronis Backup & Recovery 10 Universal Restore)	

2.14.4	Acronis Active Restore	59
2.15 L	Inderstanding centralized management	61
2.15.1	Basic concepts	61
2.15.2	Setting up centralized data protection in a heterogeneous network	62
2.15.3	Grouping the registered machines	66
2.15.4	Policies on machines and groups	68
2.15.5	Backup policy's state and statuses	73
2.15.6	·	
2.15.7	-0	
2.15.8	·	
•	ons	
3.1 C	Console options	
3.1.1	Startup page	
3.1.2	Pop-up messages	
3.1.3	Time-based alerts	•
3.1.4	Number of tasks	
3.1.5	Fonts	
	Management server options	
3.2.1	Logging level	
3.2.2	Log cleanup rules	
3.2.3	Event tracing	
3.2.4 3.2.5	Domain access credentials	
3.2.6	VM protection options	
3.2.7	Online backup proxy	
	Nachine options	
3.3.1	Machine management	
3.3.2	Event tracing	
3.3.3	Log cleanup rules	
3.3.4	Online backup proxy	
3.3.5	Customer Experience Program	
3.4 D	Default backup and recovery options	
3.4.1	Default backup options	
3.4.2	Default recovery options	
4 Vault	rs	
4.1 C	Centralized vaults	136
4.1.1	Working with the "Centralized vault" view	
4.1.2	Actions on centralized vaults	
4.1.3	Tape libraries	
4.2 P	ersonal vaults	166
4.2.1	Working with the "Personal vault" view	
4.2.2	Actions on personal vaults	
	Common operations	
4.3.1	Operations with archives stored in a vault	
4.3.1	Operations with backups	
4.3.3	Deleting archives and backups	
4.3.4	Filtering and sorting archives	
5 Sche	duling	173
	Paily schedule	
J.1 L	any serieuare	

	5.2 V	Veekly schedule	176
	5.3 N	Ionthly schedule	178
	5.4 A	t Windows Event Log event	180
		dvanced scheduling settings	
		pon an Acronis Drive Monitor alert	
		onditions	
	5.7.1	User is idle	
	5.7.2	Location's host is available	
	5.7.3	Fits time interval	
	5.7.4	User logged off	
	5.7.5	Time since last backup	
6	Direc	t management	188
	6.1 A	dministering a managed machine	188
	6.1.1	Dashboard	188
	6.1.2	Backup plans and tasks	190
	6.1.3	Log	201
	6.2 C	reating a backup plan	203
	6.2.1	Why is the program asking for the password?	
	6.2.2	Backup plan's credentials	
	6.2.3	Label (Preserving machine properties in a backup)	
	6.2.4	Source type	208
	6.2.5	Items to back up	209
	6.2.6	Access credentials for source	
	6.2.7	Exclusions	211
	6.2.8	Archive	
	6.2.9	Simplified naming of backup files	
	6.2.10		
	6.2.11	•	
	6.2.12		
	6.2.13		
	6.3 R	ecovering data	
	6.3.1	Task credentials	
	6.3.2	Archive selection	
	6.3.3	Data type	
	6.3.4	Content selection	
	6.3.5	Access credentials for location	
	6.3.6 6.3.7	Destination selection	
	6.3.8	When to recover	
	6.3.9	Universal Restore	
	6.3.10		
	6.3.11	•	
	6.3.12	,	
	6.3.13	., ,	
	6.3.14	·	
	6.4 V	alidating vaults, archives and backups	
	6.4.1	Task credentials	
	6.4.2	Archive selection	
	6.4.3	Backup selection	
	6.4.4	Location selection	
	6.4.5	Access credentials for source	

6.4	1.6	When to validate	255
6.5	М	ounting an image	255
6.5	5.1	Archive selection	256
6.5	5.2	Backup selection	257
6.5	5.3	Access credentials	257
6.5		Volume selection	
6.6	M	anaging mounted images	258
6.7	Ex	porting archives and backups	259
6.7	7.1	Task credentials	261
6.7	7.2	Archive selection	262
6.7	7.3	Backup selection	
6.7		Access credentials for source	
6.7		Location selection	
6.7		Access credentials for destination	
6.8	Ac	ronis Secure Zone	
6.8		Creating Acronis Secure Zone	
6.8		Managing Acronis Secure Zone	
6.9		ronis Startup Recovery Manager	
6.10	Вс	ootable media	269
6.1	l0.1	How to create bootable media	
6.1	L0.2	Connecting to a machine booted from media	
_	10.3	Working under bootable media	
_	10.4	List of commands and utilities available in Linux-based bootable media	
_	LO.5 LO.6	Recovering MD devices and logical volumes Acronis PXE Server	
		sk management	
		Basic precautions	
	l1.1 l1.2	Running Acronis Disk Director Lite	
_	11.3	Choosing the operating system for disk management	
	11.4	"Disk management" view	
6.1	1.5	Disk operations	289
6.1	1.6	Volume operations	295
6.1	L1.7	Pending operations	302
6.12	Co	llecting system information	302
7 Ce	ntra	alized management	303
7.1		Iministering Acronis Backup & Recovery 10 Management Server	
		Dashboard	
7.1 7.1		Backup policies	
7.1		Physical machines	
7.1		Virtual machines	
7.1	L.5	Storage nodes	
7.1	L.6	Tasks	339
7.1	L.7	Log	
7.1	8 .1	Reporting	345
7.2	Co	onfiguring Acronis Backup & Recovery 10 components	351
7.2	2.1	Parameters set through administrative template	351
7.2	2.2	Parameters set through GUI	
7.2	2.3	Parameters set through Windows registry	366
7.3	Cr	eating a backup policy	367
7.3	3.1	Policy credentials	370

	7.3.2	Items to back up	370
	7.3.3	Access credentials for source	375
	7.3.4	Exclusions	375
	7.3.5	Archive	377
	7.3.6	Access credentials for location	378
	7.3.7	Backup scheme selection	
	7.3.8	Archive validation	388
8	Online	e backup	389
8	3.1 In	troduction to Acronis Backup & Recovery 10 Online	389
	8.1.1	What is Acronis Backup & Recovery 10 Online?	389
	8.1.2	What data can I back up and recover?	390
	8.1.3	How long will my backups be kept in the online storage?	390
	8.1.4	How to secure my data?	390
	8.1.5	How to back up virtual machines to the online storage?	390
	8.1.6	Backup and recovery FAQ	
	8.1.7	Initial Seeding FAQ	392
	8.1.8	Large Scale Recovery FAQ	397
	8.1.9	Subscription lifecycle FAQ	399
8	3.2 W	here do I start?	401
8	3.3 Ch	noosing a subscription	401
8	3.4 Ac	ctivating online backup subscriptions	402
	8.4.1	Activating subscriptions	403
	8.4.2	Reassigning an activated subscription	403
8	3.5 Co	onfiguring proxy settings	404
8	3.6 Lir	mitations of the online storage	405
8	3.7 Te	rminology reference	406
9	Glossa	ary	408
		•	

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 Advanced Server Virtual Edition 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
- Scalability from a single machine to an enterprise
- Support for both Windows and Linux environments
- Centralized management for distributed workstations and servers
- Dedicated servers for storage resource optimization.

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

Leveraging existing IT infrastructure

Data deduplication to reduce storage consumption and network bandwidth utilization Flexible deduplication mechanism allowing deduplication of backup data both at the source and at the storage

Improved support for robotic tape libraries

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

Policy-based centralized management

Applying backup policies to groups of machines

Static and dynamic machine grouping

Grouping either physical or virtual machines

Easy work with virtual environments

Backup and recovery of virtual machines without installing backup software on each individual machine

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

Restricted remote access to a backup agent

Secure communication between the product components

Using third-party certificates for authentication of the components

Data encryption options for both data transmission and storage

Backup of remote machines to a centralized storage node behind firewalls.

1.2 Getting started

Direct management

- 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.

Linux

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

/usr/sbin/acronis_console

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. 29)".

For understanding of the GUI elements see "Using the management console (p. 11)".

For how to enable non-root users to start the console under Linux see "Privileges for local connection (p. 81)".

For how to enable remote connection to a machine running Linux see "Privileges for remote connection in Linux (p. 82)".

Centralized management

We recommend that you first try to manage the single machine using direct management as described above.

To start with centralized management:

- 1. Install Acronis Backup & Recovery 10 Management Server (p. 20).
- Install Acronis Backup & Recovery 10 Agents on the machines that need data protection. When
 installing the agents, register each of the machines on the management server. To do so, enter
 the server's IP or name and the centralized administrator's credentials in one of the installation
 wizard's windows.

- 3. Install Acronis Backup & Recovery 10 Management Console (p. 23) on the machine from which you prefer to operate. We recommend that you use the console that installs on Windows if you have a choice between Windows and Linux console distributions. Install Acronis Bootable Media Builder.
- 4. Start the console. Create the bootable media.
- 5. Connect the console to the management server.

The simplified way of centralized management

Backup

Using the **Back up** control, select the machine which you want to back up and then create a backup plan (p. 410) on the machine. You can create backup plans on multiple machines in turn.

Recovery

Using the **Recover** control, select the machine where the data recovery is required and create a recovery task on the machine. You can create recovery tasks on multiple machines in turn.

To recover the entire machine or the operating system that fails to start, use the bootable media (p. 411). You cannot control operations under bootable media using the management server, but you can disconnect the console from the server and connect it to the machine booted from the media.

Managing plans and tasks

To manage the plans and tasks existing on the registered machines, select **Machines > All machines** in the **Navigation** tree and then select each machine in turn. The **Information** pane below shows the state and the details of plans and tasks existing on each machine and enables you to start, stop, edit, and delete the plans and tasks.

You can also use the **Tasks** view that displays all tasks existing on the registered machines. The tasks can be filtered by machines, backup plans and other parameters. Refer to the context help for details.

Viewing log

To view the centralized log, collected from the registered machines, select **Log** in the **Navigation** tree. The log entries can be filtered by machines, backup plans and other parameters. Refer to the context help for details.

Creating centralized vaults

If you opt for storing all backup archives in a single or a few networked locations, create centralized vaults in these locations. After a vault is created, you can view and administer its content by selecting **Vaults > Centralized > 'Vault name'** in the **Navigation** tree. The shortcut to the vault will be deployed to all the registered machines. The vault can be specified as a backup destination in any backup plan created by you or by the registered machines' users.

The advanced way of centralized management

To make the best use of the centralized management capabilities offered by Acronis Backup & Recovery 10, you can opt for:

Using deduplication

- 1. Install Acronis Backup & Recovery 10 Storage Node (p. 21) and add it to the management server.
- 2. Create the deduplicating managed vault on the storage node.

- 3. Install the Acronis Deduplication add-on to the agent on all machines that will back up to the deduplicating vault.
- 4. Ensure that the backup plans you create use the managed vault as destination for the backup archives.

Creating a backup policy rather than backup plans

Set up a centralized backup policy and apply it to the **All machines** group. This way you will deploy backup plans on each machine with a single action. Select **Actions > Create backup policy** from the top menu and then refer to the context help.


Grouping the machines registered on the management server

Group the registered machines by appropriate parameters, create several policies and apply each policy to the appropriate group of machines. For more information please refer to "Grouping the registered machines (p. 66)".

The comprehensive example of advanced centralized management is provided in the "Setting up centralized data protection in a heterogeneous network (p. 62)" section.

1.2.1 Using the management console

As soon as the console connects to a managed machine (p. 417) or to a management server (p. 418), 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. 12) 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. 13) 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. 15) 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.

Views for a management server

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

- [Management server name]. Root of the tree also called a Welcome view. Displays the name of the management server the console is currently connected to. Use this view for quick access to the main operations, available on the management server.
 - **Dashboard**. Use this view to estimate at a glance whether the data is successfully protected on the machines registered on the management server.
 - Sackup policies. Use this view to manage backup policies existing on the management server.
 - Physical machines. Use this view to manage machines registered on the management server.
 - **Virtual machines**. Use this view to manage virtual machines from the registered physical machines and from the registered machines with the agent for ESX/ESXi.
 - **Vaults**. Use this view to manage centralized vaults and archives stored in there: create new managed and unmanaged vaults, rename and delete the existing ones.
 - Storage nodes. Use this view to manage storage nodes. Add a storage node to be able to create centralized vaults that will be managed by the node.
 - **Tasks**. Use this view to manage tasks, run, edit, stop and delete tasks, monitor their states, examine task history.
 - Log. Use this view to examine the history of centralized management operations, such as creating a managed entities group, applying a policy, managing a centralized vault; as well as the history of operations logged in the local logs of the registered machines and the storage nodes.

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]**.
 - If the console has ever been connected to Acronis Management Server, the shortcut is added automatically as **AMS [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. 16). 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. 16) 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 (<a> - for the **Navigation** pane; <a> - 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. 12).



"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.



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 full list of Acronis Backup & Recovery 10 components with a brief description of their functionality.

Acronis Backup & Recovery 10 includes the following main types of components.

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.

Components for centralized management

These components, delivered with the advanced editions, provide the centralized management capability. Usage of these components is not licensed.

Console

The console provides Graphical User Interface and remote connection to the agents and other Acronis Backup & Recovery 10 components. 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.

1.3.1.2 Deduplication

This add-on enables the agent to back up data to deduplicating vaults managed by Acronis Backup & Recovery 10 Storage Node.

1.3.1.3 Agent for Hyper-V

Acronis Backup & Recovery 10 Agent for Hyper-V protects virtual machines residing on a Hyper-V virtualization server. The agent allows for backing up virtual machines from the host without having to install agents on each virtual machine. The agent installs on Windows 2008 Server x64 (any edition) or Microsoft Hyper-V Server 2008 as an add-on to Acronis Backup & Recovery 10 Agent for Windows.

Acronis Backup & Recovery 10 Agent for Hyper-V cannot be installed with Acronis Online Backup license.

Integration services (p. 53) have to be installed on the guest systems.

1.3.2 Agent for Linux

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

Disk backup

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

File backup

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

1.3.2.1 Deduplication

This add-on enables the agent to back up data to deduplicating vaults managed by Acronis Backup & Recovery 10 Storage Node.

1.3.3 Agent for ESX/ESXi

Acronis Backup & Recovery 10 Agent for ESX/ESXi protects virtual machines residing on a VMware ESX or ESXi virtualization server. The agent allows for backing up virtual machines from the host without having to install agents on each virtual machine.

The agent is delivered as a virtual appliance.

1.3.4 Components for centralized management

This section lists the components included in the Acronis Backup & Recovery 10 editions that provide the centralized management capability. Besides these components, Acronis Backup & Recovery 10 Agents have to be installed on all machines that need data protection.

1.3.4.1 Management Server

Acronis Backup & Recovery 10 Management Server is the central server that drives data protection within the enterprise network. The 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. 417) using backup policies (p. 410) and grouping
- enterprise-wide monitoring and reporting functionality
- the ability to create centralized vaults (p. 412) for storing enterprise backup archives (p. 409)
- the ability to manage storage nodes (p. 419).

If there are multiple management servers on the network, they operate independently, manage different machines and use different centralized vaults for storing archives.

The management server's databases

The management server uses three Microsoft SQL databases:

- The configuration database that stores the list of registered machines and other configuration information, including backup policies created by the administrator.
- The synchronization database used for synchronization of the management server with registered machines and storage nodes. This is a database with rapidly changing operational data.
- The reporting database that stores the centralized log. This database may grow large. Its size depends on the logging level you set.

The configuration and synchronization databases should reside on the same Microsoft SQL Server (called an operational server) preferably installed on the same machine as the management server. The reporting database can be configured on the same or different SQL server.

When installing a management server, you can select for both operational and reporting servers what server to use. The following options are available:

- Microsoft SQL Server 2005 Express that comes with the installation package and installs on the same machine. In this case, an SQL server instance with three databases will be created on the machine.
- 2. Microsoft SQL Server 2008 (any edition) previously installed on any machine.
- 3. Microsoft SQL Server 2005 (any edition) previously installed on any machine.

VMware vCenter integration

This feature provides the capability to view virtual machines managed by a VMware vCenter Server in the management server GUI, view the backup status of these machines in the vCenter, and automatically register virtual machines created by Acronis Backup & Recovery 10 in the vCenter.

Integration is available in all Acronis Backup & Recovery 10 advanced editions; a license for Virtual Edition is not required. No software installation is required on the vCenter Server.

This feature also enables automatic deployment and configuration of Agent for ESX/ESXi to any ESX/ESXi server, that is not necessarily managed by the vCenter.

1.3.4.2 Storage Node

Acronis Backup & Recovery 10 Storage Node is a server aimed to optimize usage of various resources (such as the corporate storage capacity, the network bandwidth, or the managed machines' CPU load) required for the enterprise data protection. This goal is achieved through organizing and managing the locations that serve as dedicated storages of the enterprise backup archives (managed vaults).

The storage nodes enable creating highly scalable and flexible, in terms of the hardware support, storage infrastructure. Up to 20 storage nodes can be set up, each being able to manage up to 20 vaults. The administrator controls the storage nodes centrally from the Acronis Backup & Recovery 10 Management Server (p. 418). Direct console connection to a storage node is not possible.

Setting up the storage infrastructure

Install the storage nodes, add them to the management server (the procedure is similar to the managed machine registration (p. 419)) and create centralized vaults (p. 412). When creating a

centralized vault, specify the path to the vault, the storage node that will manage the vault, and the management operations to be performed on the vault.

A managed vault can be organized:

- on the hard drives local to the storage node
- on a network share
- on a Storage Area Network (SAN)
- on a Network Attached Storage (NAS)
- on a tape library locally attached to the storage node.

The management operations are as follows.

Storage node-side cleanup and validation

Archives, stored in unmanaged vaults, are maintained by the agents (p. 409) that create the archives. This means that each agent not only backs up data to the archive, but also executes service tasks that apply to the archive, the retention rules and validation rules specified by the backup plan (p. 410). To relieve the managed machines of unnecessary CPU load, execution of the service tasks can be delegated to the storage node. Since the tasks' schedule exists on the machine the agent resides on, and therefore uses that machine's time and events, the agent has to initiate the storage node-side cleanup (p. 420) and the storage node-side validation (p. 420) according to the schedule. To do so, the agent must be online. Further processing is performed by the storage node.

This functionality cannot be disabled in a managed vault. The next two operations are optional.

Deduplication

A managed vault can be configured as a deduplicating vault. This means that identical data will be backed up to this vault only once to minimize the network usage during backup and storage space taken by the archives. For more information, please see the "Deduplication (p. 76)" section in the User Guide.

Encryption

A managed vault can be configured so that anything written to it is encrypted and anything read from it is decrypted transparently by the storage node, using a vault-specific encryption key stored on the node server. 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 this specific storage node.

If the archive is already encrypted by the agent, the storage node-side encryption is applied over the encryption performed by the agent.

1.3.4.3 PXE Server

Acronis PXE Server allows for booting machines into Acronis bootable components through the network.

The network booting:

- Eliminates the need to have a technician onsite to install the bootable media (p. 411) into the system that has to be booted
- During group operations, reduces the time required for booting multiple machines as compared to using physical bootable media.

1.3.4.4 License Server

The server enables you to manage licenses of Acronis products and install the components that require licenses.

For more information about Acronis License Server please see "Using Acronis License Server".

1.3.5 Management Console

Acronis Backup & Recovery 10 Management Console is an administrative tool for remote or local access to Acronis Backup & Recovery 10 agents, and in the product editions that include the centralized management capability, to the Acronis Backup & Recovery 10 Management Server.

The console has two distributions for installation on Windows and installation on Linux. While both distributions enable connection to any Acronis Backup & Recovery 10 agent and Acronis Backup & Recovery 10 Management Server, we recommend that you use the console for Windows if you have a choice between the two. The console that installs on Linux has limited functionality:

- remote installation of Acronis Backup & Recovery 10 components is not available
- the Active Directory-related features, such as browsing the AD, are not available.

1.3.6 Bootable Media Builder

Acronis Bootable Media Builder is a dedicated tool for creating bootable media (p. 411). There are two media builder distributions for installation on Windows and installation on Linux.

The media builder that installs on Windows can create bootable media based on either Windows Preinstallation Environment, or Linux kernel. The Universal Restore (p. 19) 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.

The media builder that installs on Linux creates bootable media based on Linux kernel.

The Deduplication (p. 19) add-on enables you to create bootable media with the back up to a deduplicating vault functionality. This add-on can be installed to either of the media builder distributions.

1.3.7 Acronis Wake-on-LAN Proxy

Acronis Wake-on-LAN Proxy enables Acronis Backup & Recovery 10 Management Server to wake up for backup machines located in another subnet. Acronis Wake-on-LAN Proxy installs on any server in the subnet where the machines to be backed up are located.

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 License Server

- 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 except for Vista Home Basic and Vista Home Premium (x86, x64)
- Windows 7 SP1 all editions except for the Starter and Home 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

Acronis Backup & Recovery 10 Management Console

- 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

Acronis Backup & Recovery 10 Management Server and Acronis Backup & Recovery 10 Storage Node

Windows XP Professional SP3 (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 except for Vista Home Basic and Vista Home Premium (x86, x64)
- Windows 7 SP1* all editions except for the Starter and Home 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*
- * Acronis Backup & Recovery 10 Storage Node handles tape libraries and autoloaders by using Removable Storage Management (RSM). Since Windows 7, Windows Server 2008 R2 and Windows MultiPoint Server 2010 do not support RSM, a storage node installed in these operating systems does not support tape libraries and autoloaders.

Acronis Backup & Recovery 10 Agent for ESX/ESXi

- VMware ESX Infrastructure 3.5 Update 2+
- VMware ESX/ESXi 4.0 and 4.1

The Agent for ESX/ESXi is delivered as a virtual appliance.

The Agent for ESX/ESXi supports all VMware ESXi licenses except for the free license. This is because the agent uses the Remote Command Line appliance, and free VMware ESXi restricts access to this appliance to read-only access. The agent works during the VMware ESXi evaluation period. Once a free VMware ESXi serial key is entered, the Agent for ESX/ESXi will stop functioning.

Acronis Backup & Recovery 10 Agent for Hyper-V

- Windows Server 2008/2008 R2 (x64) with Hyper-V
- Microsoft Hyper-V Server 2008/2008 R2

This agent installs on a Hyper-V host as an add-on to Acronis Backup & Recovery 10 Agent for Windows.

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 except for Vista Home Basic and Vista Home Premium (x86, x64)
- Windows 7 SP1 all editions except for the Starter and Home 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

Acronis Backup & Recovery 10 Agent for Linux

- Linux with kernel 2.4.20 or later (including 2.6.x kernels) and glibc 2.3.2 or later
- Various 32-bit and 64-bit Linux distributions, including:
 - Red Hat Enterprise Linux 4.x and 5.x
 - Red Hat Enterprise Linux 6
 - Ubuntu 9.04 (Jaunty Jackalope), 9.10 (Karmic Koala) and 10.04 (Lucid Lynx)
 - Fedora 11 and 12
 - SUSE Linux Enterprise Server 10 and 11
 - Debian 4 (Lenny) and 5 (Etch)
 - CentOS 5
- Agent for Linux is in fact a 32-bit executable. For authentication, the agent uses system libraries, 32-bit versions of which are not always installed by default with 64-bit distributions. When using the agent on a 64-bit RedHat based distribution (such as RHEL, CentOS, Fedora), or on a 64-bit SUSE distribution, make sure that the following 32-bit packages are installed in the system:

pam.i386 libselinux.i386 libsepol.i386

These packages should be available in the repository of your Linux distribution.

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

apt-get install rpm

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 operating systems

Component	Memory (above the OS and running applications)	Disk space required during installation or update	Disk space occupied by the component(s)	Additional
The components installed in	Windows			
Complete installation	300 MB	2.7 GB	1.7 GB including SQL Express Server	
Agent for Windows	120 MB	700 MB	260 MB	
Agent for Hyper-V (add-on)	See "Agent for Windows"	50 MB	20 MB	

Bootable Media Builder	80 MB	700 MB	300 MB	CD-RW or DVD- RW drive
Management Console	30 MB	950 MB	450 MB	Screen resolution 1024*768 pixels or higher
Management Server	40 MB	250 MB 400 MB for SQL Express Server	250 MB 400 MB for SQL Express Server	
Wake-on-LAN Proxy	Negligible	30 MB	5 MB	
Storage Node	100 MB	150 MB	When using a tape library, space required for tapes database: approx. 1 MB per 10 archives	Recommended hardware: 4 GB RAM High speed storage such as hardware RAID
License Server	Negligible	25 MB	25 MB	
PXE Server	5 MB	80 MB	15 MB	
The components installed in	ı Linux			
Complete installation	160 MB	400 MB	250 MB	
Agent for Linux	65 MB	150 MB	70 MB	
Bootable Media Builder	70 MB	240 MB	140 MB	
Management Console	25 MB	100 MB	40 MB	
The components installed o	n VMware ESX(i) ser	ver		
Agent for ESX/ESXi Virtual Appliance	512 MB (the Virtual Appliance memory setting)	5 GB	5 GB	CPU reservation: minimum 300 MHz recommended In a vCenter cluster, a shared storage is required

Network interface card or virtual network adapter is a common requirement for all the components.

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.com/support/

Product Updates

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

2 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. 409) on the machine which becomes a managed machine (p. 417) from this point on.
- 2. To be able to manage the machine using Graphical User Interface, install Acronis Backup & Recovery 10 Management Console (p. 413) 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. 411).
- 4. Connect the console to the managed machine.
- 5. Create a backup plan (p. 410).
 - To do so, you have to specify, at the very least, the data to be protected and the location where the backup archive (p. 409) will be stored. This will create a minimal backup plan consisting of one task (p. 420) that will create a full backup (p. 409) 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. 35); perform archive maintenance operations such as backup validation (p. 421) or deleting outdated backups (archive cleanup (p. 412)). 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. 421). 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 User privileges on a managed machine

Windows

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.

Linux

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

- Back up and recover any data or the entire machine, having full control over all Acronis Backup & Recovery 10 agent operations and log files on the machine.
- Manage local backup plans and tasks owned by any user registered in the operating system.

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

2.3 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.

2.4 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. 409) 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

2.5 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 (p. 153).)

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".



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.




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.



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	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	Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α
2		В				В				В				В				В				В				В				В	
3				С								С								С								С			
4								D																D							
5																Е															

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	Α		Α		Α		Α		Α		Α		Α		Α	
2		В				В				В				В		
3				С								С				
4								D								
5																Е

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)		Α		Α		Α		Α		Α		Α		Α		Α
2 (Differential)			В				В				В				В	
3 (Differential)					С								С			
4 (Differential)									D							
5 (Full)	Е															

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)		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α		Α	
2 (Differential)			В				В				В				В				В				В		
3 (Differential)					С								С								С				
4 (Differential)									D																D
5 (Full)	Е																Е								

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 (p. 160).


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. 225) 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. 135) 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. 416) 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. 415), or more exactly, on a disk group (p. 414). 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	Dynamic volume	Dynamic volume
	Type as of the	Type as of the	Type as of the
	target	target	target
Unallocated space (disk group)	Dynamic volume	Dynamic volume	N/A
Commercial opinion (men. 8. c.a.p.)	Type as of the	Simple	
	source		
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. 287) section.

2.9 Backing up LVM volumes and MD devices (Linux)

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

2.9.1 Backing up logical volumes

Acronis Backup & Recovery 10 Agent for Linux can access, back up and recover such volumes when running in Linux with 2.6.x kernel or a Linux-based bootable media.

Backup (GUI)

In Acronis Backup & Recovery 10 GUI, logical volumes appear under **Dynamic & GPT Volumes** at the end of the list of volumes available for backup.

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

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

Recovery

When recovering logical volumes, you have two options:

 Recovering volume contents only. The type or other properties of the target volume will not change.

This option is available both in the operating system and under bootable media.

This option is useful in the following cases:

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

A system, recovered from a logical volume backup to a basic MBR disk, cannot boot because its kernel tries to mount the root file system at the logical volume. To boot the system, change the loader configuration and /etc/fstab so that LVM is not used and reactivate your boot loader (p. 248).

- When recovering a basic or logical volume to a previously created logical volume. Such is the
 case when you create the structure of logical volumes manually by using the lvm utility.
- Recovering both the structure of logical volumes and their contents.

Such is the case when recovering on bare metal or on a machine with different volume structure. The structure of logical volumes can be automatically created at the time of recovery, if it has been saved in the backup (p. 49).

This option is available only under bootable media.

For detailed instructions on how to recover logical volumes, see Recovering MD devices and logical volumes (p. 281).

Helpful link:

http://tldp.org/HOWTO/LVM-HOWTO/

2.9.2 Backing up MD devices

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

You can back up active (mounted) MD devices in the same way as logical volumes. The MD devices appear at the end of the list of volumes available for backup.

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

When recovering MD devices under bootable media, the structure of MD devices can be automatically created if it has been saved in the backup (p. 49). For detailed information about recovering MD devices under bootable media, see Recovering MD devices and logical volumes (p. 281).

For information about assembling MD devices when performing recovery in Linux, see Assembling MD devices for recovery (Linux) (p. 249).

2.9.3 Saving the volume structure information

For the structure of MD devices and logical volumes to be automatically created at the time of recovery, you need to save the volume structure information in either of these ways:

- When creating a backup plan for the disk-level backup, go to Backup options > Advanced settings, and then select the Save software RAID and LVM metadata along with backups check box. (It is selected by default.)
- Before performing the first disk backup on a source machine, run the following command: trueimagecmd --dumpraidinfo

Either operation saves the machine's logical volume structure to the /etc/Acronis directory. Make sure that the volume with this directory is selected for backup.

2.9.4 Selecting logical volumes and MD devices in command line

Let's assume that the system has four physical disks: Disk 1, Disk 2, Disk 3 and Disk 4.

- A RAID-1 volume is configured on two basic volumes: sdb1, sdd1
- A logical volume is configured on two basic volumes: sdb2, sdd2
- Disk 1 includes Acronis Secure Zone, which normally is not backed up.

A list of a volumes can be obtained with the following command:

trueimagecmd --list

Num	Partition	Flags	Start	Size	Туре
Diek	1 (cda):				
1-1	1 (sda):	Pri,Act	63	208813	Ext2
1-1		Pri,ACC	417690	12289725	
					ReiserFS
1-3	sda3	Pri	24997140	1052257	Linux Swap
	Unallocated 7	D	27101655	2698920	Unallocated
1-4	Acronis Secure Zone	Pri	32499495	522112	FAT32
	Unallocated		33543720	5356	Unallocated
	2 (sdb):			101050	=
	sdb1	Pri - ·	62	124969	Ext2
2-2	sdb2	Pri	250001	125000	None
	Unallocated		500001	8138607	Unallocated
Disk	3 (sdc):				
	Table		0		Table
	Unallocated		1	1048575	Unallocated
Disk	4 (sdd):				
4-1	sdd1	Pri	62	124969	Ext2
4-2	sdd2	Pri	250001	125000	None
	Unallocated		500001	798575	Unallocated
Dynar	nic & GPT Volumes:				
DYN1	VolGroup00-LogVol00			245760	Ext3
		Disk: 3	250385	245760	
		Disk: 5	250385	245760	
DYN2	md0			124864	Ext2
		Disk: 5	62	249728	
		Disk: 3	62	249728	

The logical volume, DYN1, occupies basic volumes 2-2 and 4-2. The RAID-1 volume, DYN2, occupies basic volumes 2-1 and 4-1.

To back up the logical DYN1 volume, run the following command (here, the name of the backup is assumed to be /home/backup.tib):

```
trueimagecmd --partition:dyn1 --filename:/home/backup.tib --create
```

To back up the RAID-1 volume DYN2, run the following command:

```
trueimagecmd --partition:dyn2 --filename:/home/backup.tib --create
```

To back up all three hard disks with volumes, select the volumes 1-1, 1-2, 1-3, DYN1 and DYN2:

```
trueimagecmd --partition:1-1,1-2,1-3,dyn1,dyn2 --filename:/home/backup.tib --
create
```

If you select Disk 3, volume 2-1 or volume 2-2, the program will create a raw (sector-by-sector) backup.

2.10 Backing up hardware RAID arrays (Linux)

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

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

2.11 Backing up virtual machines

Acronis Backup & Recovery 10 Advanced Server Virtual Edition allows for backing up virtual machines from the host.

Preparation

On Windows 2008 Server x64 (any edition) or Microsoft Hyper-V Server 2008:

- Install the Agent for Hyper-V on the Hyper-V host.
- Integration services (p. 53) have to be installed on the guest systems.

On VMware ESX/ESXi:

- Install the Agent for ESX/ESXi on the ESX or ESXi host. The agent is delivered as a virtual appliance.
- VMware Tools (p. 53) have to be installed on the guest systems.

Virtual machines backup

Once the agent is installed on the host and the required services are installed on the guests, you can:

- back up a virtual machine or multiple virtual machines residing on the server without having to install the agent on each virtual machine
- recover a virtual machine to the same, another, or new virtual machine residing on the same server or on another virtualization server where the agent for virtual machines is installed. The virtual machine configuration, stored in a virtual machine backup, will be suggested by default at recovering the backup content to a new virtual machine
- back up and recover individual disks and volumes of a virtual machine.

A virtual machine can be online (running), offline (stopped), suspended, or switch between the three states during backup.

A virtual machine has to be offline (stopped) during the recovery to this machine. The machine will be automatically stopped before recovery. You can opt for manual stopping of machines (p. 133).

Virtual machine backup vs. the machine's volumes backup

Backing up a virtual machine means backing up all the machine's disks plus the machine configuration. With this source type, you can back up multiple machines. This comes in handy when having small (in terms of virtual disk size) but numerous legacy servers such as those resulting from workload consolidation. A separate archive will be created for each machine.

Backing up volumes within a virtual machine is similar to backing up a physical machine's volumes. With this source type, you select the machine and then select the disks/volumes to back up. This comes in handy when the operating system and applications, such as a database server, run on a virtual disk, but the data, such as a database, is stored on a large capacity physical disk added to the same machine. You will be able to use different backup strategies for the virtual disk and the physical storage. The virtual machine configuration will be also backed up.

Limitations

A Hyper-V virtual machine that uses at least one pass-through disk (a physical disk, either local or SAN-LUN, attached to the virtual machine) cannot be backed up from the host. To back up such machine or its disks, install Agent for Windows or Agent for Linux on the machine.

An online (running) ESX/ESXi virtual machine that has an independent disk or an RDM disk attached in the physical compatibility mode cannot be backed up from the host. To back up such machine or its disks, either stop the machine or install Agent for Windows or Agent for Linux on the machine.

Virtual machine backup vs. physical machine backup

Backing up an entire virtual machine or its volumes yields a standard disk backup (p. 414). With Acronis Backup & Recovery 10 Agent for Windows or Acronis Backup & Recovery 10 Agent for Linux, you can mount its volumes, recover individual files from this backup, and recover disks and volumes from the backup to a physical machine.

Similarly, you can recover disks or volumes from a physical machine backup created with the Agent for Windows or the Agent for Linux, to a new or existing virtual machine using either of the agents for virtual machines. Hence, physical to virtual and virtual to physical machine migration becomes available.

Guest operating systems

The following guest operating systems are supported.

Microsoft Windows platform:

- Microsoft Windows 2000
- Microsoft Windows XP
- Microsoft Windows Server 2003
- Microsoft Windows Server 2003 R2
- Microsoft Vista
- Microsoft Windows Server 2008
- Microsoft Windows Server 2008 R2
- Microsoft Windows 7

Linux platform.

Guest HDD

The following virtual disk configurations are supported.

Partitioning style: MBR

Volume types: basic and dynamic volumes.

Dynamic volumes (LDM in Windows and LVM in Linux) are supported to the same extent as on physical machines. The LDM/LVM structure needs to be created prior to the recovery if you want to retain the LDM/LVM. To do so, you will have to boot the target virtual machine using bootable media (p. 411) or its ISO image and use Acronis Disk Director Lite for LDM reconstruction or Linux command line tools for LVM reconstruction. Another option is to recover dynamic volumes as basic.

Troubleshooting

Agent: Agent for Hyper-V

Issue: Backup of an online virtual machine fails because of a Volume Shadow Copy Service (VSS)

error. The error can be seen in the Application Event Log (Event ID = 8193).

Cause: This happens because there is no registry key:

Solution: Add this key to the registry. To do so, create and run the following script (xxx.reg):

```
[HKEY_CLASSES_ROOT\Wow6432Node\CLSID\{F2C2787D-95AB-40D4-942D-298F5F757874}]
@="PSFactoryBuffer"
[HKEY_CLASSES_ROOT\Wow6432Node\CLSID\{F2C2787D-95AB-40D4-942D-298F5F757874}\InProcServer32]
@=hex(2):25,00,73,00,79,00,73,00,74,00,65,00,6d,00,72,00,6f,00,6f,00,74,00,25,\
00,5c,00,53,00,79,00,73,00,57,00,4f,00,57,00,36,00,34,00,5c,00,76,00,73,00,\
73,00,5f,00,70,00,73,00,2e,00,64,00,6c,00,6c,00,00,00
"ThreadingModel"="Both"
```

2.11.1 How to install Hyper-V Integration Services

To install the Hyper-V Integration Services:

- 1. Run the guest operating system.
- 2. Select Action > Insert Integration Services Setup Disk.
- 3. The server connects the ISO image of the setup disk to the machine. Follow the onscreen instructions.

2.11.2 How to install VMware Tools

To install the VMware Tools:

- 1. Run the VMware Infrastructure/vSphere Client.
- 2. Connect to the ESX server.
- 3. Select the virtual machine and run the guest operating system.
- 4. Right click the virtual machine and select **Guest > Install/Upgrade VMware Tools**.
- 5. Follow the onscreen instructions.

2.12 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 (p. 21). Storage nodes ensure fully automatic operation of tape libraries and autoloaders (p. 143).

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.

2.12.1 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 dev	vice attached	to a machine
			ABR10	ABR10	ABR10	ABR10
			Bootable	Agent for	Agent for	Storage
			Media	Windows	Linux	Node
Tape written on	Bootable Media	ATIE 9.1	+	+	+	+
a locally		ATIE 9.5	+	+	+	+
attached tape		ATIE 9.7	+	+	+	+
device (tape		ABR10	+	+	+	+
drive or tape	Agent for	ATIE 9.1	+	+	+	+
library) by	Windows	ATIE 9.5	-	-	-	+
		ATIE 9.7	-	-	-	+
		ABR10	+	+	+	+
	Agent for Linux	ATIE 9.1	+	+	+	+
		ATIE 9.5	+	+	+	+
		ATIE 9.7	+	+	+	+
		ABR10	+	+	+	+
Tape written on	Backup Server	ATIE 9.1	+	+	+	+
a tape device		ATIE 9.5	-	_	-	+
through		ATIE 9.7	-	-	-	+
	Storage Node	ABR10	-	-	-	+

2.12.2 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. 123).

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. 420) is proposed at every disk or volume recovery, even when recovering a Linux or non-system Windows volume.
- 6. Acronis Active Restore (p. 408) 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.

2.13 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"

2.14 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.

2.14.1 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. 120) 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. 421). Once created on a managed machine, the zone is always present in the list of **Personal vaults**. Centralized backup plans (p. 412) can use Acronis Secure Zone as well as local plans (p. 417).

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. 166) 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.

2.14.2 Acronis Startup Recovery Manager

A modification of the bootable agent (p. 411) 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.

2.14.3 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.

2.14.4 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

- 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.

2.15 Understanding centralized management

This section contains an overview of centralized data protection with Acronis Backup & Recovery 10. Please be sure you have an understanding of how data is protected on a single machine (p. 29) before reading this section.

2.15.1 Basic concepts

Applying backup policies and tracking their execution

To protect data on a single machine, you install on the machine an agent (p. 409) or multiple agents for various data types you want to protect. You connect the console to the machine and create a backup plan (p. 410) or multiple backup plans.

What if you have to manage hundreds of machines? It takes time to create a backup plan on each machine, while the plans may be quite similar – you need to back up, say, the system drive and the users' documents. Tracking the plans' execution on each machine separately is also time-consuming.

To be able to propagate the management operations to multiple machines, you install Acronis Backup & Recovery 10 Management Server (p. 418) and register (p. 419) the machines on the server. After that you can create groups of machines and thus manage multiple machines as a whole. You can protect all of them or your selection by setting up a common backup plan, which is called a backup policy (p. 410).

Once you apply the policy to a group of machines, the management server deploys the policy to each of the machines. On each machine the agents find the items to back up and create corresponding centralized backup plans (p. 412). You will be able to monitor the policies' statuses on a single screen and navigate, if required, to each machine, plan or task to see their status and log entries. The management server also enables you to monitor and manage the agent's locally originated activities.

Since you connect the console to the management server rather than to each machine and perform all management operations through the central management unit, this way of management is called centralized management (p. 412).

Centralized management does not rule out the direct management (p. 413) of each machine. You can connect the console to each machine and perform any direct management operation. However, centralized backup plans can be managed through the management server only, since a well-thought out policy functions automatically and rarely requires human intervention.

Using the management server, you can create one or more centralized archive storages (centralized vaults (p. 412)), which will be shared by the registered machines. A centralized vault can be used by any backup policy as well as by any backup plan created on the registered machines using direct management.

Organizing a managed archive storage

What should the capacity of your centralized vault be? What if transferring sizeable backups to the vault will cause network congestion? Does backup of an online production server affect the server performance? To ensure that the centralized backup will not slow down business processes in your company and to minimize the resources required for the data protection, you install Acronis Backup & Recovery 10 Storage Node (p. 419) and configure it to manage a centralized vault or multiple centralized vaults. Such vaults are called managed vaults (p. 417).

The storage node helps the agent deduplicate (p. 413) backups before transferring them to managed vaults and deduplicates the backups already saved in the vaults. Deduplication results in reducing backup traffic and saving storage space. The storage node also undertakes operations with archives (such as validation and cleanup), which otherwise are performed by the agent, and thus relieves the managed machines from unnecessary computing load. Last but not least, Acronis Backup & Recovery 10 Storage Node enables using a tape library as a centralized vault for storing backup archives.


More than one storage node, each managing a number of vaults, can be set up and controlled centrally from the Acronis Backup & Recovery 10 Management Server.

For more detailed information about storage nodes please refer to Acronis Backup & Recovery 10 Storage Node (p. 21).

2.15.2 Setting up centralized data protection in a heterogeneous network

Assume that the network infrastructure includes servers (1, 2, 9) and workstations (3, 5-8) running Windows and Linux. You also have a VMware ESX server (4) that hosts two guest systems.

You have to protect each server as a whole, the users' data on the workstations, and the virtual machines. You want to be able to track the health of the data protection, be sure that the backup archives do not store duplicated information and that the obsolete backups are deleted from the storage in a timely manner. These goals can be achieved by regular backup of the desired data items to a centralized vault with deduplication.



Setting up the Acronis infrastructure

- 1. Install Acronis Backup & Recovery 10 Management Console [Console] on the machine which you prefer to operate from (3). The console enables you to access and manage other Acronis components through Graphical User Interface.
- 2. Install Acronis Backup & Recovery 10 Management Server [AMS] on one of the Windows servers (2). The management server is your single entry point to the Acronis infrastructure.
- 3. Install Acronis Backup & Recovery 10 Agent on each of the machines to back up the machine's disks, volumes or files.
 - Agent (W) Agent for Windows
 - Agent (L) Agent for Linux.

When installing the agents, register each of the machines on the management server. To do so, enter the server's name or IP address and the server's administrator credentials in the appropriate window of the installation wizard. Or, alternatively, add the machines to the management server later using their names or IP addresses.

- 4. Install Acronis Backup & Recovery 10 Agent for ESX/ESXi [Agent (ESX)] on the ESX server (4) to back up the virtual machines from the host. The agent is delivered as a virtual appliance.
- 5. Install Acronis Backup & Recovery 10 Storage Node [ASN] on one of the Windows servers (9). The storage node enables you to organize the infrastructure for storing backup archives and to use the deduplication functionality. The node can be installed together with the management server if the host is capable enough.

When installing the storage node, register it on the management server in the same way as you register the agents.

Installation tips

- Both AMS and ASN can be installed on a workstation operating system as well.
- There can be multiple storage nodes on the network. Each of the nodes can manage up to 20 local or remote vaults.
- Multiple Acronis Backup & Recovery 10 components can be installed on a machine with a single installation procedure.
- In an Active Directory domain, you can deploy the components using the Group Policy.

Setting up the storage node

Before using the storage node, make sure that all users that will back up to the node's vaults have Windows accounts on the node.

- If the node is included in an Active Directory domain, all the domain users will be able to back up to the node; and all the domain administrators will become node administrators.
- In a workgroup, create a local user account for each user that will back up to the node. Members of the Administrators group become node administrators. You can add more accounts later as required.
- 1. Run the console, connect to the management server.
- 2. Create a managed vault as described in Operations with centralized vaults (p. 138). Enable deduplication when creating a managed vault.

Setting up groups and policies

The detailed explanation of when and why you need to organize groups of machines can be found in the Grouping the registered machines (p. 66) section. Here are some scenarios supported by the aforementioned Acronis Backup & Recovery 10 implementation.

2.15.2.1 Protecting the servers

You will most likely create individual backup plans on each of the servers depending on their roles. But it is necessary to perform a full backup of the entire server at least once. You might want to back up the server during a maintenance window or backup window, after installing or updating software, before relocation, etc. In our example, there is no need to back up entire servers on a regular basis. You can manually delete old backups since they are not numerous.

- 1. Create a policy that backs up [All Volumes] to the managed vault on the storage node. Choose **Back up later,** manual start and **Full** backup type.
- 2. Create a static group named, say, S_1. Add all the servers to this group. (A storage node can be added in case the managed vault is not on the local node's drives. Otherwise the archive storage will be backed up to itself).
- 3. Apply the policy to the S_1 group. Make sure that the policy has been successfully deployed to each of the servers. The policy deployment state has to change from **Deploying** to **Deployed** and its status has to be **OK**. To see the resulting backup plans on each of the servers:
 - a. navigate to the All machines group or the S 1 group
 - b. select the server
 - c. select the **Backup plans and tasks** tab on the **Information** pane.

When you need and have the opportunity to back up any of the servers, navigate to the backup plan as described above, select the plan and run it.

2.15.2.2 Protecting the workstations

Here is how to set up the most popular schedule: weekly full backup and daily incremental backup of users' default document folders. In addition, we will retain only backups from the last 7 days.

- 1. Create a policy that backs up [All Profiles Folder] to the managed vault on the storage node. This will back up the folder where user profiles are located (for example, C:\Documents and Settings in Windows XP). Choose the **Custom** backup scheme.
 - a. Schedule full backup as follows: **Weekly**, Every 1 week on: Sunday, Execute the task once at 12:00:00 AM. Advanced settings: Wake-on-LAN: On. You may also want to distribute the backup start time within the time window to optimize the network usage and the storage node CPU load.
 - b. Schedule incremental backup as follows: **Weekly**, Every 1 week on: Workdays, Execute the task once at 08:00:00 PM. Also set the advanced settings as required.
 - c. Set up the retention rules as follows: **Delete backups older than**: 7 days. **When deleting a backup that has dependencies:** Consolidate the backups. Leave the default settings for the remaining retention rules. In **Apply retention rules**, set **After backup**.
- 2. Create a dynamic group named, say, W_1. Specify **%Windows%XP%** and **%Windows%Vista%** as the criteria. This way, any workstation that will be registered on the management server later, will be added to this group and protected by the same policy.
- 3. Apply the policy to the W_1 group. Make sure that the policy has been successfully deployed to each of the workstations. The policy deployment state has to change from **Deploying** to **Deployed** and its status has to be **OK**. To see the resulting backup plans on each of the workstations:
 - a. navigate to the All machines group or the W_1 group
 - b. select the workstation
 - c. select the **Backup plans and tasks** tab on the **Information** pane.

You can also see the resulting tasks, created on the workstations, in the Tasks view.

4. Use the **Dashboard** or the **Tasks** view to track the daily activities related to the policy. Once you ascertain that all tasks run as specified, you can only check the policy status in the **Backup policies** view.

To protect data on a daily basis, you can also use the GFS or Tower of Hanoi backup schemes.

2.15.2.3 Protecting the virtual machines

Acronis Backup & Recovery 10 Agent for ESX/ESXi provides the flexibility to protect virtual machines in multiple ways:

- Connect the console to the virtual appliance (Agent for ESX/ESXi) and create a backup plan that will back up all or some of the virtual machines.
- Connect the console to the virtual appliance (Agent for ESX/ESXi) and create an individual backup plan for each machine. The plan will back up the volumes you specify.
- Register the virtual appliance (Agent for ESX/ESXi) on the management server. All virtual machines, except for the virtual appliance, will appear in the All virtual machines group. You can group these machines and apply any policy that backs up disks or volumes to them.

• Install Agent for Windows or Agent for Linux on each virtual machine. Register the machines on the management server. The machines will be considered as physical machines. You can apply a backup policy to these machines or create a backup plan on each machine separately. If any of the machines meets membership criteria set for a dynamic group of physical machines, the machine will be protected by the policy applied to this group.

Advanced product editions other than Virtual Edition (Acronis Backup & Recovery 10 Advanced Server, Advanced Server SBS Edition and Advanced Workstation) allow using only the last of the above methods.

2.15.3 Grouping the registered machines

As soon as a machine is registered (p. 419) on the management server, the machine appears in the **All machines** built-in group (p. 411). By applying a backup policy to this group, you protect all the registered machines. The thing is that a single policy may not be satisfactory because of the different roles of the machines. The backed up data is specific for each department; some data has to be backed up frequently, other - twice a year; so you may want to create various policies applicable to different sets of machines. In this case consider creating custom groups.

2.15.3.1 Static and dynamic groups

You can explicitly specify which machines the custom group has to include, for example, let's say, select each of the accountants' machines. Once you apply the accounting department policy to the group, the accountants' machines become protected. If a new accountant is hired, you will have to add the new machine to the group manually. Such groups are called static (p. 419) since their content never changes unless the administrator explicitly adds or deletes a machine.

Manual operation is not required though, if the accounting department forms a separate Active Directory organization unit. You specify the accounting OU as the group membership criterion. If a new accountant is hired, the new machine will be added to the group as soon as it is added to the OU, and thus will be protected automatically. Such groups are called dynamic (p. 415) since their content changes automatically.

2.15.3.2 Dynamic grouping criteria

Acronis Backup & Recovery 10 Management Server offers the following dynamic membership criteria:

- Operating system (OS)
- Active Directory organization unit (OU)
- IP address range.

Multiple criteria can be specified for a dynamic group. For example, a set of criteria "OS equals Windows 2000, OS equals Windows 2003, OU equals Accounting" is interpreted as "all machines running Windows 2000 or Windows 2003 and belonging to the Accounting organizational unit".

The **All machines** group can be thought of as a dynamic group with the single built-in criterion: include all the registered machines.

2.15.3.3 Using custom groups

Grouping helps the administrator to organize data protection by company departments, by Active Directory organizational units, by various populations of users, by the site locations and the like. To

make the best use of the AD OU criterion, consider reproducing the Active Directory hierarchy in the management server. Grouping by the IP address range enables taking account of the network topology.

The groups you create can be nested. The management server is capable of maintaining up to 500 groups in total. A machine can be a member of more than one group.


Besides physical machines, you can group virtual machines (p. 328) hosted on registered virtualization servers. Virtual machines have their own grouping criteria depending on their properties.

2.15.3.4 Example

The diagram below presents an example of group hierarchy.

Six machines are registered on the management server:

- 1 the international sales manager's laptop (Windows Vista)
- 2 the server that holds the corporate database and the shared document storage (Windows Server 2008)
- 3, 4, 5, 6 the salesmen's machines (Windows XP) from the "Sales department" AD organization unit.



An example of group hierarchy

The backup policy on the server has to differ from that on the workstations. The administrator creates the G1 dynamic group that contains machines with the server operating systems, and applies a backup policy to the group. Any server, that is added to the network and registered on the management server, will appear in this group and the policy will be applied to it automatically.

To protect the salesmen's workstations with a different policy, the administrator creates the G2 dynamic group using the AD OU criterion. Any change in the OU membership of a machine will be reflected in the G2 membership. The appropriate policy will be applied to the new OU members and revoked from machines deleted from the OU.

The international sales manager's laptop is not included in the OU but it has some of the data the sales machines have. To back up this data, the administrator has to add the laptop to G2 "by force". This can be done by creating a static group (G3) and moving the static group into the dynamic one. The policy applied to the parent group (G2) will be applied to the child group (G3), but members of G3 are not considered as members of G2 and so its dynamic nature is considered intact.

In real life, the administrator would most likely prefer to protect the manager's machine by applying the policy directly to the machine, without including it in any group, so this case is just an illustration of nesting different types of groups. With multiple group members, nesting of the groups comes in handy.

2.15.3.5 Operations with custom groups

You create empty groups in the generic root (Physical machines or Virtual machines) or within existing groups and populate them by manual adding machines (static groups) or by adding criteria of dynamic group membership. You can also

- edit a group, that is:
 - change the group name
 - change the group description
 - change the dynamic membership criteria
- transform a static group into a dynamic one by adding membership criteria
- transform a dynamic group into a static one with two options:
 - keep the group members
 - remove the group members
- move a group from the root to another group (any group type to any group type)
- move a group from the parent group to the root
- move a group from one parent group to another (any group type to any group type)
- delete a group, that is, disjoin the group members that remain in the group of all machines anyway.

Operations with groups to which backup policies are applied will result in changing the policies on the member machines. If a machine is not available or reachable at the moment, the action becomes pending and will be performed as soon as the machine becomes available.

For information on how to perform the operations please see Operations with groups (p. 323).

2.15.4 Policies on machines and groups

This section helps you understand the automatic deployment and revoking policies performed by the management server when a policy or a number of policies are applied to machines and nested groups of machines in various combinations; when a policy is revoked from machines and groups; when a machine or a group is moved from one group to another.

Operations with groups to which backup policies are applied will result in changing the policies on the member machines. On any hierarchy change, that is, when moving, removing, creating groups; adding machines to static groups; or when machines enter a group based on dynamic criteria, a huge number of inheritance changes may occur. Please familiarize yourself with this section to be sure that your actions yield the desired result and to understand the result of the automated Acronis Backup & Recovery 10 Management Server operations.

What is applying, deploying and revoking?

Applying a policy establishes the correspondence between the policy and one or more machines. This process takes place inside the management server's database and does not take much time.

Deploying a policy transfers the established correspondence to the machines. Physically, a bundle of tasks is created on each machine according to the configuration provided by the policy.

Revoking a policy is the reverse action to the aggregate of applying and deploying. Revoking removes the correspondence between the policy and one or more machines and then removes the tasks from the machines.


If a machine is not available or not reachable at the moment, the change will be propagated on the machine when it becomes available. This means that deploying a policy to multiple machines is not a momentary action. The same is true for revoking. These two processes may be durable and so the management server tracks and displays personal statuses for each machine that it works with, as well as the policy's cumulative status.

2.15.4.1 A policy on a machine or a group

In the diagrams below, each numbered scheme illustrates the result of the respectively numbered action.


The container stands for a group; the colored circle stands for a machine with applied policy; the dark colored circle stands for a machine with two applications of the same policy; the white circle stands for a machine to which no policy is applied.


Policy on a machine



- 1. A policy can be applied to a machine.
- 2. A policy can be revoked from a machine.

Policy on a group






1. A policy can be applied to a group.

- 2. A policy can be revoked from a group.
- 3. A policy applied to a group cannot be revoked from a machine.
- 4. To revoke the policy from the machine, remove the machine from the group.

The same policy on a group and on a machine




- 1. The same policy can be applied to a group and to a machine. Nothing changes on the machine at the second application of the same policy, but the server remembers that the policy has been applied twice.
- 2. A policy, revoked from the group, remains on the machine.
- 3. A policy, revoked from the machine, remains on the group and therefore on the machine.
- 4. To completely revoke the policy from the machine, revoke it from both the group and the machine.


2.15.4.2 Operations with a machine

This section is a simplified illustration of what happens with the policies on a machine when the machine is moved, copied, or deleted from a group.

In the diagram below, the container stands for a group; the one-color circle stands for a machine with one applied policy; the two-color circle stands for a machine with two applied policies; the white circle stands for a machine with no policy applied.

- 1. Here is the initial state: two custom groups contain different machines. A policy is applied to one group; another policy is applied to another group. The next schemes illustrate results of the specified actions.
- 2. **Move to another group:** Machine #3 is moved from one group to another. The "orange" policy is revoked; the "blue" policy is applied to the machine.
- 3. **Add to another group**: Machine #3 is added to another group. It becomes a member of both groups. The "blue" policy is applied, but the "orange" policy remains on the machine.
- 4. **Remove from the group**: Machine #3 is removed from the group. The "orange" policy is revoked from the machine. The machine remains in the **All machines** group.






2.15.4.3 Inheritance of policies

Policy inheritance can be easily understood if we assume that a machine can be a member of only one group besides the **All machines** group. Let's start from this simplified approach.

In the diagram below, the container stands for a group; the two-color circle stands for a machine with two applied policies; the three-color circle stands for a machine with three applied policies and so on.



Besides the **All machines** group, we have the custom G1 group in the root and the custom G2 group, which is G1's child.

The "green" policy, applied to the **All machines** group, is inherited by all machines.

The "orange" policy, applied to G1, is inherited by the G1 members and all its child groups, both immediate and indirect.

The "blue" policy, applied to G2, is inherited only by the G2 members since G2 does not have child groups.

The "violet" policy is applied straight to machine #4. It will exist on machine #4 irrespectively of this machine's membership in any group.

Let's assume we create the G3 group in the root. If no policies are applied to the group, all its members are supposed to be "green". But if we add, say, the #1 machine to G3, the machine will bear both "orange" and "green" policies, in spite of the fact that G3 has nothing to do with the "orange" policy.

That's why it is difficult to track the policies' inheritance from the top of the hierarchy if the same machine is included in multiple groups.

In real life, it's much easier to view the inheritance from the machine's side. To do so, navigate to any group that contains the machine, select the machine and then select the **Backup policies** tab on the **Information** pane. The **Inheritance** column shows whether a policy is inherited or applied directly to the machine. Click **Explore inheritance** to view the inheritance order of the policy. In our example, the policy names, the **Inheritance** column and the inheritance order will be as follows:

For machine	Name of the policy	Inheritance	Inheritance order
#1 or #2 or #3	"green"	Inherited	All machines -> #1 or #2 or #3
	"orange"	Inherited	G1 -> #1 or #2 or #3
#4	"green"	Inherited	All machines -> #4
	"orange"	Inherited	G1 -> G2 -> #4
	"blue"	Inherited	G2 -> #4
	"violet"	Applied directly	

#5 or #6 "green" Inherited All machines -> #5 or #6

"orange" Inherited G1 -> G2 -> #5 or #6

"blue" Inherited G2 -> #5 or #6

2.15.5 Backup policy's state and statuses

Centralized management presumes that the administrator can monitor the health of the entire product infrastructure using a few easily understandable parameters. The state and status of a backup policy are included in such parameters. Issues, if any, arise from the very bottom of the infrastructure (tasks on managed machines) to the cumulative policy status. The administrator checks the status at a glance. If the status is not OK, the administrator can navigate down to the issue details in a few clicks.

This section helps you understand the policies' states and statuses displayed by the management server.

2.15.5.1 Policy deployment state on a machine

To see this parameter, select any group, containing the machine, in the tree, then select the machine, and then select the **Backup policies** tab on the **Information** pane.

Once you apply a policy to a machine or a group of machines, the server deploys the policy to the machines. On each of the machines, the agent creates a backup plan. While the policy is transferred to the machine and the backup plan is being created, the policy's deployment state on the machine is **Deploying**.

Once the backup plan is successfully created, the policy state on the machine becomes **Deployed**.


You may need to modify the policy for some reason. Once you confirm the changes, the management server updates the policy on all machines the policy was deployed to. While the changes are transferred to the machine and the agent updates the backup plan, the policy state on the machine is **Updating**. Once the policy is updated, its state becomes **Deployed** again. This state means that the policy is functioning and no changes are currently being made to it.

A policy that was modified while being deployed remains in the **Deploying** state. The management server just starts to deploy the modified policy from the beginning.

You may need to revoke the policy from the machine or from the group the machine is included in. Once you confirm the changes, the management server revokes the policy from the machine. While the changes are transferred to the machine and the agent deletes the backup plan from it, the policy state on the machine is **Revoking**.

You may change grouping conditions or the machine may change its properties so that the machine leaves one group and is included into another. This may result in revoking one policy and deploying another policy. In this case, the first policy's state on the machine will be **Revoking** and the second policy's state will be **Deploying**. The policies can appear in the GUI simultaneously or one after another.

Backup policy state diagram



2.15.5.2 Policy status on a machine

To see this parameter, select any group of machines in the tree, then select the machine, and then select the **Backup policies** tab on the **Information** pane.

In each of the states, the backup policy can have one of the following statuses: **Error**; **Warning**; **OK**. While the policy is in the **Deployed** state, its status reflects how successfully the policy is executed. While the policy is in any other state, its status reflects how successfully the policy is being modified.

Policy status when data to back up is not found on a machine

A backup policy can be applied to a machine that does not have data meeting the selection rules (p. 419). No error or warning will be logged during the policy deployment because it is assumed that the data may appear in the future. A backup plan is created as usual and the policy state is changed to **Deployed**.

If no data to back up is found at the time when the backup task starts, the task will fail and the policy status will turn to **Error**. If at least one of the data items is found, the backup task will succeed with a warning. The policy status will change accordingly.

The backup tasks will start on schedule as specified by the policy and produce a similar result until all data items appear on the machine or the policy is edited to exclude the non-existent data items.

Examples

Assume, the selection rule states that the policy has to back up volumes D: and F:. The policy is applied to both Linux and Windows machines. Once the first backup is started, the policy gets the **Error** status on the Linux machines and on the Windows machines that do not have such volumes. The policy gets the **Warning** status on Windows machines that have either a D: or F: volume, unless an event that will result in an error occurs.

The policy that has to back up the [System] and the /dev/sda1 volumes, will get the **Warning** status on the Windows machines (since /dev/sda is not found) and on the Linux machines that have the /dev/sda1 volume (since the [System] volume is not found). The policy will get the **Error** status on Linux machines that do not have a SCSI device.

The following table provides details.

State	Status	Description			
Deploying Error Warning		The deployment log has errors, for example, disk space runs out			
		The deployment log has warnings: the machine went offline during the deployment; cannot connect for N days			
	ОК	The deployment log does not have errors and warnings			
Deployed	Error	The status of the corresponding backup plan is Error			
Warning		The status of the corresponding backup plan is Warning			
	ок	The status of the corresponding backup plan is OK			
Updating Error		The updating log has errors: cannot delete the locked task, the Acronis service is stopped			
	Warning	The updating log has warnings			
		The updating log does not have errors and warnings			
		The revoking log has errors			
	Warning	The revoking log has warnings			
OK The revoking log does not have errors and warnings					

In addition to the deployment state and status as related to a specific machine, the backup policy has the deployment state and status on a group of machines and the cumulative deployment state and status of the policy.

2.15.5.3 Policy deployment state on a group

To see this parameter, select **Machines** in the tree, then select the group, and then select the **Backup policies** tab on the **Information** pane.

This state is defined as a combination of deployment states of the policy on the machines included in the group and its child groups.

For example, you applied the policy to the group consisting of machines A and B. While the deployment takes place on both machines, the policy's state on the group will be "Deploying". If the deployment completes on one of the machines while it continues on the other, the state will be "Deploying, Deployed". When the deployment completes on both machines, the state will be "Deployed".

2.15.5.4 Policy status on a group

To see this parameter, select **Machines** in the tree, then select the group, and then select the **Backup policies** tab on the **Information** pane.

This status is defined as the most severe status of the policy on the machines included in the group and its child groups. If the policy is currently not applied to any machine, its status is "OK".

2.15.5.5 Cumulative state and status of a policy

In addition to the deployment state and status as related to a specific machine or group, the backup policy has a cumulative deployment state and a cumulative status.

The cumulative state of a backup policy

To see this parameter, select **Backup policies** in the tree. The **Deployment state** column displays the cumulative deployment state for each policy.

This state is defined as a combination of deployment states of the policy on all machines the policy is applied to (directly or through inheritance). If the policy is currently not applied to any machine, it does not have a deployment state and the column shows "Not applied".

For example, you applied the policy to machine A. The policy was successfully deployed. Then you modify the policy and immediately apply it to the group consisting of machines B and C. The policy has to be updated on A and deployed to B and C. While the processes take place, the policy's cumulative state may look like "Updating, Deploying", then change to "Updating, Deployed" or "Deployed, Deploying" and will normally end up with "Deployed".

The cumulative status of a backup policy

To see this parameter, select **Backup policies** in the tree. The **Status** column displays the cumulative status for each policy.

This status is defined as the most severe status of the policy on all machines the policy is applied to. If the policy is not applied to any machine, its status is "OK".

2.15.6 Deduplication

This section describes deduplication, a mechanism designed to eliminate data repetition by storing identical data in archives only once.

2.15.6.1 Overview

Deduplication is the process of minimizing storage space taken by the data by detecting data repetition and storing the identical data only once.

For example, if a managed vault where deduplication is enabled contains two copies of the same file—whether in the same archive or in different archives—the file is stored only once, and a link to that file is stored instead of the second file.

Deduplication may also reduce network load: if, during a backup, a file or a disk block is found to be a duplicate of an already stored one, its content is not transferred over the network.

Deduplication is performed on disk blocks (block-level deduplication) and on files (file-level deduplication), for disk-level and file-level backups respectively.

In Acronis Backup & Recovery 10, deduplication consists of two steps:

Deduplication at source

Performed on a managed machine during backup. Acronis Backup & Recovery 10 Agent uses the storage node to determine what data can be deduplicated, and does not transfer the data whose duplicates are already present in the vault.

Deduplication at target

Performed in the vault after a backup is completed. The storage node analyses the vault's archives and deduplicates data in the vault.

When creating a backup plan, you have the option to turn off deduplication at source for that plan. This may lead to faster backups but a greater load on the network and storage node.

Deduplicating vault

A managed centralized vault where deduplication is enabled is called a *deduplicating vault*. When you create a managed centralized vault, you can specify whether to enable deduplication in it. A deduplicating vault cannot be created on a tape device.

Deduplication database

Acronis Backup & Recovery 10 Storage Node managing a deduplicating vault, maintains the deduplication database, which contains the hash values of all items stored in the vault—except for those that cannot be deduplicated, such as encrypted files.

The deduplication database is stored in the folder which is specified by the **Database path** in the **Create centralized vault** view when creating the vault. Deduplication database can be created in a local folder only.

The size of the deduplication database is about one percent of the total size of archives in the vault. In other words, each terabyte of new (non-duplicate) data adds about 10 GB to the database.

In case the database is corrupted or the storage node is lost, while the vault retains archives and the service folder containing metadata, the new storage node rescans the vault and re-creates the database.

2.15.6.2 How deduplication works

Deduplication at source

When performing a backup to a deduplicating vault, Acronis Backup & Recovery 10 Agent reads items being backed up—disk blocks for disk backup or files for file backup—and calculates a fingerprint of each block. Such a fingerprint, often called a *hash value*, uniquely represents the item's content within the vault.

Before sending the item to the vault, the agent queries the deduplication database to determine whether the item's hash value is the same as that of an already stored item.

If so, the agent sends only the item's hash value; otherwise, it sends the item itself.

Some items, such as encrypted files or disk blocks of a non-standard size, cannot be deduplicated, and the agent always transfers such items to the vault without calculating their hash values. For more information about restrictions of file-level and disk-level deduplication, see Deduplication restrictions (p. 80).

Deduplication at target

After backup to a deduplicating vault is completed, the storage node runs the **indexing task** to deduplicate data in the vault as follows:

1. It moves the items (disk blocks or files) from the archives to a special file within the vault, storing duplicate items there only once. This file is called the **deduplication data store**. If there are both

disk-level and file-level backups in the vault, there are two separate data stores for them. Items that cannot be deduplicated remain in the archives.

2. In the archives, it replaces the moved items with the corresponding references to them.

As a result, the vault contains a number of unique, deduplicated items, with each item having one or more references to it from the vault's archives.

The indexing task may take considerable time to complete. You can see this task's state in the **Tasks** view on the management server.

Compacting

After one or more backups or archives have been deleted from the vault—either manually or during cleanup—the vault may contain items which are no longer referred to from any archive. Such items are deleted by the **compacting task**, which is a scheduled task performed by the storage node.

By default, the compacting task runs every Sunday night at 03:00. You can re-schedule the task as described in Actions on storage nodes (p. 337), under "Change the compacting task schedule". You can also manually start or stop the task from the **Tasks** view.

Because deletion of unused items is resource-consuming, the compacting task performs it only when a sufficient amount of data to delete has accumulated. The threshold is determined by the **Compacting Trigger Threshold** (p. 352) configuration parameter.

2.15.6.3 When deduplication is most effective

The following are cases when deduplication produces the maximum effect:

- When backing up in the full backup mode similar data from different sources. Such is the case when you back up operating systems and applications deployed from a single source over the network.
- When performing incremental backups of similar data from different sources, provided that the changes to the data are also similar. Such is the case when you deploy updates to these systems and apply the incremental backup.
- When performing incremental backups of data that does not change itself, but changes its location. Such is the case when multiple pieces of data circulate over the network or within one system. Each time a piece of data moves, it is included in the incremental backup which becomes sizeable while it does not contain new data. Deduplication helps to solve the problem: each time an item appears in a new place, a reference to the item is saved instead of the item itself.

Deduplication and incremental backups

In case of random changes to the data, de-duplication at incremental backup will not produce much effect because:

- The deduplicated items that have not changed are not included in the incremental backup.
- The deduplicated items that have changed are not identical anymore and therefore will not be deduplicated.

2.15.6.4 Deduplication best practices

Follow these recommendations when using deduplication:

- When creating a deduplicating vault, place the vault and its deduplication database on different disks. This will make deduplication faster, because deduplication involves extensive simultaneous use of both the vault and the database.
- Indexing of a backup requires that the vault have free space with a minimum size of 1.1 multiplied by the size of the archive the backup belongs to. If there is not enough free space in the vault, the indexing task will fail and start again after 5–10 minutes, on the assumption that some space has been freed up as a result of cleanup or of other indexing tasks. The more free space there is in the vault, the faster your archives will reduce to the minimum possible size.
- When backing up multiple systems with similar content, back up one of the similar systems first, so that Acronis Backup & Recovery 10 Storage Node indexes all the system's files as potential deduplication items. This will lead to faster backup processes and less network traffic (because of effective deduplication at source), regardless of whether the backups are performed simultaneously or not.
 - Before starting the subsequent backups, make sure that the **indexing task has finished** deduplication of the first backup and is now idle. You can view the state of the indexing task in the list of tasks on Acronis Backup & Recovery 10 Management Server.

2.15.6.5 Deduplication ratio

The deduplication ratio shows the size of archives in a deduplicating vault in relation to the size they would occupy in a non-deduplicating vault.

For example, suppose that you are backing up two files with identical content from two machines. If the size of each file is one gigabyte, then the size of the backups in a non-deduplicating vault will be approximately 2 GB, but this size will be just about 1 GB in a deduplicating vault. This gives a deduplication ratio of 2:1, or 50%.

Conversely, if the two files had different content, the backup sizes in non-deduplicating and duplicating vaults would be the same (2 GB), and the deduplication ratio would be 1:1, or 100%.

What ratio to expect

Although, in some situations, the deduplication ratio may be very high (in the previous example, increasing the number of machines would lead to ratios of 3:1, 4:1, etc.), a reasonable expectation for a typical environment is a ratio between 1.2:1 and 1.6:1.

As a more realistic example, suppose that you are performing a file-level or disk-level backup of two machines with similar disks. On each machine, the files common to all the machines occupy 50% of disk space (say, 1 GB); the files that are specific to each machine occupy the other 50% (another 1 GB).

In a deduplicating vault, the size of the first machine's backup in this case will be 2 GB, and that of the second machine will be 1 GB. In a non-deduplicating vault, the backups would occupy 4 GB in total. As a result, the deduplication ratio is 4:3, or about 1.33:1.

Similarly, in case of three machines, the ratio becomes 1.5:1; for four machines, it is 1.6:1. It approaches 2:1 as more such machines are backed up to the same vault. This means that you can buy, say, a 10-TB storage device instead of a 20-TB one.

The actual amount of capacity reduction is influenced by numerous factors such as the type of data that is being backed up, the frequency of the backup, and the backups' retention period.

2.15.6.6 Deduplication restrictions

Block-level deduplication restrictions

During a disk backup to an archive in a deduplicating vault, deduplication of a volume's disk blocks is not performed in the following cases:

- If the volume is a compressed volume
- If the volume's allocation unit size—also known as cluster size or block size—is not divisible by 4 KB

Tip: The allocation unit size on most NTFS and ext3 volumes is 4 KB and so allows for block-level deduplication. Other examples of allocation unit sizes allowing for block-level deduplication include 8 KB, 16 KB, and 64 KB.

If you protected the archive with a password

Tip: If you want to protect the data in the archive while still allowing it to be deduplicated, leave the archive non-password-protected and encrypt the deduplicating vault itself with a password, which you can do when creating the vault.

Disk blocks that were not deduplicated are stored in the archive as they would be in a non-deduplicating vault.

File-level deduplication restrictions

During a file backup to an archive in a deduplicating vault, deduplication of a file is not performed in the following cases:

- If the file is encrypted and the **In archives, store encrypted files in decrypted state** check box in the backup options is cleared (it is cleared by default)
- If the file is less than 4 KB in size
- If you protected the archive with a password

Files that were not deduplicated are stored in the archive as they would be in a non-deduplicating vault.

Deduplication and NTFS data streams

In the NTFS file system, a file may have one or more additional sets of data associated with it—often called *alternate data streams*.

When such file is backed up, so are all its alternate data streams. However, these streams are never deduplicated—even when the file itself is.

2.15.7 Privileges for centralized management

This section describes the users' privileges that are required to manage a machine locally and remotely, to manage a machine registered on Acronis Backup & Recovery 10 Management Server, and to access and manage Acronis Backup & Recovery 10 Storage Node.

2.15.7.1 Types of connection to a managed machine

There are two types of connection to a managed machine: local connection and remote connection.

Local connection

The local connection is established between Acronis Backup & Recovery 10 Management Console on a machine and Acronis Backup & Recovery 10 Agent on the same machine.

To establish a local connection

• On the toolbar, click **Connect**, then point to **New connection**, and then click **This machine**.

Remote connection

A remote connection is established between Acronis Backup & Recovery 10 Management Console on one machine and Acronis Backup & Recovery 10 Agent on another machine.

You might need to specify logon credentials to establish a remote connection.

To establish a remote connection

- 1. On the toolbar, click **Connect**, then point to **New connection**, and then click **Manage a remote** machine.
- 2. In **Machine**, type or select the name or IP address of the remote machine to which you want to connect; or click **Browse** to select the machine from the list.
- 3. To specify credentials for connection, click **Options** and then type the user name and password in the **User name** and **Password** boxes respectively. In Windows, if you leave the **User name** box empty, the credentials under which the console is running will be used.
- 4. To save the password for the specified user name, select the **Save password** check box; the password will be saved in a secure storage on the machine where the console is running.

2.15.7.2 Privileges for local connection

Windows

Local connection on a machine running Windows can be established by any user who has the "Log on locally" user right on the machine.

Linux

Establishing a local connection on a machine running Linux, and managing such machine, requires the root privileges on it.

To establish a local connection as the root user

1. If you are logged on as the root user, run the following command:

/usr/sbin/acronis_console

Otherwise, run the following command:

su -c /usr/sbin/acronis_console

2. Click Manage this machine.

To allow a non-root user to start the console

As a root user, add the name of the non-root user whom you want to allow to start the console, to the file /etc/sudoers—for example, by using the visudo command.

Caution: As a result of this procedure, the non-root user will not only be allowed to start the console with the root privileges, but also may be able to perform other actions as the root user.

To establish a local connection as a non-root user

- 1. Make sure that the root user has allowed you to start the console, as described in the previous procedure.
- Run the following command: sudo /usr/sbin/acronis_console
- 3. Click Manage this machine.

2.15.7.3 Privileges for remote connection in Windows

To establish a remote connection to a machine running Windows, the user must be a member of the Acronis Remote Users security group on that machine.

After remote connection is established, the user has management rights on the remote machine as described in User rights on a managed machine (p. 33).

Note: On a remote machine running Windows Vista with enabled User Account Control (UAC)—and which is not part of a domain—only the built-in Administrator user can back up data and perform disk management operations. To overcome the restriction, include the machine into a domain or disable UAC on the machine (by default, UAC is enabled). The same applies to machines running Windows Server 2008 and Windows 7.

For information about Acronis security groups and their default members, see Acronis security groups (p. 83).

2.15.7.4 Privileges for remote connection in Linux

Remote connections to a machine running Linux—including those performed by the root user—are established according to authentication policies, which are set up by using Pluggable Authentication Modules for Linux, known as Linux-PAM.

For the authentication policies to work, we recommend installing the latest version of Linux-PAM for your Linux distribution. The latest stable source code of Linux-PAM is available at Linux-PAM source code Web page.

Remote connection as the root user

Remote connections by the root user are established according to the Acronisagent authentication policy, which is automatically set up during the installation of Acronis Backup & Recovery 10 Agent for Linux, by creating the file /etc/pam.d/Acronisagent with the following content:

```
#%PAM-1.0
auth required pam_unix.so
auth required pam_succeed_if.so uid eq 0
account required pam_unix.so
```

Remote connection as a non-root user

Since accessing the system as the root user should be restricted, the root user can create an authentication policy to enable remote management under non-root credentials.

The following are two examples of such policies.

Note: As a result, the specified non-root users will be able to connect to the machine remotely as if they were root users. A security best practice is to make sure that the user accounts are hard to compromise—for example, by requiring that they have strong passwords.

Example 1

This authentication policy uses the pam_succeed_if module and works with Linux distributions with kernel version 2.6 or later. For an authentication policy which works with kernel version 2.4, see the next example.

Perform the following steps as the root user:

- Create the Acronis_Trusted group account, by running the following command: groupadd Acronis Trusted
- 2. Add the names of the non-root users, whom you want to allow to connect to the machine remotely, to the **Acronis_Trusted** group. For example, to add the existing user user_a to the group, run the following command:

```
usermod -G Acronis Trusted user a
```

3. Create the file /etc/pam.d/Acronisagent-trusted with the following content:

```
#%PAM-1.0

auth required pam_unix.so

auth required pam_succeed_if.so user ingroup Acronis_Trusted

account required pam_unix.so
```

Example 2

The above authentication policy might not work on Linux distributions with kernel version 2.4— including Red Hat Linux and VMware ESX™ 3.5 Upgrade 2—because the pam_succeed_if.so module is not supported there.

In this case, you can use the following authentication policy.

- 1. As the root user, create the file /etc/pam.d/Acronis_trusted_users
- 2. Add the names of the non-root users whom you want to allow to manage the machine, to this file, one user name per line. For example, if you want to add the users user_a, user_b, and user_c, add the following three lines to the file:

```
user_a
user_b
user_c
```

If necessary, also add the root user to the file.

3. Create the file /etc/pam.d/Acronisagent-trusted with the following content:

```
#%PAM-1.0
auth required pam_unix.so
auth required pam_listfile.so item=user sense=allow
file=/etc/pam.d/Acronis_trusted_users onerr=fail
account required pam_unix.so
```

2.15.7.5 Acronis security groups

On a machine running Windows, Acronis security groups determine who can manage the machine remotely and act as Acronis Backup & Recovery 10 Management Server administrator.

These groups are created when Acronis Backup & Recovery 10 Agents or Acronis Backup & Recovery 10 Management Server are being installed. During installation, you can specify what users to include in each group.

Acronis Backup & Recovery 10 Agents

When Acronis Backup & Recovery 10 Agent for Windows is being installed on a machine, the **Acronis Remote Users** group is created (or updated).

A user who is a member of this group can manage the machine remotely by using Acronis Backup & Recovery 10 Management Console, according to the management rights described in Users' privileges on a managed machine (p. 33).

By default, this group includes all members of the Administrators group.

Acronis Backup & Recovery 10 Management Server

When Acronis Backup & Recovery 10 Management Server is being installed on a machine, two groups are created (or updated):

Acronis Centralized Admins

A user who is a member of this group is a management server administrator. Management server administrators can connect to the management server by using Acronis Backup & Recovery 10 Management Console; they have the same management rights on the registered machines as users with administrative privileges on those machines—regardless of the contents of Acronis security groups there.

To be able to connect to the management server *remotely*, an administrator of the management server must also be a member of the Acronis Remote Users group.

No user—even a member of the Administrators group—can be an administrator of the management server without being a member of the Acronis Centralized Admins group.

By default, this group includes all members of the Administrators group.

Acronis Remote Users

A user who is a member of this group can connect to the management server remotely by using Acronis Backup & Recovery 10 Management Console—provided that the user is also a member of the Acronis Centralized Admins group.

By default, this group includes all members of the Administrators group.

On a domain controller

If a machine is a domain controller in an Active Directory domain, the names and default contents of Acronis security groups are different:

- Instead of Acronis Remote Users and Acronis Centralized Admins, the groups are named DCNAME \$ Acronis Remote Users and DCNAME \$ Acronis Centralized Admins respectively; here, DCNAME stands for the NetBIOS name of the domain controller. Each dollar sign is surrounded by a single space on either side.
- Instead of explicitly including the names of all members of the Administrators group, the Administrators group itself is included.

Tip: To ensure proper group names, you should install Acronis components in a domain controller after you have set up the domain controller itself. If the components were installed before you set up the domain controller, create the groups DCNAME **\$ Acronis Remote Users** and DCNAME **\$ Acronis Centralized Admins** manually, and then include the members of Acronis Remote Users and Acronis Centralized Admins in the newly created groups.

2.15.7.6 Management server administrator rights

Normally, the Acronis Backup & Recovery 10 Management Server administrator operates on a registered machine on behalf of the Acronis Managed Machine Service (also known as the Acronis service) on that machine and has the same privileges as the service has.

Alternatively, when creating a backup policy, the management server administrator has the option to explicitly specify a user account under which the centralized backup plans will run on the registered machines. In this case, the user account must exist on all the machines to which the centralized policy will be deployed. This is not always efficient.

To be a management server administrator, the user must be a member of the Acronis Centralized Admins group on the machine where the management server is installed.

2.15.7.7 User privileges on a storage node

The scope of a user's privileges on Acronis Backup & Recovery 10 Storage Node depends on the user's rights on the machine where the storage node is installed.

A regular user, such as a member of the Users group on the storage node, can:

- Create archives in any centralized vault managed by the storage node
- View and manage archives owned by the user

A user who is a member of the Administrators group on the storage node can additionally:

- View and manage any archive in any centralized vault managed by the storage node
- Create centralized vaults to be managed by the storage node—provided that the user is also an Acronis Backup & Recovery 10 Management Server administrator
- Re-schedule the compacting task, as described in Operations with storage nodes (p. 337), under "Change the compacting task schedule"

Users with these additional privileges are also called storage node administrators.

Recommendations on user accounts

To allow users to access the centralized vaults managed by a storage node, you must ensure that those users have a right to access the storage node from the network.

If both the users' machines and the machine with the storage node are in one Active Directory domain, you probably do not need to perform any further steps: all users are typically members of the Domain Users group and so can access the storage node.

Otherwise, you need to create user accounts on the machine where the storage node is installed. We recommend creating a separate user account for each user who will access the storage node, so that the users are able to access only the archives they own.

When creating the accounts, follow these guidelines:

- For users whom you want to act as storage node administrators, add their accounts to the **Administrators** group.
- For other users, add their user accounts to the Users group.

Additional right of machine administrators

A user who is a member of the Administrators group on a machine can view and manage any archives created *from that machine* in a managed vault—regardless of the type of that user's account on the storage node.

Example

Suppose that two users on a machine, UserA and UserB, perform backups from this machine to a centralized vault managed by a storage node. On the storage node, let these users have regular (non-administrative) accounts UserA SN and UserB SN, respectively.

Normally, UserA can access only the archives created by UserA (and owned by UserA_SN), and UserB can access only the archives created by UserB (and owned by UserB_SN).

However, if UserA is a member of the Administrators group on the machine, this user can additionally access the archives created from this machine by UserB—even though UserA's account on the storage node is a regular one.

2.15.7.8 Rights for Acronis services

The Acronis Backup & Recovery 10 Agent for Windows, Acronis Backup & Recovery 10 Management Server, and Acronis Backup & Recovery 10 Storage Node components run as services. When installing any of these components, you need to specify the account under which the component's service will run.

For each service, you can either create a dedicated user account (recommended in most cases) or specify an existing account of a local or domain user—for example: .\LocalUser or DomainName\DomainUser.

If you choose to create dedicated user accounts for the services, the setup program will create the following user accounts:

- For the Acronis Backup & Recovery 10 Agent for Windows service, Acronis Agent User
- For the Acronis Backup & Recovery 10 Management Server service, AMS User
- For the Acronis Backup & Recovery 10 Storage Node service, ASN User

The newly created accounts are given the following privileges:

- All three accounts are assigned the Log on as a service user right.
- The Acronis Agent User user account is assigned the Adjust memory quotas for a process and Replace a process level token user rights.
- The Acronis Agent User and ASN User user accounts are included in the Backup Operators group.
- The AMS User user account is included in the Acronis Centralized Admins group.

The setup program will assign the above listed user rights to any existing account you specify for a corresponding service.

If you choose to specify an existing user account for the agent service or the storage node service, make sure that this account is a member of the **Backup Operators** group, before proceeding with the installation.

If you choose to specify an existing user account for the management server service, this account will be added to the **Acronis Centralized Admins** group automatically.

If the machine is part of an Active Directory domain, make sure that the domain's security policies do not prevent the accounts described in this section (whether existing or newly created) from having the above listed user rights.

Important: After the installation, do not specify a different user account for a component's service. Otherwise, the component may stop working.

The newly created user accounts are also granted access to the registry key

HKEY_LOCAL_MACHINE\SOFTWARE\Acronis (called Acronis registry key) with the following rights:

Query Value, Set Value, Create Subkey, Enumerate Subkeys, Notify, Delete, and Read Control.

In addition, there are two Acronis services which run under a system account:

- The Acronis Scheduler2 Service provides scheduling for Acronis components' tasks. It runs under the Local System account and cannot run under a different account.
- The Acronis Remote Agent Service provides connectivity among Acronis components. It runs under the Network Service account and cannot run under a different account.

2.15.8 Communication between Acronis Backup & Recovery 10 components

This section describes how Acronis Backup & Recovery 10 components communicate with each other using secure authentication and encryption.

This section also provides information on configuring communication settings, selecting a network port for communication, and managing security certificates.

2.15.8.1 Secure communication

Acronis Backup & Recovery 10 provides the capability to secure the data transferred between its components within a local area network and through a perimeter network (also known as demilitarized zone, DMZ).

There are two mechanisms which ensure secure communication between Acronis Backup & Recovery 10 components:

- Secure authentication provides secure transfer of certificates needed to establish a connection, by using the Secure Sockets Layer (SSL) protocol.
- Encrypted communication provides secure transfer of information between any two
 components—for example, between Acronis Backup & Recovery 10 Agent and Acronis Backup &
 Recovery 10 Storage Node—by encrypting the data being transferred.

For instructions on how to set up secure authentication and data encryption settings, see Configuring communication options (p. 88).

For instructions on how to manage SSL certificates used for secure authentication, see SSL certificates (p. 91).

Note: The components of earlier Acronis products, including those of the Acronis True Image Echo family, cannot connect to the Acronis Backup & Recovery 10 components, regardless of the secure authentication and data encryption settings.

2.15.8.2 Client and server applications

There are two stakeholders of the secure communication process:

- Client application, or client, is an application that tries to establish connection.
- **Server application**, or server, is an application to which the client tries to connect.

For example, if Acronis Backup & Recovery 10 Management Console is connecting to Acronis Backup & Recovery 10 Agent on a remote machine, the former is the client and the latter is the server.

An Acronis component can act as a client application, a server application, or both, as shown in the following table.

Component name	Can be client	Can be server
Acronis Backup & Recovery 10 Management Console	Yes	No
Acronis Backup & Recovery 10 Agent	Yes	Yes
Acronis Backup & Recovery 10 Management Server	Yes	Yes
Acronis Backup & Recovery 10 Storage Node	Yes	Yes
Acronis PXE Server	No	Yes
Acronis Backup & Recovery 10 Bootable Agent	Yes	Yes

2.15.8.3 Configuring communication settings

You can configure communication settings, such as whether to encrypt transferred data, for Acronis Backup & Recovery 10 components installed on one or more machines, by using Acronis Administrative Template. For information on how to load the Administrative Template, see How to load Acronis Administrative Template (p. 351).

When applied to a single machine, the Administrative Template defines the communication settings for all the components on the machine; when applied to a domain or an organizational unit, it defines the communication settings for all the components on the machines in that domain or organizational unit.

To configure communication settings

- 1. Click Start, then click Run, and then type gpedit.msc
- 2. In the **Group Policy** console, expand **Computer Configuration**, then expand **Administrative Templates**, and then click Acronis.
- 3. In the Acronis pane to the right, double-click a communication option that you want to configure. The Administrative Template contains the following options (each option is explained later in this topic):
 - Remote Agent ports
 - Client Encryption options
 - Server Encryption options

- 4. For the new communication settings to take effect, restart all running Acronis components—preferably, by restarting Windows. If restart is not possible, make sure you do the following:
 - If Acronis Backup & Recovery 10 Management Console is running, close it and start it again.
 - If other Acronis components, such as Acronis Backup & Recovery 10 Agent for Windows or Acronis Backup & Recovery 10 Management Server are running, restart their correspondent services from the **Services** snap-in in Windows.

Remote Agent ports

Specifies the port that the component will use for incoming and outgoing communication with other Acronis components.

Select one of the following:

Not configured

The component will use the default TCP port number 9876.

Enabled

The component will use the specified port; type the port number in the **Server TCP Port** box.

Disabled

The same as Not configured.

For details about the network port and instructions on how to specify it in Linux and a bootable environment, see Network port configuration (p. 91).

Client Encryption options

Specifies whether to encrypt the transferred data when the component acts as a client application, and whether to trust self-signed SSL certificates.

Select one of the following:

Not configured

The component will use the default settings, which is to use encryption if possible and to trust self-signed SSL certificates (see the following option).

Enabled

Encryption is enabled. In **Encryption**, select one of the following:

Enabled

Data transfer will be encrypted if encryption is enabled on the server application, otherwise it will be unencrypted.

Disabled

Encryption is disabled; any connection to a server application which requires encryption will not be established.

Required

Data transfer will be performed only if encryption is enabled on the server application (see "Server Encryption options"); it will be encrypted.

Authentication parameters

Selecting the **Trust self-signed certificates** check box allows the client to connect to the server applications that use self-signed SSL certificates such as certificates created during the installation of Acronis Backup & Recovery 10 components—see SSL certificates (p. 91).

You should keep this check box selected, unless you have a Public Key Infrastructure (PKI) in your environment.

In Use Agent Certificate Authentication, select one of the following:

Do not use

The use of SSL certificates is disabled. Any connection to a server application which requires the use of SSL certificates will not be established.

Use if possible

The use of SSL certificates is enabled. The client will use SSL certificates if their use is enabled on the server application, and will not use them otherwise.

Always use

The use of SSL certificates is enabled. The connection will be established only if the use of SSL certificates is enabled on the server application.

Disabled

The same as Not configured.

Server Encryption options

Specifies whether to encrypt the transferred data when the component acts as a server application.

Select one of the following:

Not configured

The component will use the default setting, which is to use encryption if possible (see the following option).

Enabled

Encryption is enabled. In **Encryption**, select one of the following:

Enabled

Data transfer will be encrypted if encryption is enabled on the client application, otherwise it will be unencrypted.

Disabled

Encryption is disabled; any connection to a client application which requires encryption will not be established.

Required

Data transfer will be performed only if encryption is enabled on the client application (see "Client Encryption options"); it will be encrypted.

Authentication parameters

In Use Agent Certificate Authentication, select one of the following:

Do not use

The use of SSL certificates is disabled. Any connection to a client application which requires the use of SSL certificates will not be established.

Use if possible

The use of SSL certificates is enabled. The server will use SSL certificates if their use is enabled on the client application, and will not use them otherwise.

Always use

The use of SSL certificates is enabled. The connection will be established only if the use of SSL certificates is enabled on the client application.

Disabled

The same as **Not configured**.

2.15.8.4 Network port configuration

Acronis Backup & Recovery 10 components use the 9876/TCP network communication port by default. The server listens to this port for incoming connection. This port is also used as default by the Acronis client. During component installation you might be asked to confirm the port opening or to open the port manually, in case you are using a firewall other than Windows Firewall.

After installation, you can change the ports at any time to match your preferable values or for the purpose of security. This operation requires the restart of Acronis Remote Agent (in Windows) or the Acronis_agent (in Linux) service.

After the port is changed on the server side, connect to the server using the <Server-IP>:<port> or the <Server-hostname>:<port> URL notation.

Note: If you use network address translation (NAT), you can also configure the port by setting up port mapping.

Configuring the port in the operating system

Windows

To be able to change the ports' numbers, load and configure the Administrative Template, provided by Acronis, as described in Configuring communication settings (p. 88), under "Remote Agent ports".

Linux

Specify the port in the /etc/Acronis/Policies/Agent.config file. Restart the Acronis agent daemon.

Configuring the port in a bootable environment

While creating Acronis bootable media, you have the option to pre-configure the network port that will be used by the Acronis Backup & Recovery 10 Bootable Agent. The choice is available between:

- The default port (9876)
- The currently used port
- New port (enter the port number)

If a port has not been pre-configured, the agent uses the default port number.

2.15.8.5 SSL certificates

Acronis Backup & Recovery 10 components use Secure Sockets Layer (SSL) certificates for secure authentication.

SSL certificates for the components can be one of the two types:

- Self-signed certificates, such as certificates automatically generated during the installation of an Acronis component.
- Non-self-signed certificates, such as certificates issued by a third-party Certificate Authority (CA)—for example, by a public CA such as VeriSign[®] or Thawte™—or by your organization's CA.

Certificate path

All Acronis components installed on a machine, when acting as a server application, use an SSL certificate called the server certificate.

In Windows, the certificate path and the server certificate's file name are specified in the registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\Encryption\Server**. The default path is %SystemDrive%\Program Files\Common Files\Acronis\Agent.

For self-signed certificates, the certificate thumbprint (also known as fingerprint or hash) is used for future host identification: if a client has previously connected to a server by using a self-signed certificate and tries to establish connection again, the server checks whether the certificate's thumbprint is the same as the one used before.

Self-signed certificates

On machines running Windows, if the certificate location contains no server certificate, a self-signed server certificate is automatically generated and installed during the installation of any Acronis component except Acronis Backup & Recovery 10 Management Console.

If the machine is renamed after its self-signed certificate was generated, the certificate cannot be used and you will need to generate a new one.

To generate a new self-signed certificate

- 1. Log on as a member of the Administrators group.
- 2. In the Start menu, click Run, and then type: cmd
- 3. Run the following command (note quotation marks):
 "%CommonProgramFiles%\Acronis\Utils\acroniscert" --reinstall
- 4. Restart Windows, or restart the running Acronis services.

Non-self-signed certificates

You have the option to use trusted third-party certificates or certificates created by your organization's CA as an alternative to self-signed certificates, by using Acronis Certificate Command-line Utility.

To install a third-party certificate

- 1. Click **Start**, then click **Run**, and then type: **certmgr.msc**
- 2. In the Certificates console, double-click the name of the certificate that you want to install.
- 3. In the **Details** tab, in the list of fields, click **Thumbprint**.
- 4. Select and copy the field's value, called a certificate thumbprint—a string such as **20 99 00 b6 3d 95 57 28 14 0c d1 36 22 d8 c6 87 a4 eb 00 85**
- 5. In the Start menu, click Run, and then type the following in the Open box: "%CommonProgramFiles%\Acronis\Utils\acroniscert.exe" --install "20 99 00 b6 3d 95 57 28 14 0c d1 36 22 d8 c6 87 a4 eb 00 85" (Note quotation marks; substitute the sample thumbprint shown here with that of your certificate.)

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. 417) or to the management server (p. 418).

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. 419).

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 Management server options

The management server options enable you to adjust the behavior of the Acronis Backup & Recovery 10 Management Server.

To access the management server options, connect the console to the management server and then select **Options > Management server options** from the top menu.

3.2.1 Logging level

This option defines whether the management server has to collect log events from the registered machines to the centralized log that is stored in a dedicated database and is available in the **Log** view. You can set the option for all the events at once or select the event types to be collected. If you completely disable collection of the log events, the centralized log will contain only the management server's own log.

The preset is: **Collect logs** for **All events**.

Use the **Types of events to log** combo-box to specify the types of events that will be collected:

- All events all events (information, warnings and errors) occurred on all the machines registered on the management server will be recorded to the centralized log
- Errors and warnings warnings and errors will be recorded to the centralized log
- Errors only only errors will be recorded to the centralized log.

To disable collection of the log events, clear the **Collect logs** check box.

3.2.2 Log cleanup rules

This option specifies how to clean up the centralized event log stored in the management server's reporting database.

This option defines the maximum size of the reporting database.

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.

Even if you remove the log size limit, logging events to an SQL Server Express database will stop after the log size reaches 4 GB, because SQL Express Edition has the 4 GB per database limit. Set the maximum log size to approximately 3.8 GB if you want to use the maximum capacity of the SQL Express database.

This parameter can also be set by using Acronis Administrative Template (p. 355).

3.2.3 Event tracing

You can configure the management server to log events in the Application Event Log of Windows, besides the management server's own log.

You can configure the management server to send Simple Network Management Protocol (SNMP) objects to a specified SNMP manager.

3.2.3.1 Windows event log

This option defines whether the management server has to record its own 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 recorded.

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.

3.2.3.2 SNMP notifications

This option defines whether the management server has to send its own log events 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. 55)".

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.

3.2.4 Domain access credentials

This option determines the user name and password that the management server will use to access the domain.

The preset is: No credentials

The management server needs domain access credentials when working with a dynamic group that is based on the **Organizational unit** criterion (p. 325). When you are creating such group and no credentials are given by this option, the program will ask you for credentials and save them in this option.

It is sufficient to specify the credentials of a user who is a member of the **Domain Users** group on the domain.

3.2.5 Acronis WOL Proxy

This option works in combination with the **Use Wake-On-LAN** (p. 182) advanced scheduling setting. Use this option if the management server has to wake up for backup machines located in another subnet.

When the scheduled operation is about to start, the management server sends out magic packets to wake up the appropriate machines. (A magic packet is a packet that contains 16 contiguous copies of the receiving NIC's MAC address). The Acronis WOL Proxy, installed in the other subnet, transfers the packets to machines located in that subnet.

The preset is: Disabled.

To bring this option into use:

- 1. Install Acronis WOL Proxy on any server in the subnet where the machines to be woken are located. The server has to provide continuous services availability. With multiple subnets, install Acronis WOL Proxy in every subnet where you need to use the Wake-On-LAN functionality.
- 2. Enable Acronis WOL Proxy in the Management server options as follows:
 - a. Select the **Use the following proxies** check box.
 - b. Click **Add**, and then enter the name or IP address of the machine where the Acronis WOL Proxy is installed. Provide access credentials for the machine.
 - c. Repeat this step if there are several Acronis WOL Proxies.
- 3. When scheduling a backup policy, enable the Use Wake-On-LAN setting.

You also have the ability to delete proxies from the list. Please keep in mind that any change to this option affects the entire management server. If you delete a proxy from the list, the Wake-On-LAN functionality in the corresponding subnet will be disabled for all policies, including the policies already applied.

3.2.6 VM protection options

These options define the management server behavior as related to backup and recovery of virtual machines hosted on virtualization servers.

3.2.6.1 VMware vCenter integration

This option defines whether to show virtual machines managed by a VMware vCenter Server in the management server and show the backup status of these machines in the vCenter.

Integration is available in all Acronis Backup & Recovery 10 advanced editions; a license for Virtual Edition is not required. No software installation is required on the vCenter Server.

On the management server side

When integration is enabled, the vCenter's **VMs and Templates** inventory view appears in the management server's GUI under **Navigation** > **Virtual machines**.

From the management server's standpoint, this is a dynamic group of virtual machines. The group name matches the vCenter Server name or IP address, whatever was specified when configuring integration. The group content is synchronized with the vCenter Server and cannot be changed on the management server side. In case of an occasional inconsistency, right click the group and select **Refresh**.

The virtual machines managed by the vCenter Server also appear in the **All virtual machines** group. You can view virtual machine properties and power state; create groups of virtual machines and add virtual machines to existing groups.

Backup and recovery of a virtual machine is not possible, unless Acronis Backup & Recovery 10 Agent for ESX/ESXi is deployed (p. 330) to the virtual machine's host. Such machines appear as not manageable (grayed out).

Once the agent is deployed to an ESX/ESXi host (this requires a license for Acronis Backup & Recovery 10 Advanced Server Virtual Edition), the virtual machines from this host are ready for applying a backup policy or individual backup. Such machines appear as manageable.

If Agent for Windows or Agent for Linux is installed in a guest system, but there is no Agent for ESX/ESXi on its host, the virtual machine appears as not manageable under **Virtual machines**. Such machine has to be managed as a physical one.

On the vCenter Server side

When integration is enabled, the vCenter Server will store and show information about when and how successful each virtual machine was backed up. The same information is displayed in the **Status** and the **Last backup** columns on the management server.

Backup status - the most severe status of all backup plans and backup policies on the machine. For more information, see "Backup plan statuses (p. 191)" and "Policy status on a machine (p. 74)".

Last backup - how much time has passed since the last successful backup.

You can see this information in the virtual machine summary (**Summary > Annotations**) or on the **Virtual Machines** tab for every host, datacenter, folder or entire vCenter Server (for example, **View > Inventory > Hosts and Clusters > select** host **> Virtual Machines**).

3.2.7 Online backup proxy

This option is effective only for connection to Acronis Online Backup Storage over the Internet.

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

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

Proxy settings for the agent and the management server are configured separately, even if both are installed on the same machine.

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**.

3.3 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.

3.3.1 Machine management

This option defines whether the machine has to be managed centrally by the Acronis Backup & Recovery 10 Management Server.

To be able to use this option, you must be logged on as a member of the **Administrators** group on the machine.

You have the opportunity to register the machine on the management server when installing an Acronis Backup & Recovery 10 agent. If the machine is not registered, selecting **Centralized management** here will initiate the registration (p. 419). Or you can add the machine to the management server on the server side. Any of the three registration methods require the server administrator privileges.

Selecting **Stand-alone management** on a registered machine will result in the machine stopping communication with the server. On the management server, the machine appears as **Withdrawn**. The management server administrator can delete the machine from the server or register the machine once again.

The preset is: Stand-alone management.

To set up centralized management on the machine:

- 1. Select Centralized management.
- 2. Specify the Management Server IP/Name.
- 3. Specify the user name and password of the management server administrator on prompt.
- 4. In the **Machine's registration address**, select how the machine will be registered on the management server: by its name (recommended) or by its IP address.
- 5. Click **OK** and the machine will be registered on the management server.

To disable centralized management, select **Stand-alone management**.

3.3.2 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 that occur during backup or during recovery, in the Default backup and recovery options (p. 103). 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.

3.3.2.1 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. 103). 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.

3.3.2.2 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. 103). 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. 55)".

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. 101).

3.3.2.3 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 lmmib2.dll that can be found on the installation disc of your operating system.

Linux

To receive SNMP messages on a machine running Linux, the net-snmp (for RHEL and SUSE) or the snmpd (for Debian) package has to be installed.

SNMP can be configured using the **snmpconf** command. The default configuration files are located in the /etc/snmp directory:

- /etc/snmp/snmpd.conf configuration file for the Net-SNMP SNMP agent
- /etc/snmp/snmptrapd.conf configuration file for the Net-SNMP trap daemon.

3.3.3 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 (p. 359).

3.3.4 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.

3.3.5 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 (p. 362). A setting defined by a Group Policy cannot be changed using the product GUI unless the Group Policy is disabled on the machine.

3.4 Default backup and recovery options

3.4.1 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, Linux, 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		Agent for Linux		Bootable media (Linux-based or PE-based)	
	Disk backup	File backup	Disk backup	File backup	Disk backup	File backup
Archive protection (p. 105) (password + encryption)	+	+	+	+	+	+
Source files exclusion (p. 106)	+	+	+	+	+	+

Pre/Post backup commands (p. 107)	+	+	+	+	PE only	PE only
Pre/Post data capture commands (p. 109)	+	+	+	+	-	-
Multi-volume snapshot (p. 111)	+	+	-	-	-	-
File-level backup snapshot (p. 110)	-	+	-	+	-	-
Use VSS (p. 111)	+	+	-	-	-	-
Compression level (p. 112)	+	+	+	+	+	+
Backup performance:						
Backup priority (p. 113)	+	+	+	+	-	-
HDD writing speed (p. 113)	Dest: HDD					
Network connection speed (p. 113)	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share
Fast incremental/differential backup (p. 117)	+	-	+	-	+	-
Backup splitting (p. 117)	+	+	+	+	+	+
File-level security (p. 118):		1	1	ı	1	ı
Preserve files' security settings in archives	-	+	-	-	-	-
In archives, store encrypted files in decrypted state	-	+	-	-	-	-
Media components	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	-	-
Error handling (p. 119):						
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Ignore bad sectors	+	+	+	+	+	+
Dual destination (p. 120)	Dest: local	Dest: local	Dest: local	Dest: local	-	-
Task start conditions (p. 120)	+	+	+	+	-	-
Task failure handling (p. 121)	+	+	+	+	-	-
Tape support (p. 122)	Dest: managed vault on a tape library					
Additional settings (p. 123):						

Overwrite data on a tape without prompting for user confirmation	Dest: Tape					
Dismount media after backup has finished	Dest: removable media	Dest: removable media	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	Dest: removable media	Dest: removable media
Reset archive bit	-	+	-	-	-	+
Restart the machine automatically after backup is finished	-	-	-	-	+	+
Deduplicate backup only after transferring it to the vault	Dest: dedup. vault					
Use FTP in Active mode	Dest: FTP server					
Save software RAID and LVM metadata along with backups	-	-	+	-	-	-
Notifications:						
E-mail (p. 114)	+	+	+	+	-	-
Win Pop-up (p. 115)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 116)	+	+	-	-	-	-
SNMP (p. 116)	+	+	+	+	-	-

3.4.1.1 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.

3.4.1.2 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.

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 (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							
		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 /home/user/Finance		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.

3.4.1.3 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	Backup	Post-backup
command		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.		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.

3.4.1.4 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.

	<> Backup>				
Pre-backup	Pre-data	Data	Post-data		Post-backup
command	capture	capture	capture		command
	command		command		

If the Volume Shadow Copy Service (p. 111) 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. 107), 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	Selected	Selected	Cleared	Cleared

capture until the command execution is complete				
	Res	sult		
	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		
	Res	ult				
	Preset Continue the backup only after the command is successfully executed. Delete the TIB file and temporary files and fail the task if the command execution fails.		N/A	Continue the backup concurrently with the command execution and irrespective of the command execution result.		

3.4.1.5 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.

3.4.1.6 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. 110) 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.

3.4.1.7 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. 109) 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.

3.4.1.8 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.

3.4.1.9 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.

3.4.1.10 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 failed

The **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.

3.4.1.11 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. 99).
- 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. 55)".

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. 99).
- 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.

3.4.1.12 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.

3.4.1.13 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

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.

3.4.1.14 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. 129) 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**.

3.4.1.15 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. 411) 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.

3.4.1.16 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.

3.4.1.17 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.

3.4.1.18 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. 173) and Conditions (p. 184).


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



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.

3.4.1.19 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.

3.4.1.20 Tape support

These options are effective when the backup destination is a managed vault located on a tape library.

Tape support options enable you to specify how the backup tasks will distribute backups among the tapes.

Some combinations of tape options might degrade usage efficiency of both the whole tape library and each tape. If you are not forced to modify these options by some specific needs, leave them unchanged.

An archive can occupy several tapes. In such cases a so-called **tape set** is used for keeping the data backups.

Tape set is a logical group of one or more tapes which contain backups of the specific protected data. A tape set can contain backups of other data as well.

Separate tape set is a tape set which contains only backups of the specific protected data. Other backups cannot be written to a separate tape set.

(For the backup policy/plan to be created) Use a separate tape set

The preset is: Disabled.

If you leave this option unchanged, then the backups, belonging to the policy or plan being created, might be written onto tapes containing backups written by different backup policies and comprising of data from different machines. Similarly, backups from other policies might be written onto the tapes containing this policy's backups. You will not have a problem with such tapes, as the program manages all the tapes automatically.

When this option is enabled, the backups, belonging to the policy or plan being created, will be located on a separate tape set. Other backups will not be written to this tape set.

If the console is connected to the management server

The **Use a separate tape set** option has more precise definitions. So for the backup policy to be created you can use a separate tape set for all machines or for each single machine.

The **A single tape set for all machines** option is selected by default. Generally this option ensures more efficient usage of tapes, than the **A separate tape set for each single machine** option. However the second one can be useful, for example, when there are special requirements to store the tapes with backups from a specific machine off-site.

When the **Use a separate tape set** option is enabled, there might be a case when the backup has to be written onto a tape that is currently out of the tape library device. Define what to do in this case.

- Ask for user interaction the backup task will enter the Need Interaction state and wait for the tape, with the required label, to be loaded into the tape library device.
- Use a free tape the backup will be written onto a free tape, so the operation will be paused only if there is no free tape in the library.

Always use a free tape

If you leave the options below unchanged, then each backup will be written onto the tape specified by the **Use a separate tape set** option. With some of the options below enabled, the program will add new tapes to the tape set every time when a full, incremental or differential backup is created.

■ For each full backup

The preset is: Disabled.

When this option is enabled, each full backup will be written onto a free tape. The tape will be loaded to a drive especially for this operation. If the **Use a separate tape set** option is enabled, only incremental and differential backups of the same data will be appended to the tape.

For each differential backup

The preset is: Disabled.

When this option is enabled, each differential backup will be written onto a free tape. This option is available only when using free tape for each full backup is selected.

For each incremental backup

The preset is: Disabled.

When this option is enabled, each incremental backup will be written onto a free tape. This option is available only when using free tape for each full and differential backup is selected.

3.4.1.21 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 (p. 76).

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. 281).

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.

3.4.2 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, Linux, bootable media)
- The type of data being recovered (disk, file)
- The operating system being recovered from the disk backup (Windows, Linux)

The following table summarizes the availability of the recovery options.

	Agent for Windows		Agent for Linux		Bootable media (Linux-based or PE-based)	
	Disk recovery	File recovery (also from a disk backup)	Disk recovery	File recovery (also from a disk backup)	Disk recovery	File recovery (also from a disk backup)
Pre/Post recovery commands (p. 127)	+	+	+	+	PE only	PE only
Recovery priority (p. 128)	+	+	+	+	-	-
File-level security (p. 129):				<u>'</u>		
Recover files with their security settings	-	+	-	+	-	+
Error handling (p. 132):		<u>'</u>		'		
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Additional settings (p. 132):					ı	
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	-	Windows recovery	-
Notifications:		·		,		

E-mail (p. 129)	+	+	+	+	-	-
Win Pop-up (p. 130)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 130)	+	+	-	-	-	-
SNMP (p. 131)	+	+	+	+	-	-

3.4.2.1 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	Selected	Selected	Cleared	Cleared	

complete								
Result								
	Preset Perform the recovery only after the command is successfully executed. Fail the task if the command execution failed.	Perform the recovery after the command is executed despite execution failure or success.	N/A	Perform the recovery concurrently with the command execution and irrespective of the command execution result.				

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.

3.4.2.2 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.

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

3.4.2.4 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 failed

The 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 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.

3.4.2.5 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. 99).
- 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. 55)".

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. 99).
- **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.

3.4.2.6 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.

3.4.2.7 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 an 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.

3.4.2.8 VM power management

These options are effective for virtual machines residing on the virtualization servers.

These options are available only if any Acronis agent for virtual machines is installed on the virtualization server.

Power off target virtual machines when starting recovery

The preset is: **On**.

Recovery to an existing virtual machine is not possible if the machine is online, and so the machine is powered off automatically as soon as the recovery task starts. Users will be disconnected from the machine and any unsaved data will be lost.

Clear the check box for this option if you prefer to power off virtual machines manually before the recovery.

Power on the target virtual machine when recovery is completed

The preset is: Off.

After a machine is recovered from a backup to another machine, there is a chance the existing machine's replica will appear on the network. To be on the safe side, power on the recovered virtual machine manually, after you take the necessary precautions.

Select the check box for this option if automatic powering on of the virtual machine is required.

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.

Centralized and personal vaults

A centralized vault is a networked location allotted by the management server administrator to serve as storage for the backup archives. A centralized vault can be managed by a storage node (managed vault) or be unmanaged.

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.

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.

Centralized. This group is available when the console is connected to a managed machine or to a management server. Expand this group to display a list of centralized vaults added by the management server administrator.

Click any centralized vault in the vaults tree to open the detailed view of this vault (p. 137) and to take actions on the vault (p. 138), archives (p. 169) and backups (p. 170) stored in there.

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. 167) and to take actions on the vault (p. 168), archives (p. 169) and backups (p. 170) stored in there.

4.1 Centralized vaults

A centralized vault is a networked location allotted by the management server administrator to serve as storage for the backup archives. A centralized vault can be managed by a storage node or be unmanaged. The total number and size of archives stored in a centralized vault is limited by the storage size only.

As soon as the management server administrator commits to creating a centralized vault, the vault path and name are distributed to all machines registered on the server. The shortcut to the vault appears on the machines in the **Vaults > Centralized** group. Any backup plan 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, as well as other archives stored in the vault, will be managed by the storage node.

Managed vaults

The managed vault is a centralized vault managed by a storage node. The storage node performs storage node-side cleanup (p. 420) and storage node-side validation (p. 420) for each archive stored in the managed vault. When creating a managed vault, an administrator can specify additional operations that the storage node will perform (deduplication (p. 76), encryption). Management operations cannot be canceled or disabled. They will be performed for all archives stored in the vault unless the vault is deleted.

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.

Accessing managed vaults

To be able to back up to a managed vault, a user must have an account on the machine where the storage node is installed. The scope of a user's privileges in a vault depends on the user's rights on the storage node. A user who is a member of the Users group can view and manage his/her own archives. Members of the Administrators group can view and manage any archive stored on the storage node. A user who is a member of the Administrators group on a managed machine can view and manage archives created by any user of this machine.

To learn more about privileges depending on the user rights, see the User privileges on a storage node (p. 85) section.

Unmanaged vaults

An unmanaged vault is a centralized vault that is not managed by a storage node. To access an unmanaged vault, a user has to have access privileges for the location from the network.

Any user that has permission to read/write files in an unmanaged vault can:

- Back up data to the unmanaged vault
- Recover data from any backup located in the unmanaged vault.
- View and manage all the archives located in the unmanaged vault.

4.1.1 Working with the "Centralized vault" view

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

Vault toolbar

The toolbar contains operational buttons that let you perform operations with the selected centralized vault. See the Actions on centralized vaults (p. 138) 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. The pie chart is not available if the vault is located on a tape library.

- 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.

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

- [for managed vaults only] the name of the storage node that manages 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
- [for managed vaults only] deduplication (p. 76) state (On, Off)
- [for managed vaults only] encryption state (Yes, No)

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. 169)

Operations with backups (p. 170)

Filtering and sorting archives (p. 172)

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 centralized vaults

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** bar (on the **Actions and tools** pane) and from the **[Vault name] actions** item of the main menu.

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

То	Do
Create a managed or an	1. Click Create.
unmanaged vault	2. In the Type field, select the vault type: Managed or Unmanaged
	The procedure of creating centralized vaults is described in-depth in the following sections:
	■ Create a managed centralized vault (p. 139)
	■ Create an unmanaged centralized vault (p. 142)
Edit a managed or an	1. Select the vault.
unmanaged vault	2. Click Edit .
	Depending on the vault you select (managed or unmanaged), the respective Edit page will be opened:
	■ The Edit managed vault page lets you change the vault's name, encryption password (if the vault is encrypted) and information in the Comments field.
	■ The Edit unmanaged vault page lets you edit the vault's name and information in the Comments field.
Validate a vault	1. Select the vault.
	2. Click Validate .
	You will be taken to the Validation (p. 251) page with an already pre-selected vault as a source. The vault validation checks all the archives in this vault.

Delete a vault	1. Select the vault.					
	2. Click X Delete.					
	You'll be asked whether to keep the archives stored in the vault, or delete the vault along with all the archives. The plans and tasks that use this vault will fail.					
	If you choose to keep the archives for a managed vault, the vault will be detached from the storage node. Later on, you'll be able to attach this vault to the same or to another storage node.					
Explore an unmanaged	1. Select the unmanaged vault.					
vault	2. Click Explore.					
	The vault will be available for examination with the standard file manager program.					
Attach the managed	Click Attach.					
vault that was deleted without removing its content.	The procedure of attaching a managed vault to a storage node is described in-depth in the Attaching a managed vault (p. 142) section.					
Change user credentials for accessing a vault	Click Change user . Changing user credentials is available for vaults that reside on shared storages only.					
Refresh a vault's	Click C Refresh.					
information	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.					
Actions on a tape library o	n a managed vault					
Define tape labels and	Click Manage tapes.					
perform inventorying of a tape library on a managed vault	In the Tape Management window, define labels for tapes and refresh the inventory. For more details, see the Managing tape library (p. 148) section.					
Rescan tapes in a	Click Rescan tapes.					
managed vaults	Rescan reads information about the content of user-selected tapes and updates the storage node database.					
	This operation is described in-depth in the Rescan (p. 149) section.					
	1					

4.1.2.1 Creating a managed centralized vault

To create a managed centralized vault, perform the following steps

Vault

Name

Specify a unique name for the vault. Creation of two centralized vaults with the same name is prohibited.

Comments

[Optional] Enter the distinctive description of the vault being created.

Type

Select the Managed type.

Storage node

Select the Acronis Backup & Recovery 10 Storage Node that will manage the vault. You may need to enter access credentials for the storage node.

Path (p. 140)

Specify where the vault will be created. Managed centralized vaults can reside on a network share, SAN, NAS, or on a hard drive local to the storage node.

Database path (p. 141)

Specify a local folder on the storage server to create a vault-specific database. This database will store the metadata required for cataloguing the archives and performing deduplication.

Deduplication

[Optional] Select whether to enable archive deduplication in the vault. Deduplication minimizes storage space taken by the archives and backup traffic. It reduces the size of archives in the vault by eliminating redundant data such as duplicate files or disk blocks.

Deduplication is not possible on tape devices.

To learn more about how deduplication works, see the Deduplication (p. 76) section.

Compression

[Optional] Select whether to compress the deduplication data stores. This setting is available only if deduplication is enabled.

Encryption (p. 141)

[Optional] Select whether to protect the vault with encryption. Anything written to the vault will be encrypted and anything read from it will be decrypted transparently by the storage node, using a vault-specific encryption key stored on the storage node.

After you have performed all the required steps, click **OK** to commit creating the managed vault.

Vault path

To specify the path where the managed vault will be created

- 1. Enter the full path to the folder in the **Path** field or select the desired folder in the folders tree. Managed vaults can be organized:
 - on the hard drives local to the storage node
 - on a network share
 - on a Storage Area Network (SAN)
 - on a Network Attached Storage (NAS)
 - on a tape library locally attached to the storage node.

To create a new folder for the vault in the selected location, click General Create folder.

2. Click OK.

A vault can be created in an empty folder only.

We do not recommend creating a deduplicating managed vault on a FAT32 volume. The reason is that such vault stores all deduplicated items in two potentially large files. Because the maximum file size in the FAT file systems is limited to 4 GB, the storage node may stop working when this limit is reached.

The folder permissions must allow the user account under which the storage node's service is running (by default, **ASN User**) to write to the folder. When assigning permissions, specify the user account explicitly (not just **Everyone**).

Vault database path

To specify the path where the vault's database will be created

1. In the **Local folders** of the storage node, select the desired folder or enter the full path to the folder in the **Path** field.

To create a new folder for the database, click Greate folder.

2. Click OK.

When choosing a folder for the vault's database, follow these considerations:

- The folder must reside on a fixed drive. Please do not try to place the database on external detachable drives.
- The folder size may become large—one estimate is 200 GB per 8 TB of used space, or about 2.5 percent.
- The folder permissions must allow the user account under which the storage node's service is running (by default, ASN User) to write to the folder. When assigning permissions, specify the user account explicitly (not just Everyone).

Vault encryption

If you protect a vault with encryption, anything written to the vault will be encrypted and anything read from it will be decrypted transparently by the storage node, using a vault-specific encryption key stored on the node. If 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.

This encryption has nothing to do with the archive encryption specified by the backup plan and performed by an agent. If the archive is already encrypted, the storage node-side encryption is applied over the encryption performed by the agent.

To protect the vault with encryption

- 1. Select the **Encrypt** 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:
 - AES 128 the vault contents will be encrypted using the Advanced Encryption Standard (AES) algorithm with a 128-bit key
 - AES 192 the vault contents will be encrypted using the AES algorithm with a 192-bit key
 - AES 256 the vault contents 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 archives stored in the vault and the more secure the archives 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; the password hash is used for verification purposes. With this two-level security, the archives are protected from any unauthorized access, but recovering a lost password is not possible.

4.1.2.2 Creating an unmanaged centralized vault

To create an unmanaged centralized vault, perform the following steps.

Vault

Name

Specify a unique name for the vault. The creation of two centralized vaults with the same name is prohibited.

Comments

Enter the distinctive description of the vault.

Type

Select the **Unmanaged** type.

Path (p. 142)

Specify where the vault will be created.

After you have performed all the required steps, click **OK** to commit creating the unmanaged centralized vault.

Vault path

To specify the path where the unmanaged vault will be created

- 1. Enter the full path to the folder in the Path field or select the desired folder in the folders tree. Unmanaged vaults can be organized:
 - Acronis Online Backup Storage
 - on a network share
 - on a Storage Area Network (SAN)
 - on a Network Attached Storage (NAS)
 - on FTP and SFTP servers.

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 create a new folder for the vault, click Greate folder.

A vault can be created in an empty folder only.

2. Click OK.

4.1.2.3 Attaching a managed vault

A vault managed by a storage node can be attached to another storage node. You might need to do so when retiring storage node hardware, when the storage node is lost or when balancing loads between storage nodes. As a result, the first node stops managing the vault. The second node scans archives in the vault, creates and fills up the database corresponding to the vault, and starts managing the vault.

When deleting a managed vault, you have the option to retain archives contained in the vault. The location resulting from such deletion can also be attached to the same or another storage node.

Personal or centralized unmanaged vaults cannot be attached.

To attach a managed vault to a storage node, perform the following steps.

Vault

Storage node

Select the Acronis Backup & Recovery 10 Storage Node that will manage the vault.

Path

Specify the path to the location where the archives are stored.

Database path

Specify a local folder on the storage server to create a vault-specific database. This database will store the metadata required for cataloguing the archives and performing deduplication.

Password

For the vault that was encrypted, provide the encryption password.

After you have performed all the required steps, click **OK** to commit to attaching the vault. This procedure may last for quite a while since the storage node has to scan the archives, write the metadata in the database, and deduplicate the archives if the vault was originally deduplicating.

4.1.3 Tape libraries

This section describes in detail how to use robotic tape devices as vaults for storing backup archives.

A tape library (robotic library) is a high-capacity storage device that contains the following:

- one or more tape drives
- multiple (up to several thousand) slots to hold tape cartridges
- one or more loaders (robotic mechanisms) intended for relocating the tape cartridges between the slots and the tape drives
- barcode readers (optional).

4.1.3.1 Overview

Acronis Backup & Recovery 10 provides full support of a tape library through Acronis Backup & Recovery 10 Storage Node. The storage node should be installed on the machine a tape library is attached to. Storage node can simultaneously use more than one tape library for keeping archives.

To manage a tape library media, the storage node uses the Windows Removable Storage Manager (RSM). See the RSM Media Pools (p. 145) section for more information.

A dedicated database of the storage node keeps information of the backup content written onto the tapes. So some operations (for example, Cleanup (p. 412)) can be performed quite fast without accessing the media. It is possible to view the content of a backup archive located on a tape through the console, even if a tape library is turned off, due to content information stored in the database. To create an incremental or differential backup of data, the program uses the database instead of loading, mounting, rewinding and reading a tape with the full data backup. However, a tape should be read, for example, to validate (p. 421) a backup or to recover data from a backup.

A tape library can be locally attached to a machine the agent is installed on, but only in the case the library is considered as a single tape drive. The agent can use such device to write and read data backups, but the backup's format differs from the format of the backups on the tapes written through the storage node. To get information about the readability of the archives on tapes, written

by different components of other versions of the product by means of Acronis Backup & Recovery 10, see the Tape compatibility table (p. 54) section.

Acronis Backup & Recovery 10 enables you to set up distribution of backups by media. For example, a separate tape set can be used to back up some specific data, and the backups of all other data will be written onto any currently mounted tape, which does not belong to the tape set. See the Tape support (p. 122) section for more information.

The backup schemes (Grandfather-Father-Son (p. 37), Tower of Hanoi (p. 40)) considerably assist you with creating effective schedule and retention rules for backups on a tape library. In combination with the tape options, the backup schemes enable you to reuse, in automatic mode, the tapes that are considered as free after backup deletion. See the Tape rotation (p. 151) section for more information.

4.1.3.2 Hardware

A tape library (robotic library) is a high-capacity storage device that contains the following:

- one or more tape drives
- multiple (up to several thousand) slots to hold tape cartridges
- one or more loaders (robotic mechanisms) intended for relocating the tape cartridges between the slots and the tape drives
- barcode readers (optional).

Each tape may have a special label attached to the side of a cartridge and comprise of:

- a barcode to scan by a special reader that is usually mounted on a loader
- a readable barcode digital value.

Such labels are used for tape identification in a tape library or especially in off-site storage.

If all cartridges in a tape library have barcodes, the library is ready to be automatically managed by software.

Tape libraries are a cost-effective solution for data storages with huge capacity. Moreover, tape is perfect for archiving because cartridges can be stored off-site for enhanced data security. However reading even a small amount of data from a tape library takes much more time (from several seconds to several minutes) than from other types of data storages. The best practice of tape usage is "LESS requests to write/read LARGER amount of data". So systematic access to very large quantities of data is more suitable for a tape library than random access to small portions of data.

4.1.3.3 Limitations

Limitations of tape library usage are the following:

- 1. The consolidation (p. 413) operation is not possible for archives located on tapes. Deletion of a single separate backup is impossible from a tape. It is possible to delete all the backups stored on a tape. However, after this operation all the incremental and differential backups, stored on other tapes and based on the deleted backups, cannot be used for data recovery. In a Custom backup plan's retention rules the If deletion of a backup affects other backups > Consolidate the backup option is disabled. Only the Postpone the deletion option is available.
- 2. Deduplication (p. 413) is not available for archives located on tape storage devices.
- 3. File recovery from a disk backup stored on tape is possible, but can take a very long time.

- 4. A tape with backups written by the storage node cannot be read on a tape device, locally attached to a machine, the agent is installed on, because of a difference in tape format. To get information about the readability of the archives on tapes, written by different components of other versions of the product by means of Acronis Backup & Recovery 10, see the Tape compatibility table (p. 54) section.
- 5. Barcode printers are not used.

4.1.3.4 RSM Media Pools

Acronis Backup & Recovery 10 uses Windows Removable Storage Manager (RSM) to manage tape cartridges belonging to tape libraries.

To separate access to media by different programs the RSM uses so called Media Pools that are logical media groups. There are two categories of media pools in the manager: **System** and **Application**.

System media pools include **Free** pool, **Import** pool and **Unrecognized** pool. The **System** pools hold media that are not currently used by applications. The **Free** pool holds media that are considered as free and can be used by applications. The **Import** and **Unrecognized** pools are temporary pools for media that are new in certain library.

Through RSM an application can get its own pools with proper names, move media from the **Free** pool into its own pools, use its own pools' media for correct purpose, return media to the **Free** pool, etc.

Acronis Backup & Recovery 10 Storage Node manages the tapes belonging to the **Acronis** pool.

If you fill tape library slots with unused tapes, all the tapes will be included into the **Free** pool automatically.

If a tape was used previously, the RSM tries to detect the registered application the tape is concerned to. If the application is not found, the RSM will move the tape into the **Unrecognized** pool. If the application is not found, but the RSM database has no information about the tape, it will be moved into the **Import** pool. If the RSM database has the information, the tape moves into its own pool of the application.

Acronis Backup & Recovery 10 Storage Node provides the RSM to detect the tapes written by Acronis True Image Echo, Acronis True Image 9.1 product families and by components of Acronis Backup & Recovery 10. The storage node will locate all tapes written in "Acronis" format into the **Acronis** pool at the Inventory (p. 149) operation.

Acronis Backup & Recovery 10 components don't use the **Unrecognized** pool. To utilize a tape from this pool forcibly, move the tape to the **Free** pool using the Removable storage snap-in (**Control panel > Administrative tools > Computer management > Removable storage > Media pools**).

If a tape has moved into the **Free** pool, it is considered as free and will be accessible to write by any application. So the tape data will be lost.

If all the backups are deleted from a tape, it will not return to the **Free** pool. It remains in the **Acronis** pool as a free tape to be reused. So if a storage node needs a new tape, it finds a free tape first in the **Acronis** pool, then in the **Free** pool.

Thereinafter Acronis Backup & Recovery 10 Storage Node deals only with the tapes belonging to the **Acronis** pool.

4.1.3.5 Getting started with a tape library

If you have a tape library device attached to a machine with Acronis Backup & Recovery 10 Storage Node installed, all you need to do to back up onto the tape library is to create an archive vault on the device under storage node management.

Prerequisites

A tape library device has to be installed on a machine running Windows in accordance with the device manufacturer's installation instructions.

If Removable Storage Manager (RSM) is present in your version of Windows, it must be activated.

In Microsoft Windows XP and Microsoft Windows Server 2003:

Removable Storage Manager is part of the operating system and is activated initially.

To activate Removable Storage Manager in Microsoft Windows Server 2008:

- 1. Click Administrative Tools > Server Manager > Features > Add Feature.
- 2. Select the Removable Storage Manager check box.

To activate Removable Storage Manager in Microsoft Windows Vista:

- 1. Click Control Panel > Programs > Programs and Features > Turn Windows features on or off.
- 2. Select the **Removable Storage Management** check box.

Fill the library slots with tape cartridges. If a tape does not get a barcode or its barcode is corrupted, you can define the tape label for identification purposes later.

You should have Acronis Backup & Recovery 10 Management Server and Acronis Backup & Recovery 10 Management Console installed on local or remote machines, as well as Acronis Backup & Recovery 10 Storage Node, installed on the machine with the tape library device, and registered in the management server.

Tape library as a managed vault

To enable data protection operations using a tape library you have to create a managed vault on the tape library. You can create a vault from the **Centralized vaults** view of the console. See the Creating a managed centralized vault (p. 139) section for more information.

But the simplest way is to create a vault from the **Storage Nodes** view. In addition to that select the storage node, the tape library is attached to, and then click **Create vault**. The **Create centralized vault** page will be displayed with the pre-selected parameters. All you need to do is to specify the vault **Name** before you click **OK**.

Once the vault is created, it is accessible from the **Centralized vaults** view of the console. Then the tape library can be used for backing up.

Acronis Backup & Recovery 10 allows creating only one vault per tape device.

If all cartridges in a tape library have barcodes, and the RSM **Free** pool contains enough tapes for a chosen backup scheme, the library is ready to fully work automatically.

You can start working with the vault even though all the tape library slots are empty. If there are no available tapes in the tape library slots during the backup operation, the **Tasks Need Interaction** window asks you to load a tape.

If the tape barcode cannot be read, another **Tasks Need Interaction** window asks you to label a tape.

Actions on a tape library vault

If a tape library vault is selected on the **Navigation** pane of the console, the **Centralized vaults** page toolbar will contain the following two actions that are used for tape libraries only:

- Manage tapes displays the Tape Management window allowing you to refresh information on the library slots, inventory tapes in the slots, and define labels for the tapes. If you have a new label assigned to the tape, the action enables you to eject the tape temporarily to make the same label outside the cartridge.
- Rescan tapes displays the Tape Rescanning window, which is useful for selecting slots and launching the Rescan (p. 149) procedure to read some special information on the content of the specified tapes.

Also the Edit, Delete, Validate, and Refresh functions are allowed on a tape library vault.

It should be noted, these functions have some specific features for a tape library. So the **Edit** operation enables you to substitute a tape library device without the **Rescan** operation. The **Delete** operation clears all the information on the selected tape library vault from the storage node database, i.e. the operation deletes the content data of all the tapes, when ever the data is used by the storage node on the tape library device.

At the **Delete** operation, the vault content will be deleted from the storage node database without accessing the tapes. The plans and tasks that use this vault will fail.

The backup archives, belonging to a deleting centralized vault on a tape library, will be deleted as well, but these archives might be recovered by any storage node through the **Rescan** operation.

Actions with archives on tapes in a library

The following are common functions for archive data management for a backup archive selected in the **Centralized vaults** view of the console, when the current vault is a tape library: **Validate**, **Delete**, **Delete all archives**. Deletion in the storage node database is performed without access to tapes. A backup archive deleted from a tape library vault can be restored after the deletion by the Rescan (p. 149) operation, which is performed for all the tapes keeping the archive's data.

The **Rescan** operation for a tape, where a backup was deleted from, can recover the backup, as it recreates information on the content of the backup in the storage node database.

If all the backups are deleted from a tape, it is considered as free. So the deleted backups will be irrevocably lost after the first writing to the tape.

Backing up to tape library

At creating a backup policy/plan with a tape library destination, you set up the backing up in the same way as with other storage devices. The only difference is the additional Tape support (p. 122) options that can be set up during the backup policy/plan creation. These options enable you to specify how the created backup policy/plan should use tapes from the tape library, however the options' presets increase usage efficiency of both whole tape library and each tape.

To view and change the tape options, select **Options > Default backup and recovery options > Default backup options > Tape support** from the top menu.

To change the settings of the backup policy/plan to be created click **Change...** in the **Backup options** section on the **Create backup policy/plan** page. It opens the **Backup options** window where the **Tape support** page is contained with the pre-defined values.

When backing up to a tape and the end of the tape is reached, a free tape will be mounted automatically and the operation will continue onto the new tape.

While a backup task is running, the following tape-specific information is accessible from the console:

- number of tapes currently used by the backup operation
- labels of the tapes used by the task up to the current time in case of backup splitting
- label of the tape that is currently written.

Recovering from tape library

Data recovery from archives located on tape devices is performed in the same way as with other storage devices.

When recovering, you start creating a recovery task, select the tape device vault, and select the archive and the backup to recover data from. At task creation, the program uses the storage node database instead of accessing tapes. However, selection of data to recover (e.g. some files or specific volumes) requires reading of one or more tapes, so it might be durational.

The program finds the tapes and inserts them automatically in the right order. The **Task Need Interaction** window comes up if a required tape is not found.

Keep in mind that a data recovery operation may require access to a number of tapes. For example, data recovery from an incremental backup commonly might require loading, mounting, rewinding and reading of the following tapes containing the data backups:

- tapes storing the incremental backup selected to recover the data
- tapes storing the last full backup created before the selected incremental one
- tapes storing the last differential backup created after the last full backup but before the selected incremental one if necessary
- tapes containing all incremental backups created after the last full or differential backups before the selected incremental one if necessary.

While a recovery task is running, the following tape-specific information is accessible from the management console:

- labels of all the tapes that may be required for the operation
- label of the tape that is currently being read
- labels of the tapes that have already been read
- labels of tapes that are still waiting to be read with information of their current availability (loaded or not).

4.1.3.6 Managing a tape library

To manage a tape library the following tasks/procedures are in the product:

- Inventory (p. 149)
- Rescan (p. 149)
- Labeling (p. 150)

Any user with access to a managed vault on a tape library is able to perform these operations. However two or more users cannot manage a tape library drive simultaneously, because some operations can take minutes, hours or even days. For example, if a user launches a tape library **Rescan** task, all other users' requests to perform the same task will be canceled automatically, as it is already running on the vault.

Inventory

A storage node needs information about a tape in its own database to be able to operate with the tape. So after the vault is created, generally the next step is to inventory tapes.

Inventorying is a procedure that allows the storage node recognize tapes that are currently loaded into the tape library slots. It is relatively fast and normally requires reading the cartridge barcodes without reading the tape data. If a barcode cannot be read, the tape will be mounted to read its GUID identifier only.

The **Inventory** procedure can be run manually by a user or automatically, when access to recently added tapes is required.

To launch the procedure select the tape library vault in the **Navigation** pane of the console, click **Manage tapes** and then click **Start inventory** on the **Tape Management** window.

When inventorying is completed a user has the list of tapes currently loaded into the library.

Perform the procedure every time you load new tapes into tape library slots.

Rescan

As stated above the storage node keeps information about tapes and their contents in a dedicated database. The **Rescan** task reads information about the content of user-selected tapes and updates the database.

The task can take a long-time so it is only initiated manually. You should select each slot with a tape you want to rescan before the task launch.

Run the Rescan task:

- for tapes that are unknown for the storage node
- if the storage node database is lost or damaged
- for tapes whose content is out of date (for example, a tape content was modified through another storage node or manually).

Bear in mind, a tape might keep some backups that were deleted before the tape rescanning. So after the task is completed, all such backups will be recovered in the storage node database and become accessible for data recovery.

At rescanning a tape label should be saved in the storage node database. If a slot, selected for the procedure, contains a tape that still does not have a label, the **Rescan** task for the tape is paused to perform the Labeling (p. 150) procedure.

Labeling

When a tape required for data recovery is not found, the **Task Need Interaction** window will ask the user to bring the tape and insert it into a tape library slot. So, all the tape cartridges need a barcode or other readable labels.

If a tape does not get a label, you should define it before the tape will be used.

If you need to apply a specific label for a tape (for example, "MyWork" label for a tape dedicated to back up files from the folder C:\work) instead of a barcode label, use the **Labeling** procedure as well.

To launch the procedure, select the tape library vault in the **Navigation** pane of the console, and click **Manage tapes** on the toolbar. Then the **Tape Management** window will show a list of the library slots that contain tapes. For every tape belonging to the **Free** pool or to the **Acronis** pool, the slot data field indicates the tape label. Labels are also displayed for tapes that are in the **Imported** pool and contain backups written by Acronis (such might be the case when you bring a tape from another tape library).

By default, an unused tape with a barcode gets a label that is equal to the barcode. If a barcode is absent or corrupted, the label name will be created automatically. You can accept proposed labels or provide your own label as a plain text.

Tapes from the **Free** or the **Imported** pool can be renamed on condition that the user account used to run the storage node service (**ASN User**) has write permissions for these pools. These permissions are not assigned to **ASN User** during installation, so you might need to add them manually.

To define your own label for a tape, select a related data field, type in a new label, click **Eject tape**, write the same label on the tape cartridge (to make association with the label) and insert it back into the same slot.

Once all the required tape labels are specified press **Set labels** to store labels in the storage node database.

4.1.3.7 Tape support

These options are effective when the backup destination is a managed vault located on a tape library.

Tape support options enable you to specify how the backup tasks will distribute backups among the tapes.

Some combinations of tape options might degrade usage efficiency of both the whole tape library and each tape. If you are not forced to modify these options by some specific needs, leave them unchanged.

An archive can occupy several tapes. In such cases a so-called **tape set** is used for keeping the data backups.

Tape set is a logical group of one or more tapes which contain backups of the specific protected data. A tape set can contain backups of other data as well.

Separate tape set is a tape set which contains only backups of the specific protected data. Other backups cannot be written to a separate tape set.

(For the backup policy/plan to be created) Use a separate tape set

The preset is: **Disabled**.

If you leave this option unchanged, then the backups, belonging to the policy or plan being created, might be written onto tapes containing backups written by different backup policies and comprising of data from different machines. Similarly, backups from other policies might be written onto the tapes containing this policy's backups. You will not have a problem with such tapes, as the program manages all the tapes automatically.

When this option is enabled, the backups, belonging to the policy or plan being created, will be located on a separate tape set. Other backups will not be written to this tape set.

If the console is connected to the management server

The **Use a separate tape set** option has more precise definitions. So for the backup policy to be created you can use a separate tape set for all machines or for each single machine.

The **A single tape set for all machines** option is selected by default. Generally this option ensures more efficient usage of tapes, than the **A separate tape set for each single machine** option. However the second one can be useful, for example, when there are special requirements to store the tapes with backups from a specific machine off-site.

When the **Use a separate tape set** option is enabled, there might be a case when the backup has to be written onto a tape that is currently out of the tape library device. Define what to do in this case.

- Ask for user interaction the backup task will enter the **Need Interaction** state and wait for the tape, with the required label, to be loaded into the tape library device.
- Use a free tape the backup will be written onto a free tape, so the operation will be paused only if there is no free tape in the library.

Always use a free tape

If you leave the options below unchanged, then each backup will be written onto the tape specified by the **Use a separate tape set** option. With some of the options below enabled, the program will add new tapes to the tape set every time when a full, incremental or differential backup is created.

For each full backup

The preset is: **Disabled**.

When this option is enabled, each full backup will be written onto a free tape. The tape will be loaded to a drive especially for this operation. If the **Use a separate tape set** option is enabled, only incremental and differential backups of the same data will be appended to the tape.

For each differential backup

The preset is: **Disabled**.

When this option is enabled, each differential backup will be written onto a free tape. This option is available only when using free tape for each full backup is selected.

For each incremental backup

The preset is: **Disabled**.

When this option is enabled, each incremental backup will be written onto a free tape. This option is available only when using free tape for each full and differential backup is selected.

4.1.3.8 Tape rotation

If all backups are deleted from a tape, i.e. if information about the last backup on the tape is deleted from the storage node database, the tape is considered as empty and can be reused during a backup

cycle. The same tape rotation enables you to get by with the minimum number of cartridges and not to be buried in used tapes.

Acronis Backup & Recovery 10 enables you to achieve full automation of tape rotation while backing up onto tape libraries.

This section provides you with useful information to choose a backup scheme and tape options for tape rotation.

To calculate the number of tapes required for tape rotation schemes, you can use the method described in the Tape planning (p. 163) section.

Choosing a backup scheme

When creating a backup policy/plan with a tape library destination, the following backup schemes are available: **Back up now**, **Back up later**, **Grandfather-Father-Son**, **Tower of Hanoi**, or **Custom**. The **Simple** backup scheme is disabled, because backup consolidation is impossible for archives located on tapes.

Acronis Backup & Recovery 10 provides automation of tape rotation for **Grandfather-Father-Son**, **Tower of Hanoi**, and **Custom** backup schemes.

Grandfather-Father-Son (p. 37) (GFS) and Tower of Hanoi (p. 40) (ToH) are the most popular backup schemes to use on tape library devices. These schemes are optimized to maintain the best balance between a backup archive size, the number of recovery points available from the archive, and the quantity of required tapes for archiving.

If your backup archive must provide recovery with daily resolution for the last several days, weekly resolution for the last several weeks and monthly resolution for any time in the past, the most preferred scheme for you is the **Grandfather-Father-Son** scheme.

If the main goal is to provide data protection for the longest period with the minimal number of used tapes permanently loaded into a small tape library (e.g. autoloader), the best solution is to probably choose the **Tower of Hanoi** scheme.

The **Custom** backup scheme enables you to specify a backup schedule and retention rules to define a desired tape rotation. Use this scheme, when the **Grandfather-Father-Son** and the **Tower of Hanoi** schemes' usage is not enough. For example, if the full size of protected data is considerably less than the size of a tape, the best choice is to use the **Custom** backup scheme with regular daily/weekly/monthly full backups, some simple retention rules, and tape options by default.

Criteria of the choice

Every time you are about to design a tape rotation scheme for a backup policy/plan to be created, you ought to come from the following arguments:

- full size of the data to protect
- approximate size of the daily changes of data
- approximate size of the weekly changes of data
- requirements for the backup scheme (frequency, performance and duration of backup operations)
- requirements for keeping backups (minimal/maximal period of backup keeping; need to store tape cartridges off-site)

- capability of the tape library (number of drives, loaders, slots and available tapes; capacity of tapes)
- requirements for performing data recovery (maximal duration)

You need to analyze every argument that is relevant for your case and select the main criteria for the choice. Then choose a backup scheme and specify the tape options.

Note, that any backup scheme in combination with different tape options will have quite different results for efficient use of both tapes and devices.

Case to analyze

Suppose you need to automate a tape rotation for the case if:

- the full size of the data to protect is approximately 320 GB
- the approximate size of daily changes of data is about 16 GB
- the approximate size of weekly changes of data is no more than 40 GB
- tape capacity is 400 GB.

Let's analyze the results of a combination of GFS and ToH schemes with different tape options for the case.

All the below analyzed examples are a simplistic approach to a real case, but provide you with a general conception of backup distribution onto tapes.

Legend for the case example figures

Any daily/incremental backup (16 GB) is shown in the figures as a green rectangle: ____.

Weekly/differential backups (40 GB) are displayed as a blue rectangle:

Any full monthly backup (320 GB) is drawn in orange:

A whole tape (400 GB) is drawn as a gray rectangle:

Using the Grandfather-Father-Son tape rotation scheme

Tape rotation for the GFS backup scheme is substantially defined by the tape options specified for the backup policy/plan to be created.

Assume the GFS settings are the following:

Start backup at: 11:00:00 PM

Back up on: Workdays

Weekly/Monthly: Friday

Keep backups: Daily: 2 weeks; Weekly: 2 months; Monthly: 1 year.

The main goal is to achieve full automation of tape rotation for these settings.

Keep in mind that a monthly backup is full, a weekly backup is differential, and a daily backup is incremental in this implementation of the GFS scheme. The first backup is always full. So if the

backup policy/plan starts on Wednesday and full backups should be created on every fourth Friday, on Wednesday the first backup will be full instead of an incremental one.

There are analyzed examples showing how the GFS scheme can be combined with different tape options in the following sections:

- GFS Example 1 (p. 154). The Use a separate tape set option is selected. All the Always use a free tape options are cleared. It requires 25 tapes in rotation.
- GFS Example 2 (p. 157). The Use a separate tape set option is selected. The Always use a free tape: For each full backup option is selected. Other Always use a free tape options are cleared. It requires 16 tapes in rotation.
- GFS Example 3 (p. 159). The Use a separate tape set option is selected. All the Always use a free tape options are selected. It requires 28 tapes in rotation.

These examples demonstrate how the number of tapes required for automated rotation depends on the tape options. If a tape library does not have enough tapes for automated rotation, the Tasks **Need Interaction** window will sometimes ask you to load a free tape into the library.

GFS Example 1

Suppose, the backup plan has the following tape options:

- the Use a separate tape set option is selected
- the Always use a free tape: For each full backup option is cleared
- the Always use a free tape: For each incremental backup option is cleared
- the **Always use a free tape: For each differential backup** option is cleared.

Imagine the first backup operation is scheduled on Friday 1st of January. On that day at 11:00 PM the first full backup (320 Gb on the tape whose size is 400 Gb) is created. As the Use a separate tape set option is selected, the currently mounted tape is ejected (if it is not a free tape). Then a free tape is loaded especially for backing up the data. The tape is marked with number 01 in the figure below. In accordance with the legend described in the Case to analyze (p. 153) section, the full data backup is displayed as an orange rectangle in the figure.

The specified GFS backup scheme settings force the data to be backed up on Workdays only, so the next backup is created at the same time (11:00 PM) on Monday 4th of January. This backup is an incremental one (16 Gb) that is written onto the same tape 01, because the Always use a free tape: For each incremental backup option is cleared. The backup is drawn as a green rectangle in the figure.



The next three incremental backups are written onto tape 01 on 5th, 6th and 7th of January. As a result the free space on the tape is only 16 Gb at the moment.

On 8th of January the data differential backup (40 Gb) is recorded onto the same tape 01, as the Always use a free tape: For each differential backup option is cleared. However the tape reached the end after the first 16 Gb of the backup is written. Then the tape is dismounted and ejected from the drive into a slot by the loader. Further, a free tape is loaded into the same drive and mounted, and then the backup (last 24 Gb) is continued onto the beginning of the new tape.

The next figure demonstrates the data backup archive at the moment. The differential backup is drawn as a blue rectangle in the figure. Number 1 in the green rectangle marks the incremental backup created on Monday of the 1st week of the year.



Then the following backups are written onto tape 02:

- four incremental and one differential backup on the second week
- four incremental and one differential backup on the third week
- four incremental backups on the 4th week.

The next full backup (320 Gb) should be written on Friday of the 4th week. However tape 02 has only 104 Gb of free space at the moment. So after the tape reaches the end, the recording continues from the beginning of free tape 03.




Keep in mind, that the **Cleanup** task is launched after each backup operation for the GFS scheme. This task deletes all the outdated backups. The next figure shows dark-gray rectangles instead of the backups deleted up to the current time.




Physically the deleted backups are still on the tapes; however information about the backups is deleted from the storage node database.

Below, the figure shows the deleted backups as actual, but demonstrates tape usage during the whole year for the GFS backup scheme in combination with the specified tape options. A number in the green rectangle marks an incremental backup created on Monday of the corresponding week of the year.



Tape usage during the first year

The next figure shows the actual usage of the tapes with free space instead of the deleted backups on the first Friday of the following year. At the time the differential backup (blue rectangle) is written onto tape 24.



The full backup stored on tape 01 is deleted after the next full backup is created onto both tapes 23 and 24 on Friday of the 52nd week. As all backups of tape 01 have been deleted, the tape is considered as free and can be reused.

Further analysis of the example proves that the maximal number of tapes required to store the data backups is 25 tapes. This maximum occurs on the 16th week of the following year.

The above mentioned figures show that a data recovery requires one or two tapes for a full backup, two or three tapes for a differential backup, and one, two or three tapes for an incremental backup.

For example, if we need to recover data from a backup created on Monday of the 52nd week, the task will require the following tapes:

- Tape 23 with an incremental backup (marked with "52") and a differential backup created on Friday of the 51st week
- Tape 21 and Tape 22 that contain a full backup created on Friday of the 48th week.

The example reveals the following shortcomings of the scheme combination with the specified tape options:

- commonly any data recovery is a long process that requires loading, mounting, rewinding and reading of one (3% - for backups displayed in the "Tape usage during the first year" figure), two (65%) or three (32%) tapes
- 22 tapes are used to store 13 monthly full backups when the monthly backup size is less than the size of a tape, so keeping data is more expensive
- 25 tapes are required for full year rotation of the data backups.

GFS Example 2


Suppose, the backup plan has the following tape options:

- the Use a separate tape set option is selected
- the Always use a free tape: For each full backup option is selected
- the Always use a free tape: For each incremental backup option is cleared
- the Always use a free tape: For each differential backup option is cleared.


The example has only one difference from the previous one. That is selection of the **Always use a free tape: For each full backup** option.

Below, the figure shows the deleted backups as actual, but demonstrates tape usage during the whole year for the GFS backup scheme in combination with the specified tape options. A number in the green rectangle marks an incremental backup created on Monday of the corresponding week of the year.

If all the backups have to be kept during the year, the archive will require 28 tapes.



As the GFS backup scheme forces automatic deletion of the outdated backups, on the first Friday of the second year the tapes keep only the backups displayed in the next figure.



This figure demonstrates that the **GFS Example 2** tape rotation scheme is more suitable for the case than **GFS Example 1**. The advantages of the **GFS Example 2** tape rotation scheme for the analyzed case are the following:


- it uses 16 tapes instead of 25
- a data recovery task requires one (25%) or two (75%) tapes
- data recovery from a full backup requires only one tape that makes the data recovery from an incremental or differential backup faster.

GFS Example 3

Imagine the backup plan has the following tape options:

- the Use a separate tape set option is selected
- the Always use a free tape: For each full backup option is selected
- the Always use a free tape: For each incremental backup option is selected
- the Always use a free tape: For each differential backup option is selected.

These options define the tape rotation scheme that is classical for GFS. The figure shows the beginning of the rotation scheme that uses 8 tapes for daily backups, 6 tapes for weekly backups and 13 tapes for monthly backups (since there are 13 four-week cycles in a year) for the analyzed case. And one tape is required for the next backup. In total this rotation scheme, combined with the options requires 28 tapes.



To recover the data only one tape is required for a full backup, two tapes for a differential backup, and two or three tapes for an incremental backup.

This scheme has the following advantages:

access to any full backup requires only one tape

backup deletion frees a tape so it can be reused.

The main drawback is the large number of required tapes that is used 5-10%.

If we have to keep a daily backup for a week (4 backups) and a weekly backup for a month (4 backups), the total number of required tapes will be equal to 4+4+13+1=22.

Using the Tower of Hanoi tape rotation scheme

The ToH scheme requires fewer tapes for rotation as compared with the GFS scheme. So the ToH scheme is the best choice for small tape libraries, especially for autoloaders.

Once the ToH backup scheme is selected, it is possible to specify the scheme schedule and the number of levels.

Best practice suggests that five levels should be used if you are applying the Tower of Hanoi to weekly backups and eight levels if you are applying it to daily backups. In the first case the rotation includes 16 weekly sessions so it ensures that the roll-back period (the minimum number of days that you can go back in the archive) is 8 weeks. Tape rotation for the second case includes 128 daily sessions, i.e. it allows the roll-back period equal 64 days. The roll-back period is always half the number of sessions.

Each additional level doubles not only the number of sessions but also the oldest backup age.

Let's return to the analyzed case described in the Case to analyze (p. 153) section, and suppose the ToH settings are the following:

Schedule: Start the task every 1 day at 11:00 PM. Repeat once.

Number of levels: 5

The Tower of Hanoi scheme with five levels ensures that the roll-back period is 8 days. Let's designate the backups of the levels with numbers from 1 to 5 by letters A, B, C, D, and E respectively. Then the rotation template for the backup sequence in the archive is the following: E-A-B-A-C-A-B-A-D-A-B-A-C-A-B-A. In the five-level ToH scheme all the backups on the 1st level (A) are incremental, on the 5th level (E) – full, and other backups on levels 2, 3, and 4 (B, C, and D) are differential.

Tape rotation for the ToH scheme substantially depends on the tape options, whose default settings do not always provide optimal usage of tapes and the whole tape library.

The goal is to choose the tape options requiring the minimal number of tapes in the rotation.

There are analyzed examples showing how the ToH scheme can be combined with different tape options in the following sections:

- ToH Example 1 (p. 161). The **Use** a separate tape set option is selected. All the **Always use** a free tape options are cleared. It requires 5 tapes in rotation.
- ToH Example 2 (p. 162). The **Use a separate tape set** option is selected. The **Always use a free tape**: For each full backup option is selected. The other **Always use a free tape** options are cleared. It requires 4 tapes in rotation.
- ToH Example 3 (p. 163). The **Use a separate tape set** option is selected. All the **Always use a free tape** options are selected. It requires 7 tapes in rotation.


ToH Example 2 requires 4 tapes, which is the minimum for the case. So its tape options settings are the best in comparison with options for other examples.

ToH Example 1


Suppose, the backup plan has the following tape options:

- the Use a separate tape set option is selected
- the Always use a free tape: For each full backup option is cleared
- the Always use a free tape: For each incremental backup option is cleared
- the Always use a free tape: For each differential backup option is cleared.

The figure below shows the tapes' usage for the ToH scheme combined with the above mentioned tape options. The recurring part of the scheme contains sixteen backup sessions. The figure displays the backup archive state at the moment when the 17th session is finished.




As the Tower of Hanoi backup scheme forces presence of only one backup on each level, all the outdated backups are deleted automatically. In the next figure the deleted backups are drawn as dark-gray rectangles. Actually the deleted backup is still stored on the tapes, but the information about it is deleted from the storage node database.



The figure shows the full backup kept on tape 01 at the moment, which cannot be deleted as it is a base for actual differential (D, C, B) and incremental (A) backups stored on tape 02. The full backup deletion is postponed until all the above mentioned four backups will be deleted.

The next figure demonstrates the tapes' content at the moment before creation of the new backup on level D:



At the moment the data archive occupies four tapes, and the total size of the backups written up to the current time is maximal for the example. However, if in the future a full backup will be written at the end of a tape, the archive will occupy five tapes.

After the next backup is created on level D, tape 01 is freed and can be reused.

It is noticed that the ToH scheme combined with the specified options has the following properties for the analyzed case:

 the last figure shows that the data recovery requires loading and mounting of up to three tapes (one tape - 16%, two tapes - 72%, three tapes - 12%) as well as rewinding and reading of one (6%), two (50%) or three (44%) backups five-level scheme requires up to five tapes for this case.


ToH Example 2

Suppose, the backup plan has the following tape options:

- the Use a separate tape set option is selected
- the Always use a free tape: For each full backup option is selected
- the Always use a free tape: For each incremental backup option is cleared
- the Always use a free tape: For each differential backup option is cleared.

The only difference between **ToH Example 2** and **ToH Example 1** is that the **Always use a free tape: For each full backup** option is selected.

The first figure shows the tapes' usage for the ToH scheme combined with the above mentioned tape options. The recurring part of the scheme contains sixteen backup sessions. The figure displays the backup archive state at the moment when the 17th session is finished.




In the figure below the backups deleted at the moment are drawn as dark-gray rectangles.



The figure indicates that there are two full backups on level E because at the moment the first full backup is a base for differential backups D, C and B are a base for incremental backup A. So the full backup deletion is postponed until all the D, C, B and A backups will be deleted.

The next figure shows the tape usage at the moment before creating a new backup on level D:



At the moment the backup archive occupies four tapes. It is the maximal number of tapes required in the example.

After the next backup on level D is created, both tapes 01 and 02 are freed and can be reused.

It is noticed that the ToH scheme combined with the specified options has the following properties for the analyzed case:

- the data recovery requires access to the backups kept on one (25%) or two tapes (75%)
- five-level scheme can require up to four tapes.


So in this specific case the selection of the **Always use a free tape: For each full backup** option considerably increases the usage efficiency of the tapes in the library.

ToH Example 3

Imagine the backup plan has the following tape options:

- the Use a separate tape set option is selected
- the Always use a free tape: For each full backup option is selected
- the Always use a free tape: For each incremental backup option is selected
- the Always use a free tape: For each differential backup option is selected.

The figure shows tape rotation for the ToH scheme with these options.



Maximal number of tapes used in the rotation is seven that is more than in classical five-level ToH scheme.

Two additional tapes used for:

- 1. keeping an old full backup (postponed deletion) as it is a base for other level backups
- 2. keeping an old backup on a level until a new backup has been successfully created on the level.

The example demonstrates that the tapes' usage efficiency is reduced. Moreover, the data recovery requires access to the backups kept on one (full backups, 6%), two (differential backups, 44%) or three (incremental backups, 50%) tapes. So on average the operation takes more time than in the previous examples.

Tape planning

Once you have specified the backup scheme and tape options, you should determine the minimal number of tapes necessary to achieve full automation of tape rotation.

To simplify the tape planning lets discard the possibility that the calculated tapes might contain backups of other data. It is implied that the **Use a separate tape set** option is enabled.

To calculate the number of tapes you should take into account the following considerations:

- full backup size
- average size of incremental backups
- average size of differential backups
- compression level specified for backing up the data
- tape rotation scheme (frequency of backups, retention rules)
- tape-append options
- requirements to support off-site tape cartridge archives.

There is no common formula to calculate a number of tapes required in all possible combinations of above listed considerations. But the general way to get a number of tapes for a case includes the following steps:

- 1. Draw (or write) a chain of backups until the first backup can be deleted
- 2. Take into account the tape-append options, the chain might be sectioned onto tape sets
- 3. Calculate the number of tapes in each tape set
- 4. The sum of the calculated values gives the total number of tapes required for the case.

Tape planning: Example 1

Imagine the case with the following features:

- full backup size is **F GB**
- average size of incremental backups is I_GB
- average size of differential backups is D_GB
- compression level provides CL average reduction coefficient
- selected tape rotation scheme is Tower of Hanoi with four levels
- tape options are the following:
 - the Use a separate tape set option is selected
 - the Always use a free tape: For each full backup option is cleared
 - the Always use a free tape: For each incremental backup option is cleared
 - the Always use a free tape: For each differential backup option is cleared
- tape size is T_GB.

The Tower of Hanoi scheme with four levels (A, B, C, and D) specifies the following line of backups on the tapes before the first backup will be deleted: D (full), A, B, A, C, A, B, A, D, A, B, A, C. The specified tape options do not require using a free tape for any backup, so the backup line will be automatically split and continued on a new tape when the end of the current tape is reached. There is one tape set to calculate.

Total number of required tapes = round up ((2*F GB + 6*I GB + 5*D GB) * CL/T GB) + 1.

The above described ToH Example 1 (p. 161) is based on the five-level Tower of Hanoi backup scheme with the same tape options. Its backup line was the following: E (full), A, B, A, C, A,

Total number of required tapes = round up $((2*F_GB + 12*I_GB + 11*D_GB) * CL / T_GB) + 1 = round$ up ((2*320 + 12*16 + 11*40) * 1 / 400) + 1 = round up (3.18) + 1 = 5 (tapes).

Tape planning: Example 2

Imagine the case with the following features:

- full backup size is **F_GB**
- average size of incremental backups is I_GB
- average size of differential backups is **D_GB**
- compression level provides CL average reduction coefficient
- selected tape rotation scheme is Custom with the following settings:
 - full backup every 10 days
 - differential backup every 2 days
 - incremental backup every 1 day, every 6 hours
 - retention rules: delete backups older than 5 days
- tape options are the following:
 - the Use a separate tape set option is selected
 - the Always use a free tape: For each full backup option is selected
 - the Always use a free tape: For each incremental backup option is cleared
 - the Always use a free tape: For each differential backup option is cleared
- tape size is **T GB**.

The case defines the line of backups that consists of two sections. The figure below shows the sections at the moment before the first backup will be deleted. In the figure the full, differential and incremental backups are designated as orange, blue and green rectangles respectively.



At the moment some backups are deleted by the Cleanup task. Deletion of the outdated backups painted in dark colors is postponed, as these backups are basic for the actual backups.



As an exact correlation between the tape size and backup's size is unknown, it is impossible to determine number of the tapes that will be free after the deletion. So the calculation ignores this probability.

Tape set 01 should contain (round up (($F_GB + 4*D_GB + 5*7*I_GB$) * CL / T_GB)) tapes to store the backups. Tape set 02 needs (round up (($F_GB + 1*D_GB + 7*I_GB$) * CL / T_GB)) tapes. The sum of the calculated values gives the total number of tapes required for the case.

4.1.3.9 What if

- What if I have to move tapes with backups from one tape library to another?
 - 1. If both tape libraries are attached to the same machine with Acronis Backup & Recovery 10 Storage Node installed (i.e. the libraries are managed by the same storage node), the storage

- node database has all the required information about the content of the moved tapes. So all you need to do is to perform the inventory (p. 149) procedure for the managed vault on the library where the tapes were placed to.
- 2. If you move tapes to a tape library managed by another storage node, you should rescan (p. 149) each relocated tape to provide the storage node with information about backups contained on the tape.
- Acronis agents create backups on tapes in a format that differs from the format used by the storage node. It is the reason why it is impossible to interchange tapes between tape devices attached to a storage node and attached to a managed machine: a tape written by a storage node cannot be read by an agent in a locally attached tape device. However the storage node can read tapes written by an agent. Please refer to the tape compatibility table (p. 54) to get comprehensive information about the compatibility of tape formats in Acronis Backup & Recovery 10.
- What if I have to reinstall the storage node or attach the tape library to another machine?
 Install a storage node on the machine the tape library is attached to, create a centralized vault on the tape library, and then rescan each tape containing backups.
- What if I have lost my storage node and need to recover data from a tape?
 - If you know which tape has the data to recover, and you have a tape device with vault managed by a storage node, insert the tape cartridge into the device, go to the **Centralized vaults** view of the console, select the vault, rescan the tape, select the archive and the backup to recover data from, and create the recovery task.
 - If you don't know which tape has the data to recover, you have to rescan each tape until the data is found. Generally all the steps you need to do are the same as mentioned above, except the rescan has to be applied to a number of tapes instead one tape.
- What if I need to recover data from an Echo tape?

Use the table from the Tape compatibility table (p. 54) section to find out which Acronis Backup & Recovery 10 components can read data from your tape.

4.2 Personal vaults

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

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

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

Sharing a personal vault

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

displays the owner of each archive. To find out more about the owner concept see Owners and credentials (p. 34).

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.

4.2.1 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. 168) 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. 169)

Operations with backups (p. 170)

Filtering and sorting archives (p. 172)

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.

4.2.2 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.

То	Do
Create a personal vault	Click Create.
	The procedure of creating personal vaults is described in-depth in the Creating a personal vault (p. 169) 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	Click A Change user.
for accessing a vault	In the appearing dialog box, provide the credentials required for accessing the vault.
Create Acronis Secure	Click Create Acronis Secure Zone.
Zone	The procedure of creating the Acronis Secure Zone is described in-depth in the Creating Acronis Secure Zone (p. 265) section.
Explore a vault's	Click C Explore.
content	In the appearing Explorer window, examine the selected vault's content.
Validate a vault	Click Validate .
	You will be taken to the Validation (p. 251) page, where this vault is already preselected as a source. The vault validation checks all the archives stored in the vault.
Delete a vault	Click X 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 C 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.

4.2.2.1 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.

4.2.2.2 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. 198) 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. 198) 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.

4.3 Common operations

4.3.1 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.

То	Do
Validate an archive	Click Validate .
	The Validation (p. 251) page will be opened with the pre-selected archive as a source.
	Validation of an archive will check all the archive's backups.
Export an archive	Click Seport.
	The Export (p. 259) page will be opened with the pre-selected archive as a source. The export of an archive creates a duplicate of the archive with all its backups in the location you specify.
Delete a single archive	1. Select the archive or one of the archives you want to delete.
or multiple archives	2. Click X Delete.
	The program duplicates your selection in the Backups deletion (p. 171) 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.
Delete all archives in the vault	Please be aware that if filters have been applied to the vaults list, you see only a part of the vault content. Be sure that the vault does not contain archives you need to retain before starting the operation.
	Click Note: Delete all.
	The program duplicates your selection in the new window that has check boxes for each archive and each backup. Review the selection and correct if need be, then confirm the deletion.

4.3.2 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.

То	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	Right-click the disk backup, then select Recover as virtual machine .
as a 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. 251) 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. 259) 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	Select one of the backups you want to delete, then click X Delete.
multiple backups	The program duplicates your selection in the Backups deletion (p. 171) 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 Polete all .
	The program duplicates your selection in the Backups deletion (p. 171) window that has check boxes for each archive and each backup. Review the selection and correct if need be, then confirm the deletion.

4.3.3 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.

4.3.4 Filtering and sorting archives

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

То	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 and Linux operating systems.

Event	Windows	Linux
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 and Linux operating systems.

Condition: run the task only if	Windows	Linux
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. 120) 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 (p. 182) 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

Every: 3 day(s).
 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

Every: 3 day(s).
 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: <>	Set up how many times the task will be run during the specified time interval.
From: <> Until: <>	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 (p. 182) 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. 225) 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.

- 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. 225) 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.
From: <> Until: <>	For example, setting the task frequency to Every 1 hour From 10:00:00 AM until 10:00:00 PM allows the task to be run 12 times from 10 AM to 10 PM
	during one day.

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

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

Advanced scheduling settings (p. 182) 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. 225) 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. Fffective:

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: SystemEvent source: diskEvent 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 Advanced scheduling settings

The following advanced settings are available when setting up a daily, weekly, or monthly schedule in a backup policy.

Use Wake-On-LAN

When this setting is enabled, Acronis Backup & Recovery 10 Management Server will use the Wake-On-LAN (WOL) functionality to wake up turned-off registered machines when a backup, cleanup or validation is scheduled to start. If the backup task on each machine starts with a delay (see the next setting), the management server will wake up the machines according to those delays.

Before using this setting, make sure that you have enabled Wake-on-LAN on the registered machines. The machine's basic input/output system (BIOS) configuration, network adapter configuration, and the operating system configuration must allow waking up the machine from the powered-off state—also known as the S5 or G2 power state.

Distribute start time within the time window

When this setting is enabled, the backup task on each registered machine will start with a specific delay from the start time set in the policy. This distributes the tasks' actual start times within a time interval.

You may want to use this setting when creating a backup policy for backing up multiple machines to a network location, to avoid excessive network load.

Delay values range from zero to the specified maximum delay value, and are determined according to the chosen distribution method.

The delay value for each machine is determined when the policy is deployed to the machine, and remains the same until you edit the policy and change the maximum delay value.

The conditions, if any, will be checked at the task's actual start time on each machine.

The following examples illustrate this setting.

Example 1

Suppose that you are deploying a backup policy with the following schedule to three machines:

Run the task: **Daily** Once at: **09:00:00 AM**

Distribute start time within the time window

Maximum delay: 1 Hour(s)
Distribution method: Random

Then the task's start time on each machine may be any time between 09:00:00 AM and

09:59:59 AM—for instance:

First machine: Every day at 09:30:03 AM Second machine: Every day at 09:00:00 AM Third machine: Every day at 09:59:59 AM

Example 2

Suppose that you are deploying a backup policy with the following schedule to three machines:

Run the task: Daily

Every: 2 Hour(s) From: 09:00:00 AM Until: 11:00:00 AM

Distribute start time within the time window

Maximum delay: 1 Hour(s)
Distribution method: Random

Then the time of the task's first run on each machine may be any time between 09:00:00 AM and 09:59:59 AM; the interval between the first and the second run is exactly two hours—for instance:

First machine: Every day at 09:30:03 AM and 11:30:03 AM Second machine: Every day at 09:00:00 AM and 11:00:00 AM Third machine: Every day at 09:59:59 AM and 11:59:59 AM

To specify advanced settings

1. Connect to the management server or to a machine registered on it, and then start creating a backup policy or a backup plan.

- 2. In **How to back up**, select the Simple, Tower of Hanoi, or Custom scheme, and then click **Change** to specify a schedule for the scheme.
- 3. Under Run the task, select Daily, Weekly, or Monthly.
- 4. In the **Advanced settings** area, click **Change**.
- 5. To enable the use of the Wake-On-LAN functionality, select the **Use Wake-on-LAN** check box.
- 6. To distribute the centralized backup tasks' start times, select the **Distribute start time within the time window** check box and then specify the maximum delay value and the distribution method.

5.6 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. 225), 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.

5.7 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. 173)" 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. 120) 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. 225) 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.

5.7.1 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.

5.7.2 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.

5.7.3 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.

5.7.4 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.

5.7.5 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
8	Failed tasks: X	Resolve	Resolve will open the Backup plans and Tasks view with failed tasks, where you can examine the reason of failure.
8	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).
8	Failed to check the license for the current edition. X day(s) remaining until the software stops working.	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.
	Please make sure you have a valid license on Acronis License Server.		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.
8	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 <i>X</i> 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.
8	Trial version of product expires in <i>X</i> 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.
8	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.
8	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: Back up the system volume (and the boot volume, if it is different) Create at least one bootable media (p. 411). Create now will launch the Bootable Media Builder (p. 418).
<u> </u>	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.
<u> </u>	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. 190) 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. 188).

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. 190) view and specify your decision on each task in the **Information** panel (or in the **Task details** (p. 198) 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. 191). A backup plan state is a cumulative state of the plan's tasks. The status of a backup plan (p. 191) 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. 192). Check a task status (p. 193) 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. 197) 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. 194) 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. 200) and **Task details** (p. 198) 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:
			 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.	View the log to read the warnings -> [optionally] Perform actions to prevent the future warnings or failure.
		Otherwise, see 3.	
3	ОК	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. 120) 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:
			 Remove the reason of the failure -> [optionally] Start the failed task manually
			Edit the failed task to prevent its future failure
			 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

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	ОК	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.

То	Do	
Create a new backup plan,	Click New, then select one of the following:	
or a task	■ Backup plan (p. 203)	
	■ Recovery task	
	■ Validation task (p. 251)	
View details of a plan/task	Backup plan	
	Click View details. In the Plan Details (p. 200) window, review the plan details.	
	<u>Task</u>	
	Click View details. In the Task Details (p. 198) window, review the task details.	
View plan's/task's log	Backup plan	
	Click View log.	
	You will be taken to the Log (p. 201) view containing the list of the plan-related log entries.	
	<u>Task</u>	
	Click View log.	
	You will be taken to the Log (p. 201) view containing the list of the task-related log entries.	

Run a plan/task

Backup plan

Click Run.

In the Run Backup Plan (p. 198) window, select the task you need to be run.

Running the backup plan starts the selected task of that plan immediately in spite of its schedule and conditions.

Why can't I run the backup plan?

Do not have the appropriate privilege

Without the Administrator privileges on the machine, a user cannot run plans owned by other users.

Task

Click Run.

The task will be executed immediately in spite of its schedule and conditions.

Stop a plan/task

Backup plan

Click Stop.

Stopping the running backup plan stops all its tasks. Thus, all the task operations will be aborted.

Task

Click Stop.

What will happen if I stop the task?

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.

- 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

Backup plan

Click P Edit.

Backup plan editing is performed in the same way as creation (p. 203), except for the following **limitations**:

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).

- 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.

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.

Why can't I edit the backup plan?

- 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.

<u>Task</u>

Click P Edit.

Why can't I edit the task?

■ 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	Backup plan
	Click X Delete.
	What will happen if I delete the backup plan?
	The plan's deletion deletes all its tasks.
	Why can't I delete the backup plan?
	■ 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.
	<u>Task</u>
	Click X Delete.
	Why can't I delete the task?
	■ 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	Click C Refresh.
	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.

Filtering and sorting backup plans and tasks

То	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

- 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. 198) 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 P Fdit.
- 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. 191) of the backup plan.
- Status status (p. 191) 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. 208) 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. 202) 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. 201) 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. 203) 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.

То	Do	
Select a single log entry	Click on it.	
Select multiple log	non-contiguous: hold down CTRL and click the log entries one by one	
entries	 contiguous: 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	1. Select a log entry.
	2. Do one of the following
	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	Select a single log entry or multiple log entries.
entries to a file	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	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	1. Set filters to get a list of the log entries that satisfy the filtering criteria.
entries to a file	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	Click 🖺 Clear Log.
	All the log entries will be deleted from the log, and a new log entry will be created. It will contain information about who deleted the entries and when.

6.1.3.2 Filtering and sorting log entries

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

То	Do	
Display log entries for a	1. In the From field, select the date starting from which to display the log entries.	
given time period	2. In the To field, select the date up to which to display the log entries.	
Filter log entries by type	Press or release the following toolbar buttons:	
	≤ to filter error messages	
	⚠ to filter warning messages	
	🕠 to filter information messages	
Filter log entries by the original backup plan or managed entity type	Under the Backup plan (or Managed entity type) column header, select the backup plan or the type of managed entity from the list.	
Filter log entries by task, managed entity, machine, code, owner	Type the required value (task name, machine name, owner name, etc.) in the field below the respective column header.	
	As a result you will see that the list of log entries fully or just partly coincide with the entered value.	
Sort log entries by date and time	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.	

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:

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. 410), please familiarize yourself with the basic concepts (p. 29) 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. 206)

[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 (p. 206)

[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. 208)

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

Items to backup (p. 209)

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

Access credentials (p. 211)

[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. 211)

[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. 213)

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. 218)

[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. 219)

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. 229)

[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. 103) 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. 229)

[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. 229)

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. 229)

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

Storage (p. 229)

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. 206).

The plan you have created will be accessible for examination and managing in the **Backup plans and tasks** (p. 190) 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. 33) section.

6.2.3 Label (Preserving machine properties in a backup)

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

If you recover the machine to a VMware ESX Server using Agent for ESX/ESXi, or convert (p. 229) the backup to a ESX/ESXi virtual machine, these properties will be transferred to the virtual machine's configuration. You can view them in the virtual machine settings: **Edit settings** > **Options** > **Advanced** > **General** > **Configuration parameters**. You can select, sort and group the virtual machines with the help of these custom parameters. This can be useful in various scenarios.

Example:

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

To add a text label to a backup:

- 1. In the Create backup plan (p. 203) or Create backup policy (p. 367) page, select the Advanced view check box.
- 2. In **Label**, enter the text label or select it from the drop-down menu.

Parameters specification

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

Values of the "acronisTag.os.type" parameter

Windows NT 4	winNTGuest
Windows 2000 Professional	win2000ProGuest
Windows 2000 Server	win2000ServGuest
Windows 2000 Advanced Server	win2000ServGuest
Windows XP All Editions	winXPProGuest
Windows XP All Editions (64 bit)	winXPPro64Guest
Windows Server 2003, All Editions	win Net Standard Guest
Windows Server 2003, All Editions (64 bit)	winNetStandard64Guest
Windows 2008	winLonghornGuest
Windows 2008 (64 bit)	winLonghorn64Guest
Windows Vista	winVistaGuest
Windows Vista (64 bit)	winVista64Guest

Windows 7 windows 7Guest
Windows 7 (64 bit) windows 7_64Guest

Windows Server 2008 R2 (64 bit) windows7Server64Guest

Linux otherLinuxGuest
Linux (64 bit) otherLinux64Guest

Other Operating System otherGuest
Other Operating System (64 bit) otherGuest64

Example

```
acronisTag.label = "DEPT:BUCH; COMP:SUPERSERVER; OWNER:EJONSON" acronisTag.hostname = "superserver.corp.local" acronisTag.os.type = "windows7Server64Guest" acronisTag.os.servicepack = "1" acronisTag.os.sid = "S-1-5-21-874133492-782267321-3928949834"
```

6.2.4 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.

Entire virtual machines

Available if Acronis Backup & Recovery 10 Agent for Hyper-V (or for ESX/ESXi) is installed.

Select this option to back up one or more virtual machines residing on a virtualization server.

Backing up a virtual machine means backing up all the machine's disks plus the machine configuration. With this source type, you can back up multiple machines. This comes in handy when having small (in terms of virtual disks size) but numerous legacy servers such as those resulting from workload consolidation. A separate archive will be created for each machine.

Volumes of a virtual machine

Available if Acronis Backup & Recovery 10 Agent for Hyper-V (or for ESX/ESXi) is installed.

Select this option to back up individual disks or volumes within a virtual machine residing on a virtualization server.

With this source type, you select the machine and then select the disks/volumes to back up. This comes in handy when the operating system and applications, such as a database server, run on a virtual disk, but the data, such as a database, is stored on a large capacity physical disk added to the same machine. You will be able to use different backup strategies for the virtual disk and the physical storage.

6.2.5 Items to back up

The items to backup depend on the source type (p. 208) selected previously.

6.2.5.1 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) (p. 47)".
- 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.

6.2.5.2 Selecting files and folders

To select files and/or folders for backing up

- 1. Expand the local folder 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.

6.2.5.3 Selecting entire virtual machines

Backing up a virtual machine means backing up all the machine's disks plus the machine configuration.

To back up one or more virtual machines residing on a virtualization server

- Select the check boxes next to virtual machines that you want to back up. Selecting the check box for the virtualization server automatically selects all the virtual machines hosted on this server.
 Use the right part of the window to view details of the selected virtual machine or selected virtualization server.
- 2. Click OK.

Backing up an entire virtual machine yields a standard disk backup (p. 414). Having Acronis Backup & Recovery 10 Agent for Windows or for Linux, you can mount its volumes, recover individual files from this backup, and recover disks and volumes from the backup to a physical machine. The virtual machine configuration, stored in a virtual machine backup, will be suggested by default at recovering the backup content to a new virtual machine.

Limitations

A Hyper-V virtual machine that uses at least one pass-through disk (a physical disk, either local or SAN-LUN, attached to the virtual machine) cannot be backed up from the host. To back up such machine or its disks, install Agent for Windows or Agent for Linux on the machine.

A SAN-LUN disk attached to an ESX/ESXi virtual machine in the "Physical compatibility" mode, cannot be backed up from the host while the virtual machine is online (running). To back up such disk, either stop the machine or install Agent for Windows or Agent for Linux on the machine.

6.2.5.4 Selecting a virtual machine's disks and volumes

To back up individual disks or volumes within a virtual machine residing on a virtualization server

- Select the virtual machine whose volumes you need to back up.
 Use the right part of the window to view details on the selected virtual machine.
- 2. Click OK.
- 3. In the **Selecting disks and volumes** (p. 209) window, select the virtual machine disks or volumes. Backing up volumes within a virtual machine is similar to backing up a physical machine's volumes. The virtual machine configuration will be also backed up.

Backing up a virtual machine's volumes yields a standard disk backup (p. 414). Having Acronis Backup & Recovery 10 Agent for Windows or for Linux, you can mount its volumes, recover individual files from this backup, and recover disks and volumes from the backup to a physical machine. The virtual machine configuration, stored in a virtual machine backup, will be suggested by default at recovering the backup content to a new virtual machine.

Limitations

A Hyper-V virtual machine that uses at least one pass-through disk (a physical disk, either local or SAN-LUN, attached to the virtual machine) cannot be backed up from the host. To back up such machine or its disks, install Agent for Windows or Agent for Linux on the machine.

A SAN-LUN disk attached to an ESX/ESXi virtual machine in the "Physical compatibility" mode, cannot be backed up from the host while the virtual machine is online (running). To back up such disk, either stop the machine or install Agent for Windows or Agent for Linux on the machine.

6.2.6 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.

6.2.7 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

6.2.8 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. 401) to the online backup service and activate (p. 402) 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.com/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**.

If the automatically generated name looks like [Virtualization Server Type] [Virtual Machine Name], this means that the name contain variables. Such might be the case when you have selected virtual machines to back up. Virtualization Server Type stands for the virtualization server type (ESX, Hyper-V or other). Virtual Machine Name stands for the virtual machine name. You can add suffixes to the name but never delete the variables, since each virtual machine has to back up to a separate archive with the unique name.

Backing up to an existing archive

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

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

- 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.

6.2.9 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. 216) or all backups created during a week (p. 216). 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. 218) 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 (|).

6.2.9.1 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. 218).

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. 218).

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:

- a) When creating the first backup plan, specify **MyMachine** as the archive name and **D:** as the archive location, select the **Name backup files using the archive name...** check box, select **Full** as the backup type, and schedule the backups to run every week on Monday.
- b) When creating the second backup plan, specify the same settings as in the first backup plan, but select **Incremental** as the backup type and schedule the backups to run every week on Tuesday through Sunday.

Result:

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

Example 6. Backups within working hours

Consider the following scenario:

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

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

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

Run the task: Daily

Every: 1 Hour(s)

From: 08:00:00 AM

Until: 05:01:00 PM

Result:

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

See also "The [Date] variable" (p. 218).

6.2.9.2 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.

6.2.9.3 Backup splitting and simplified file naming

When a backup is split according to backup splitting (p. 117) 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**.

6.2.10 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.

6.2.11 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.

6.2.11.1 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. 35).

6.2.11.2 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.

6.2.11.3 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. 173) section.
Retention rule	With the simple scheme, only one retention rule (p. 43) is available. Set the retention period for the backups.

6.2.11.4 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.

	Мо	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:	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.						
	The default values for each backup type are as follows.						
	Daily: 7 days (recommended minimum)						
	Weekly: 4 weeks						
	Monthly: indefinitely						
	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.						
	We recommend setting a retention period of at least one week for daily backups.						

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 weekWeekly: 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 weekWeekly: 10 daysMonthly: 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 monthsWeekly: 6 monthsMonthly: 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.

6.2.11.5 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. 174), weekly (p. 176), or monthly (p. 178) 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.

I	1	2	3	4	5	6	7	8	9	10	11	12	13
I	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.

6.2.11.6 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	Specifies on what schedule and under which conditions to perform an incremental backup.
	If the archive contains no backups at the time of the task run, a full backup is created instead of the incremental backup.
Differential	Specifies on what schedule and under which conditions to perform a differential backup.
	If the archive contains no full backups at the time of the task run, a full backup is created instead of the differential backup.
Clean up archive	Specifies how to get rid of old backups: either to apply retention rules (p. 43) regularly or clean up the archive during a backup when the destination location runs out of space.
	By default, the retention rules are not specified, which means older backups will not be deleted automatically.
	Using retention rules
	Specify the retention rules and when to apply them.
	This setting is recommended for backup destinations such as shared folders or centralized vaults.
	When there is insufficient space while backing up
	The archive will be cleaned up only during backup and only if there is not enough space to create a new backup. In this case, the program will act as follows:
	 Delete the oldest full backup with all dependent incremental/differential backups
	If there is only one full backup left and a full backup is in progress, then delete the last full backup with all dependent incremental/differential backups
	If there is only one full backup left, and an incremental or differential backup is in progress, an error occurs saying there is a lack of available space
	This setting is recommended when backing up to a USB drive or Acronis Secure Zone. This setting is not applicable to managed vaults.
	This setting enables deletion of the last backup in the archive, in case your storage device cannot accommodate more than one backup. However, you might end up with no backups if the program is not able to create the new backup for some reason.
Apply the rules	Specifies when to apply the retention rules (p. 43).
(only if the retention rules are set)	For example, the cleanup procedure can be set up to run after each backup, and also on schedule.
	This option is available only if you have set at least one retention rule in Retention rules .
	Constitute and the form and the state of the
Cleanup schedule	Specifies a schedule for archive cleanup.
Cleanup schedule (only if On schedule is selected)	For example, the cleanup can be scheduled to start on the last day of each month.

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. 43).

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. 43).

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.

6.2.11.7 Initial seeding

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

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

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

Please refer to the "Initial Seeding FAQ (p. 392)" section for more details.

6.2.12 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. 173) section.

6.2.13 Setting up regular conversion to a virtual machine

When creating a backup plan (p. 203), 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.

6.2.13.1 Setting up a conversion schedule

A disk backup (p. 414) 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.

6.2.13.2 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. 411) or using Acronis Startup Recovery Manager (p. 57). Then, create a recovery task.

Acronis Universal Restore (p. 58) 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. 59) 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. 45) 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. 287) 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. 233)

[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. 233)

Select the archive to recover data from.

Data type (p. 234)

Applies to: disk recovery

Choose the type of data you need to recover from the selected disk backup.

Content (p. 234)

Select the backup and content to be recovered.

Access credentials (p. 235)

[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. 59) for details.

Files (p. 241)

You may have to specify credentials for the destination. Skip this step when operating on a machine booted with bootable media.

Access credentials (p. 243)

[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. 243)

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. 243)

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. 125) 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. 34) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine (p. 33) 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. 240)

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. 240) 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. 239)

Select the destination disk for each of the source disks.

NT signature (p. 237)

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. 240)

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. 240) 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. 239)

Choose the disk to recover the Master Boot Record to.

NT signature: (p. 237)

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. 239)

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:

- 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				
	Winc	dows and Linux				
By name	F.log	Excludes all files named "F.log"				
	F	Excludes all folders named "F"				
By mask (*)	*.log	Excludes all files with the .log extension				
	F*	Excludes all files and folders with names starting with "F (such as folders F, F1 and files F.log, F1.log)				
By mask (?)	F???.log	Excludes all .log files with names consisting of four symbols and starting with "F"				
		Windows				
By file path	Finance\F.log	Excludes files named "F.log" from all folders with the name "Finance"				
By folder path Finance\F\ or Finance\F		Excludes folders named "F" from all folders with the name "Finance"				
		Linux				
By file path	/home/user/Finance/F.log	Excludes the file named "F.log" located in the folder /home/user/Finance				

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

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. 242) 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. 58) 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. 408)

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. 240)** 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. 248)".

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. 281) during recovery if it has been saved (p. 49) 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, ntdetect.com, ntbootdd.sys. boot.ini is a text file that contains the loader configuration. Example:

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Microsoft Windows XP Professional"
/noexecute=optin /fastdetect
```

Windows Vista/2008

A part of the loader resides in the partition boot sector, the rest is in the files bootmgr, boot\bcd. At starting Windows, boot\bcd is mounted to the registry key HKLM \BCD000000000.

6.3.12 Assembling MD devices for recovery (Linux)

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

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

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

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

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

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

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

- If all parts of the MD device are volumes (a typical case, such as in the first example), make sure that each volume type—called partition type or system ID—is **Linux raid automount**; the hexadecimal code of this partition type is 0xFD. This will guarantee that the device will be automatically assembled following the reboot. To view or change the partition type, use a disk partitioning utility such as **fdisk**.
- Otherwise (such as in the second example), perform the recovery from bootable media. No reboot will be required in that case. In bootable media, you may need to create the MD device manually or automatically, as described in Recovering MD devices and logical volumes (p. 281).

6.3.13 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:

 $\label{thm:local_MACHINE} HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session\ Manager\Memory\ Management$

2. Add the **PoolUsageMaximum** entry to the subkey:

■ Entry type: **DWORD Value**

Base: DecimalValue: 40

3. Add the PagedPoolSize entry to the subkey:

■ Entry type: **DWORD Value**

Base: HexadecimalValue: 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.

6.3.14 Recovering the storage node

In addition to backing up data to centralized vaults managed by Acronis Backup & Recovery 10 Storage Node, you may want to perform a disk backup of the machine where the storage node itself is installed.

This section describes how to recover the storage node registered on the management server in case the storage node and the management server are installed on different machines (if they are installed on the same machine, simply recover that machine).

Consider the following scenario:

- You have a machine with the management server and a machine with the storage node.
- The storage node is registered on the management server.
- You backed up the machine with the storage node earlier, and have just recovered it—either on the same machine or on a different machine.

Before using the recovered storage node, follow these steps:

- If you have recovered the storage node on the same machine and no centralized vaults managed by the storage node have been added or removed between the backup and recovery, do nothing.
- Otherwise, do the following:
 - 1. Connect to the management server and remove the storage node from it.

Note: All vaults managed by the storage node will also be removed from the management server. No archives will be lost.

- 2. Add the storage node to the management server again, by specifiying the machine on which the recovered storage node is installed.
- 3. Re-create the necessary managed vaults.

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. 203) section.

You can access the **Validation** page from the **Vaults** (p. 135) 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. 252)

[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. 253) - in that case, you need to specify the archive.

Backup (p. 254) - specify the archive first, and then select the desired backup in this archive.

Vault (p. 254) - select a vault (or other location), which archives to validate.

Access Credentials (p. 254)

[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. 255)

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. 34) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine (p. 33) 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

- 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. 173) 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. 256)

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

Backup (p. 257)

Select the backup.

Access credentials (p. 257)

[Optional] Provide credentials for the archive location. To access this option, select the **Advanced view** check box.

Mount settings

Volumes (p. 258)

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 \bigcirc **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 **Multiple** 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.
- an entire vault that can be exported by using the command line interface. For more information, please refer to Acronis Backup & Recovery 10 Command Line Reference.

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 preselected 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. 261)

[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. 233) - in that case, you need to specify the archive only.

Backups (p. 263) - specify the archive first, and then select the desired backup(s) in this archive

Access credentials (p. 263)

[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. 263)

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. 265)

[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. 34) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine (p. 33) section.

6.7.2 Archive selection

To select an archive

- 1. Enter the full path to the location in the Path field, or select the desired folder in the folders tree.
 - If the archive is stored in 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.

For the management server: in the folders tree, select the managed vault.

- 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. If the archive is password-protected, provide the password.
 - 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.

For the management server the folders tree contains:

- Local folders group to export data onto the hard drives that are local to the storage node.
- Network folders group to export data to a network share. 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.

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. 56) 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. 266)

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. 266)

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. 266)

[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. 267) 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. 266) 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. 421). 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. 135) 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. 411), 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. 269) 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. 271) 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. 276) 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. 276) 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 AlK. To access the documentation, select **Microsoft Windows AlK -> 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. 277) for details.

6.10.1.1 Linux-based bootable media

When using the media builder, you have to specify:

[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. 272).

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. 274):
 - TCP/IP settings to be assigned to the machine network adapters.
- 6. [optional] Network port (p. 275):
 - 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. 275). 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

274

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. 243).

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.
- 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. 274) 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. 270).

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.
- 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. 274) 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. 271) 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 PEbased 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
ср	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. 281).
- Manually by using the mdadm and lvm utilities—see Creating the volume structure manually (p. 282).

6.10.5.1 Creating the volume structure automatically

Let's assume that you saved (p. 49) 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.
- 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 0qoQ41-Vk7W-yDG3-uF11-Q2AL-C0z0-vMeACu
```

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

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

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

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

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

Step 2: Starting the recovery

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

For a complete list of commands and utilities that you can use in the bootable media environment, see List of commands and utilities available in Linux-based bootable media (p. 280). 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:

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

6.10.6 Acronis PXE Server

Acronis PXE Server allows for booting machines to Acronis bootable components through the network.

Network booting:

- eliminates the need to have a technician onsite to install the bootable media into the system that must be booted
- during group operations, reduces the time required for booting multiple machines as compared to using physical bootable media.

Bootable components are uploaded to Acronis PXE Server using Acronis Bootable Media Builder. To upload bootable components, start the Bootable Media Builder (either from the management console, by selecting **Tools > Create bootable media** or as a separate component) and follow the step-by-step instructions described in the "Bootable Media Builder (p. 271)" section.

Booting multiple machines from the Acronis PXE Server makes sense if there is a Dynamic Host Control Protocol (DHCP) server on your network. Then the network interfaces of the booted machines will automatically obtain IP addresses.

6.10.6.1 Acronis PXE Server Installation

To install Acronis PXE Server:


- 1. Run the Acronis Backup & Recovery 10 setup file.
- 2. Select Acronis PXE Server from the list of **Centralized management components**.
- 3. Follow the onscreen instructions.

Acronis PXE Server runs as a service immediately after installation. Later on it will automatically launch at each system restart. You can stop and start Acronis PXE Server in the same way as other Windows services.

6.10.6.2 Setting up a machine to boot from PXE

For bare metal, it is enough that the machine's BIOS supports network booting.

On a machine that has an operating system on the hard disk, the BIOS must be configured so that the network interface card is either the first boot device, or at least prior to the Hard Drive device. The example below shows one of reasonable BIOS configurations. If you don't insert bootable media, the machine will boot from the network.



In some BIOS versions, you have to save changes to BIOS after enabling the network interface card so that the card appears in the list of boot devices.

If the hardware has multiple network interface cards, make sure that the card supported by the BIOS has the network cable plugged in.

6.10.6.3 PXE and DHCP on the same server

If Acronis PXE Server and the DHCP server are on the same machine, add to the DHCP server option 60: "Client Identifier" with string value "PXE Client". This can be done as follows:

```
C:\WINDOWS\system32>netsh
netsh>dhcp
netsh>dhcp>server \\<server_machine_name> or <IP address>
netsh dhcp>add optiondef 60 PXEClient STRING 0 comment="Option added for PXE support"
netsh dhcp>set optionvalue 60 STRING PXEClient
```

6.10.6.4 Work across subnets

To enable the Acronis PXE Server to work in another subnet (across the switch), configure the switch to relay the PXE traffic. The PXE server IP addresses are configured on a per-interface basis using IP helper functionality in the same way as DHCP server addresses. For more information please refer to: http://support.microsoft.com/default.aspx/kb/257579.

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. 287) 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. 411) 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. 302).

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. 289) initializes the new hardware added to the system
- Basic disk cloning (p. 290) transfers complete data from the source basic MBR disk to the target
- Disk conversion: MBR to GPT (p. 292) converts an MBR partition table to GPT
- Disk conversion: GPT to MBR (p. 292) converts a GPT partition table to MBR
- Disk conversion: Basic to Dynamic (p. 293) converts a basic disk to dynamic
- Disk conversion: Dynamic to Basic (p. 293) 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.
- 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. 302) 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. 408) (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. 302) 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:

- 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. 302) 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. 293) 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. 293) 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. 302) 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. 288).

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. 296) Creates a new volume with the help of the Create Volume Wizard.
- Delete Volume (p. 299) Deletes the selected volume.
- Set Active (p. 300) Sets the selected volume Active so that the machine will be able to boot with the OS installed there.
- Change Letter (p. 300) Changes the selected volume letter
- Change Label (p. 301) Changes the selected volume label
- Format Volume (p. 301) 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. 297).

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. 297).

If you click the **Next** button, you will proceed to the next page: Set the volume size (p. 298).

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. 297).

If you click the **Next** button, you will proceed to the next page: Set the volume options (p. 298).

Set the volume options

On the next wizard page you can assign the volume **Letter** (by default - the first free letter of the alphabet) and, optionally, a **Label** (by default – none). Here you will also specify the **File system** and the **Cluster size**.

The wizard will prompt you to choose one of the Windows file systems: FAT16 (disabled, if the volume size has been set at more than 2 GB), FAT32 (disabled, if the volume size has been set at more than 2 TB), NTFS or to leave the volume **Unformatted**.

In setting the cluster size you can choose between any number in the preset amount for each file system. Note, the program suggests the cluster size best suited to the volume with the chosen file system.

If you are creating a basic volume, which can be made into a system volume, this page will be different, giving you the opportunity to select the volume **Type** — **Primary** (**Active Primary**) or **Logical**.

Typically **Primary** is selected to install an operating system to a volume. Select the **Active** (default) value if you want to install an operating system on this volume to boot at machine startup. If the **Primary** button is not selected, the **Active** option will be inactive. If the volume is intended for data storage, select **Logical**.

A Basic disk can contain up to four primary volumes. If they already exist, the disk will have to be converted into dynamic, otherwise or **Active** and **Primary** options will be disabled and you will only be able to select the **Logical** volume type. The warning message will advise you that an OS installed on this volume will not be bootable.

If you use characters when setting a new volume label that are unsupported by the currently installed operation system, you will get the appropriate warning and the **Next** button will be disabled. You will have to change the label to proceed with the creation of the new volume.

If you click the **Back** button, you will be returned to the previous page: Set the volume size (p. 298).

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. 302) 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.
- 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. 302) 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. 302) 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. 302) 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. 302) 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 Centralized management

This section covers operations that can be performed centrally by using the components for centralized management. The content of this section is only applicable to advanced editions of Acronis Backup & Recovery 10.

7.1 Administering Acronis Backup & Recovery 10 Management Server

This section describes the views that are available through the navigation tree of the console connected to the management server, and explains how to work with each view.

7.1.1 Dashboard

Use the Dashboard to estimate at a glance the health of data protection on the registered machines. The Dashboard displays the summary of Acronis Backup & Recovery 10 agents' activities, lets you check for free space available in managed vaults, and rapidly identify and resolve any issues.

Alerts

The alerts section draws your attention to issues that have occurred on the management server and registered machines, in centralized vaults, and offers you ways of fixing or examining them. The most critical issues are displayed at 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
8	Failed tasks: X	View the tasks	View the tasks will open the Backup plans and Tasks view with failed tasks, where you can examine the reason of failure.
8	Tasks that need interaction: X	Resolve	When at least one task existing in the management server's database needs human interaction, the Dashboard shows an alert. Click Resolve to open the Tasks Need Interaction window where you can examine every case and specify your decision.
0	Failed to check licenses on X machine(s)	View log	Acronis Backup & Recovery 10 agent connects to Acronis License Server at the start and then every 1–5 days, as specified by the agent configuration parameters. The alert is displayed if the license check was unsuccessful on at least one agent. This might happen if the license server was unavailable, or the license key data was corrupted. Click View log to find out the cause of the unsuccessful check. If the license check does not succeed for 1-60 days (as specified by the agent configuration parameters), the agent will stop working until a successful license check.

3	Vaults with low free space: X	View vaults	The alert is displayed if at least one centralized vault has less than 10% free space. View vaults will take you to the Centralized vaults (p. 137) view where you can examine the vault size, free space, content and take the necessary steps to increase the free space.
<u> </u>	Bootable media was not created	Create now	To be able to recover an operating system when the machine fails to boot, you must:
			Back up the system volume (and the boot volume, if it is different)
			2. Create at least one bootable media (p. 411).
			Create now will launch the Bootable Media Builder (p. 418).
<u> </u>	No backups have been created for <i>X</i> day(s) on <i>Y</i>	Show list	The Dashboard warns you that no data was backed up on some of the registered machines for a period of time.
	machine(s)		To configure the length of time that is considered critical, select Options > Console options > Time-based alerts .
<u> </u>	Not connected to management server for <i>X</i> day(s): <i>Y</i> machine(s)	View the machines	The Dashboard warns you that no connection was established between some of the registered machines and the management server for a period of time, thus indicating that the machines might not be centrally managed.
			Click View the machines to open the Machines view with the list of machines filtered by the "Last connect" field.
			To configure the length of time that is considered critical, select Options > Console options > Time-based alerts .
<u> </u>	It is recommended to back up the management server to protect its configuration.	Install Acronis components	Install Acronis Backup & Recovery 10 Agent for Windows to back up the machine where the Acronis Backup & Recovery 10 Management Server resides.
	Install the agent on the management server machine and add the machine to AMS.		Click Install now to launch the installation wizard.
<u> </u>	Acronis Backup & Recovery 10 Management Server has not been backed up for X day(s)	Back up now	The alert is displayed only if Acronis Backup & Recovery 10 Agent for Windows is installed on the management server. The alert warns that no data was backed up on the management server for a period of time.
			Back up now will take you to the Create backup plan page where you can instantly configure and run the backup operation.
			To configure the length of time that is considered critical, select Options > Console options > Time-based alerts .

Activities

The stacked column chart lets you explore the daily history of the Acronis Backup & Recovery 10 agents' activities. The history is based on the log entries, collected from the registered machines and from the management server. The chart shows the number of log entries of each type (error, warning, information) for a particular day.

Statistics for the selected date are displayed to the right of the chart. All the statistics fields are interactive, i.e. if you click any field, the **Log** view will be opened with the log entries pre-filtered by this field.

At the top of the chart, you can select the activities to display depending on the presence and severity of the errors.

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

System view

The **System view** section shows summarized statistics of registered machines, tasks, backup policies, and centralized backup plans. Click the items in these sections (except for centralized backup plans) to obtain the relevant information. This will take you to the appropriate view with pre-filtered machines, tasks, or backup policies respectively. For instance, if you click **Idle** under **Tasks**, the **Tasks** view will be opened with tasks filtered by the **Idle** state.

Information presented in the **System view** section is refreshed every time the management server synchronizes with the machines. Information in other sections is refreshed every 10 minutes and every time you access the Dashboard.

Vaults

The **Vaults** section displays information about centralized managed vaults. You can sort vaults by name or by used space. In some cases information about free space in a vault might be not available, for example, if the vault is located on a tape library. If the vault itself is not available (offline), the "Vault is not available" message will be displayed.

7.1.2 Backup policies

To be able to manage and protect multiple machines as a whole, you can create a backup plan template called a "backup policy". By applying this template to a group of machines, you will deploy multiple backup plans with a single action. Backup policies exist only on the Acronis Backup & Recovery 10 Management Server.

You do not have to connect to each machine separately to check whether the data is successfully protected. Instead, check the cumulative status of the policy (p. 306) on all managed machines the policy is applied to.

To find out whether a backup policy is currently being deployed, revoked, or updated, check the deployment state (p. 306) of the policy.

Way of working with the backup policies view

- Use the toolbar's operational buttons to create new policies, apply the existing policies to machines or perform other operations with backup policies (p. 307).
- Use the **Information** pane's tabs to view detailed information about the selected policy and perform additional operations, such as revoke the policy, view details of the machine (group) the policy is applied to, etc. The panel is collapsed by default. To expand the panel, click the chevron. The content of the pane is also duplicated in the Policy details (p. 309) window.
- Use the filtering and sorting (p. 308) capabilities of the policy table for easy browsing and examination.

7.1.2.1 Backup policy deployment states

A backup policy deployment state is a combination of the policy deployment states on all machines the policy is applied to. For example, if the policy is applied to three machines and has the "Deploying" state on the 1st machine, the "Updating" state on the 2nd machine and the "Deployed" state on the 3rd machine, the state of the policy will be "Deploying, Updating, Deployed."

A backup policy deployment state on a group of machines is a combination of the policy deployment states on the machines included in the group.

For detailed information about backup policy deployment states, see the Backup policy's state and statuses (p. 73) section.

7.1.2.2 Backup policy statuses

A backup policy status is the cumulative status of the policy statuses on all machines the policy is applied to. For example, if the policy is applied to three machines and has the "OK" status on the 1st machine, the "Warning" status on the 2nd machine and the "Error" status on the 3rd machine, the status of the policy will be "Error."

A backup policy status on a group of machines is the cumulative status of the policy statuses on the machines included in the group.

The following table shows a summary of possible backup policy statuses.

	Status	How it is determined	How to handle
1	Error	The policy status on at least one machine is "Error".	View the log or identify the failed tasks to find out the reason of the failure, then do one or more of the following:
		Otherwise, see 2.	 Remove the reason of the failure -> [optionally] Start the failed task manually
			Edit the backup policy to prevent future failure
2	Warning	The policy status on at least one machine is "Warning". Otherwise, see 3.	View the log to read the warnings -> [optionally] Perform actions to prevent future warnings or failure.
3	ОК	The policy status on all	No action is required. Note that if a backup policy is not
		machines is "OK".	applied to any machine, its state is also "OK".

What to do if a policy has the Error status

- 1. To find out the reason of the failure, do one or more of the following:
 - Click the Error hyperlink to see the log entry of the latest occurred error.
 - Select the policy and click View tasks. Check the tasks that have Failed as their last result: select a task and then click View log. Select a log entry and then click View details. This approach comes in handy if the policy state is Deployed, that is, the policies' tasks already exist on the managed machines.
 - Select the policy and click View log. Check the "error" log entries to find out the reason of the failure: select a log entry and then click View details. This approach comes in handy if the policy has errors while being deployed, revoked or updated.

In the **Tasks** view, apply the **Last result -> Failed** filter if there are too many tasks. You can also sort the failed tasks by backup plans or by machines.

In the **Log** view, apply the Error of filter if there are too many log entries. You can also sort the "error" entries by backup plans, managed entities or machines.

- 2. Once the reason of the failure is clear, do one or more of the following:
 - Remove the reason of the failure. After that, you may want to start the failed task manually to maintain the backup scheme consistency, for example, if the policy uses the GFS or Tower of Hanoi backup scheme.
 - Edit the backup policy to prevent future failure.

Use the **Activities** section of the Dashboard to quickly access the "error" log entries.

What to do if a policy has the Warning status

- 1. To find out the reason of the warning, do one or more of the following:
 - Click the Warning hyperlink to see the log entry of the latest warning.
 - Select the policy and click View tasks. Check the tasks that have Succeeded with warnings as their last result: select a task and then click View log. This approach comes in handy if the policy state is Deployed, that is, the policies' tasks already exist on the managed machines.
 - Select the policy and click View log. Check the "warning" log entries to find out the reason for the warnings: select a log entry and then click View details. This approach comes in handy if the policy has warnings while being deployed, revoked or updated.

In the **Tasks** view, apply the **Last result -> Succeeded with warnings** filter if there are too many tasks. You can also sort the tasks succeeded with warnings by backup plans or by machines.

In the **Log** view, apply the Warning \triangle filter if there are too many log entries. You can also sort the "warning" entries by backup plans, managed entities or machines.

2. Once the reason of the warning is clear you might want to perform actions to prevent future warnings or failure.

Use the **Activities** section of the Dashboard to quickly access the "warning" log entries.

What to do if a policy status is OK

No action is required.

7.1.2.3 Actions on backup policies

All the operations described below are performed by clicking the corresponding items on the tasks **toolbar**. The operations can also be performed using the context menu (right-click the selected backup policy), or using the 'Backup policy name' actions bar on the Actions and tools pane.

The following is a guideline for you to perform operations with backup policies.

То	Do
Create a backup policy	Click Screate backup policy.
	The procedure of creating a backup policy is described in-depth in the Creating a backup policy (p. 367) section.
Apply policy to machines or groups	Click Apply to. In the Machines selection (p. 308) window, specify the machines (groups) the selected backup policy will be applied to. If the machine is currently offline, the policy will be deployed when the machine comes online again.

Edit a policy	Click PEdit.
	Editing policies is performed in the same way as creating (p. 367). Once the policy is edited, the management server updates the policy on all machines the policy was deployed to.
Delete a policy	Click X Delete.
	As a result, the policy will be revoked from the machines it was deployed to and deleted from the management server. If the machine is currently offline, the policy will be revoked when the machine comes online again.
View details of a policy	Click View details.
or revoke a policy	In the Policy details (p. 309) window, examine information on the selected policy. There, you can also revoke the policy from the machines or groups the policy is applied to.
View tasks of a policy	Click View tasks.
	The Tasks (p. 339) view will display a list of the tasks related to the selected policy.
View log of a policy	Click View log.
	The Log (p. 342) view will display a list of the log entries related to the selected policy.
Refresh a list of policies	Click C Refresh.
	The management console will update the list of backup policies from the management server with the most recent information. Though the list of policies is refreshed automatically based on events, the data may not be retrieved immediately from the management server due to some latency. Manual refresh guarantees that the most recent data is displayed.

Machines selection

To apply the backup policy to machines or to groups of machines

1. Choose whether to apply the selected backup policy to

Groups

In the group tree, select the group(s) the policy will be applied to. The right part of the window lists the machines of the selected group.

Individual machines

In the group tree, select the required group. Then, in the right part of the window, select the machines to apply the backup policy to.

2. Click OK.

The Acronis Backup & Recovery 10 Management Server will deploy the policy to the selected machines and machines belonging to the selected groups.

Filtering and sorting backup policies

The following is a guideline for you to filter and sort backup policies.

То	Do	
		1

Sort backup policies by any column	Click the column's header to sort the backup policies in ascending order. Click it once again to sort the backup policies in descending order.
Filter backup policies by name/owner	Type a policy's name / owner's name in the fields below the corresponding column's header. As a result you will see the list of the backup policies, whose names (or their owners' names) fully or just partly coincide with the entered value.
Filter backup policies by deployment state, status, source type, last result, schedule	In the field below the corresponding column's header, select the required value from the list.

Configuring the backup policies table

By default, the table has seven columns that are displayed, others are hidden. You can adjust presentation of the columns to your needs and preferences.

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.

7.1.2.4 Policy details

The **Policy details** window accumulates in five tabs all information on the selected backup policy and lets you perform operations with the machines and groups of machines the policy is applied to.

This information is also duplicated on the **Information** pane.

Backup policy

The tab displays information about the selected policy.

Source

The tab displays information about the type of source to be backed up and the source selection rules.

Destination

The tab displays information about the backup destination.

Settings

The tab displays information about the backup scheme used by the policy and backup options that were modified against the default settings.

Applied to

The tab displays a list of machines and groups the selected policy is applied to.

Actions

То	Do	
----	----	--

View details of the machine	Click View details.
(group).	In the Machine details (p. 317)/Group details (p. 326) window, examine all information on the selected machine (or the selected group).
View tasks of the machine	Click 😺 View tasks.
(group).	The Tasks (p. 339) view will display a list of the tasks, pre-filtered by the selected machine (group).
View log of the machine	Click View log.
(group)	The Log (p. 342) view will display a list of the log entries, pre-filtered by the selected machine (group).
Revoke policy from the	Click Revoke.
machine (group).	The management server will revoke the policy from the selected machine or group of machines. The policy itself remains on the management server.

7.1.3 Physical machines

Acronis Backup & Recovery 10 lets the administrator protect data and perform management operations on multiple machines. The administrator adds a machine to the management server using the machine's name or IP address, imports machines from Active Directory, or from text files. Once a machine is registered (p. 419) on the management server, it becomes available for grouping, applying backup policies and monitoring the activities related to data protection.

To estimate whether the data is successfully protected on a managed machine, the management server administrator checks its status. A machine's status is defined as the most severe status of all backup plans (p. 191) (both local and centralized) existing on the machine and all backup policies (p. 306) applied to the machine. It can be "OK", "Warnings" or "Errors".

Groups

The management server administrator has the ability to group machines. A machine can be a member of more than one group. One or more nested groups can be created inside any group created by the administrator.

Grouping helps organize data protection by the company departments, by the Active Directory domains or organizational units within a domain, by various populations of users, by the site locations, etc.

The main goal of grouping is protection of multiple machines with one policy. Once a machine appears in a group, the policy applied to the group is applied to the machine and the new tasks are created by the policy on the machine. Once a machine is removed from a group, the policy applied to the group will be revoked from the machine and the tasks created by the policy will be removed.

Built-in group - a group that always exists on a management server. The group cannot be deleted or renamed. A built-in group cannot include nested groups. A backup policy can be applied to a built-in group. The example of a built-in group is the **All physical machines** group, that contains all the machines registered on the management server.

Custom groups - groups created manually by the management server administrator.

■ Static groups

Static groups contain machines manually added to the group by the administrator. A static member remains in the group until the administrator removes the member from the group or deletes the corresponding managed machine from the management server.

Dynamic groups

Dynamic groups contain machines added automatically according to the criteria specified by the administrator. Once the criteria are specified, the management server starts to analyze the existing machines' properties and will analyze every newly registered machine. The machine that meets a certain dynamic criterion will appear in all groups that use this dynamic criterion.

To learn more about grouping machines, see the Grouping the registered machines (p. 66) section.

To learn more about how policies are applied to machines and groups, see the Policies on machines and groups (p. 68) section.

Way of working with machines

- First, add machines to the management server. Adding machines is available, when selecting the **Physical machines** view, or the **All physical machines** group in the **Navigation** tree.
- Select the group the required machine is in, then select the machine.
- Use the toolbar's operational buttons to take actions on the machine (p. 314).
- Use the **Information** panel's tabs to view detailed information about the selected machine and perform additional operations, such as start/stop tasks, revoke policies, explore the policy inheritance, etc. The panel is collapsed by default. To expand the panel, click the chevron. The content of the panel is also duplicated in the **Machine details** (p. 317) window.
- Use filtering and sorting (p. 322) capabilities for easy browsing and examination of the machines in question.

Way of working with groups

- In the Physical machines view, select the group.
- Use the toolbar's operational buttons to perform actions on the selected group (p. 323).
- Use the **Information** panel's tabs to view detailed information about the selected group and perform additional operations, such as revoke policies or explore policy inheritance. The panel is collapsed by default. To expand the panel, click the chevron. The content of the panel is also duplicated in the **Group details** (p. 326) window.

7.1.3.1 Actions on machines

Registering machines on the management server

Once the machine is added or imported to the **All physical machines** group, it becomes registered on the management server. Registered machines are available for deploying backup policies and for performing other centralized management operations. Registration provides a trusted relationship between the agent, residing on the machine, and the management server.

Adding and importing actions are available when you select the Physical machines view or the All physical machines group in the navigation tree.

То	Do
Add a new machine to the	Click Add a machine to AMS.
management server	In the Add machine (p. 314) window, select the machine that needs to be

	added to the management server.
Import machines from Active Directory	Click Import machines from Active Directory. In the Import machines from Active Directory (p. 314) window, specify the machines or organizational units whose machines you need to import to the management server.
Import machines from a text file	Click Import machines from file. In the Import machines from file (p. 316) window, browse for a .txt or .csv file, containing the names (or IP addresses) of machines to import to the management server.

The management console addresses to the agent and initiates the registration procedure. Because registration requires the agent's participation, it cannot take place when the machine is offline.

An additional agent installed on a registered machine becomes registered on the same management server automatically. Multiple agents are jointly registered and deregistered.

Applying policies

То	Do
Apply a backup policy to a machine	Click
machine	In the Policy selection window, specify the backup policy you need to apply to the selected machine.

Grouping actions

То	Do		
Create a custom static or dynamic group	Click Create group.		
	In the Create group (p. 324) window, specify the required parameters of the group. The new group will be created in the group, the selected machine is a		
	member of (except for the built-in All physical machines group).		
Add a machine to another	Click Add to another group.		
static group	In the Add to group (p. 316) window, specify the group to copy the selected machine to. The backup policies applied to the groups the machine is a member of will be applied to the machine.		
For machines in custom group	For machines in custom groups		
Add machines to a static	Click Add machines to group.		
group	In the Add machines to group (p. 317) window, select the machines that you need to add.		
Move a machine to another static group	Click Move to another group.		
	In the Move to group (p. 317) window, select the group to move the machine to.		
	All the backup policies applied to the group the machine was in will be revoked. The backup policies applied to the group the machine is now a member of will be deployed to the machine.		

Remove a machine from the current static group	Click Remove from group.
	The backup policies applied to the group will be revoked from the machine automatically.

Deleting the selected machine from the management server

То	Do
Delete a machine from the management server	Click Delete machine from AMS. As a result, backup policies are revoked and shortcuts to centralized vaults are deleted from the machine. If the machine is not available at the moment, these actions will be performed on the machine as soon as the machine becomes available to the management server.

Other actions

Direct management operations	
Create a backup plan on a machine	Click Backup . This operation is described in depth in the Creating a backup plan (p. 203) section.
Recover data	Click Recover. This operation is described in depth in the Recovering data section.
Connect to a machine directly	Click Connect directly. Establishes a direct connection to the managed machine. Enables to administer a managed machine and perform all the direct management operations.
Other operations	
View detailed information on a machine	Click View details. In the Machine details (p. 317) window, examine information on the machine.
View tasks existing on a machine	Click View tasks. The Tasks (p. 339) view will display a list of the tasks, existing on the machine.
View log entries of a machine	Click View log. The Log (p. 342) view will display a list of the machine's log entries.
Update all information related to the machine	Click Synchronize. The management server will query the machine and update the database with the most recent information. Along with synchronizing, the refresh operation will be performed automatically in order to update the list of the machines.
Refresh a list of machines	Click Refresh. The management console will update the list of machines from the management server with the most recent information. Though the list of machines is refreshed automatically based on events, the data may not be retrieved immediately from the management server due to some latency. Manual refresh guarantees that the most recent data is displayed.

Adding a machine to the management server

To be able to deploy backup policies from Acronis Backup & Recovery 10 Management Server to a managed machine and perform other centralized management operations, you need to register the machine on the management server.

To add a machine

- 1. In the Navigation tree, select \(\frac{1}{2} \) Physical machines.
- 2. Click Add a machine to AMS on the toolbar.
- 3. In the **IP/Name** field, enter the machine's name or its IP address, or click **Browse...** and browse the network for the machine.

Note for Virtual Edition users: When adding a VMware ESX/ESXi host, enter the IP of the virtual appliance running Acronis Backup & Recovery 10 Agent for ESX/ESXi.

4. Specify the user name and password of a user who is a member of the **Administrators** group on the machine.

Note for Virtual Edition users: When adding a VMware ESX/ESXi host, specify the user name and password for your vCenter or ESX/ESXi host.

Click **Options>>** and specify:

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

Select the **Save password** check box to store the password for future connections.

5. Click OK.

Initiating registration on the machine side

The registration procedure can be initiated on the machine side.

- Connect the console to the machine where Acronis Backup & Recovery 10 agent is installed. If prompted for credentials, specify credentials of a member of the **Administrators** group on the machine.
- 2. Select from the menu **Options** > **Machine options** > **Machine management**.
- 3. Select **Centralized management** and specify the management server where to register the machine. Refer to "Machine management (p. 99)" for details.

Importing machines from Active Directory

To import machines from Active Directory

- 1. In the Navigation tree, select Physical machines, or Im All physical machines.
- 2. Click Import machines from Active Directory on the toolbar.
- 3. In the **Search for** field, type the machine's (or the organizational unit) name, then click **Search**. You can use the asterisk (*) to substitute for zero or more characters in a machine (or an organizational unit) name.

The left part of the window displays the machine (or organizational unit) names that fully or just partly coincide with the entered value. Click the item you want to add for import, then click Add>>. The item will be moved to the right part of the window. To add all the found items, click Add all>>.

If more than 1000 matches are found, only the first 1000 items will be displayed. In this case, it is recommended that you refine your search and try again.

The right part of the window displays the items you selected for import. If required, remove the erroneously selected items by using the respective **X Remove** and **Remove** all buttons.

4. Click **OK** to start import.

Synchronizing machines with a text file

During synchronization, the management server adjusts the **All physical machines** group in accordance with the list of machines provided in a .txt or .csv file. The management server:

- Adds machines that are present in the list but are not registered
- Deletes registered machines not present in the list
- Deletes and then tries again to add registered machines that are present in the list, but their current availability (p. 317) is Withdrawn.

As a result, only those physical machines that are listed in the file will be present in the **All physical** machines group.

Text file requirements

The file should contain machine names or IP addresses, one machine per line.

Example:

```
Machine_name_1
Machine_name_2
192.168.1.14
192.168.1.15
```

Specifying an empty file leads to deletion of all physical machines from the management server.

A registered machine has to be specified by its registration address, that is, you need to provide exactly the same host name, fully qualified domain name (FQDN), or IP address as was specified when the machine was initially added to the management server. Otherwise, the machine will be deleted and added again as if it were another machine. This means all policies, both inherited and directly applied, will be revoked from the machine and its static group membership will be lost.

The registration address of each machine can be found in the **Registration address** column in any management server view that contains the machine (the column is hidden by default).

To avoid a discrepancy, you can initially import the machines from a text file. Modify this file later as required, by adding and removing machines, but do not change the names/addresses of the machines that have to remain registered.

To synchronize machines with a text file

- 1. In the Navigation tree, select Physical machines or All physical machines.
- 2. Click **Synchronize machines with text file** on the toolbar.
- 3. In the **Path** field, enter the path to a .txt or .csv file containing the list of machines, or click **Browse** and select the file in the **Browse** window.
- 4. Under **Logon settings**, specify the user name and password of a user who is a member of the Administrators group on all machines listed in the file.
- 5. Click **OK** to start synchronizing the machines.

Synchronization command line tool

Acronis Backup & Recovery 10 Management Server has a command line tool that enables you to create a batch file and schedule the synchronization task using Windows scheduler.

To synchronize machines with a text file using command line

- 1. Log on as a member of the Acronis Centralized Admins security group.
- 2. In the command prompt, change the directory to the folder where Acronis Backup & Recovery 10 Management Server has been installed—by default: C:\Program Files\Acronis\AMS.
- 3. Run the following command:

```
syncmachines [path_to_the_file] {username password}
```

where:

- [path_to_the_file] is the path to a .txt or .csv file containing the list of machines. The tool does not accept spaces in the path name.
- {username password} belong to a user who is a member of the Administrators group on all machines listed in the file. If not specified, the single sign-on mechanism is used to operate on all the machines.

Importing machines from a text file

To import machines from a file

- 1. In the Navigation tree, select Physical machines, or All physical machines.
- 2. Click is Import machines from file on the toolbar.
- 3. In the **Path** field, enter a path to the .txt or .csv file, or click **Browse** and select the file in the **Browse** window.

A .txt or .csv file should contain machine names or their IP addresses, beginning from a new line for each of the machines.

Example:

```
Machine_name_1
Machine_name_2
192.168.1.14
192.168.1.15
```

- 4. Under **Logon settings**, specify the user name and password of a user who is a member of the Administrators group on all machines that are listed in the file.
- 5. Click **OK** to start import.

Adding a machine to another group

To add the selected machine to another group

- 1. Select the group the machine will be added to.
- 2. Click OK.

The machine being added becomes a member of more than one group. As a result, the backup policies applied to the first group will remain on the machine, and the backup policies applied to the second, third, etc. group will be deployed to the machine.

Moving a machine to another group

To move the selected machine to another group

- 1. In the group tree, select the group the machine will be moved to.
- 2. Click OK.

The machine being moved leaves one group and becomes a member of another group. As a result, the backup policies applied to the first group will be revoked from the machine, and the backup policies applied to the second group will be deployed to the machine.

Adding machines to a group

To add machines to the selected group

- 1. In the groups tree, select the group whose machines you need to add.
- 2. In the right part of the window, select the machines.
- 3. To add more machines from other groups, repeat the steps 1 and 2 for each group.
- 4. Click **OK** to add machines.

Once the machines appear in the group, the policy that was applied to the group (if any), is deployed to the machines. If any of the selected machines is not available or reachable at the moment, the action will be kept in the management server as pending and will be performed as soon as the machine becomes available to the server.

Machine details

Accumulates in four tabs all information on the selected machine. Lets the management server administrator perform operations with the backup plans and tasks existing on the machine, and policies applied to the machine.

This information is also duplicated on the **Information** panel.

Machine

The tab displays the following information on the registered machine:

- Name name of the selected machine (taken from the Computer name in Windows)
- IP address IP address of the selected machine
- Status the machine's status. Determined as the most severe status (p. 191) of all backup plans (both local and centralized) existing on the machine and backup policies (p. 306) applied to the machine.
- Last connect how much time has passed since the management server last connected to the machine.
- Last successful backup how much time has passed since the last successful backup.
- Availability:
 - Online the machine is available for the management server. This means that the management server's last connection to the machine was successful. Connection is established every 2 minutes.
 - Offline the machine is unavailable for the management server: it is turned off, or its network cable is unplugged.

- Unknown this status is displayed until the first connection between the management server and the machine is established after adding the machine or starting the management server's service.
- Withdrawn the machine was registered on another management server, or the Standalone management parameter is selected in the Options > Machine options > Machine management (p. 99). As a result, it is not possible to control the machine from the current management server. However, you are able to regain control over the machine by specifying the management server address in the Machine management settings.
- **Expired** the trial period of the machine's agent has expired. To specify a full license key, use the **Change License** functionality, or run the setup program and follow its instructions.
- Installed agents full name of Acronis agents, installed on the machine.
- Operating system the operating system the machine's agent runs.
- Processor the type of CPU used in the managed machine
- CPU clock clock rate of the CPU
- RAM memory size
- Comments the machine's description (taken from the Computer description in Windows)

Backup policies

Displays a list of backup policies applied to the selected machine and lets the management server administrator perform the following operations:

То	Do
View details of a policy	Click View details.
	In the Policy details (p. 309) window, examine all information related to the selected backup policy.
View tasks of a policy	Click View tasks.
	The Tasks (p. 339) view will display a list of the tasks related to the selected backup policy.
View log of a policy	Click View log.
	The Log (p. 342) view will display a list of the log entries related to the selected backup policy.
Revoke policy from the	Click Revoke.
machine.	The management server will revoke the policy from the machine. The policy itself remains on the management server.
	In case the machine is a member of a group and the policy is applied to the group, you cannot revoke the policy from a single machine without firstly removing the machine from the group.
Examine where the applied policy has come from	Click Explore inheritance .
	The Inheritance order (p. 322) window will display the inheritance order of the policy applied to the machine.

Filtering and sorting

Filtering and sorting of the backup policies is performed in the same way as for the **Backup policies** view. See the Filtering and sorting backup policies (p. 308) section for details.

Plans and tasks

Displays a list of the plans (both local and centralized) and tasks existing on the selected machine.

Operations

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

То	Do
View details of a plan/task	Backup plan
	Click View details. In the Plan Details (p. 200) window, review the plan details.
	<u>Task</u>
	Click View details. In the Task Details (p. 198) window, review the task details.
View plan's/task's log	Backup plan
	Click View log.
	You will be taken to the Log (p. 201) view containing the list of the plan-related log entries.
	<u>Task</u>
	Click View log.
	You will be taken to the Log (p. 201) view containing the list of the task-related log entries.
Run a plan/task	Backup plan
	Click Run.
	In the Run Backup Plan (p. 198) window, select the task you need to run.
	Running the backup plan starts the selected task of that plan immediately in spite of its schedule and conditions.
	<u>Task</u>
	Click Run.
	The task will be executed immediately in spite of its schedule and conditions.
Stop a plan/task	Backup plan
	Click Stop.
	Stopping the running backup plan stops all its tasks. Thus, all the task operations will be aborted.
	<u>Task</u>
	Click Stop.
	What will happen if I stop the task?
	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 again.

- 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 when you stopped the task. To recover all the files, you will have to run the task once again.

Edit a plan/task

Backup plan

Click P Edit.

Backup plan editing is performed in the same way as creation (p. 203), except for the following **limitations**:

It is not always possible to change backup scheme properties if the created archive is not empty (i.e. contains backups).

- 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.

In all other cases the scheme can be changed, and should continue to operate as if the existing archives were created by a new scheme. For empty archives all changes are possible.

Why can't I edit the backup plan?

- The backup plan is currently running.
 - Editing of the currently running backup plan is impossible.
- The backup plan has a centralized origin.

Direct editing of centralized backup plans is not possible. You need to edit the original backup policy.

<u>Task</u>

Click P Edit.

Why can't I edit the task?

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.

Delete a plan/task	Backup plan
	Click X Delete.
	What will happen if I delete the backup plan?
	The plan's deletion deletes all its tasks.
	Why can't I delete the backup plan?
	■ The backup plan is in the "Running" state
	A backup plan cannot be deleted, if at least one of its tasks is running.
	■ 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.
	<u>Task</u>
	Click X Delete.
	Why can't I delete the task?
	■ 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.
Refresh table	Click C Refresh.
	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.

Filtering and sorting

Filtering and sorting of the backup policies is performed in the same way as in the **Backup plans and tasks** view for direct management. See the Filter and sort backup plans and tasks (p. 197) section for details.

Member of

This tab appears only if the selected machine is added to one or more custom groups and displays a list of the groups the machine is a member of.

Operations

То	Do
View details of a group	Click View details.
	You will be taken to the Group details window, where you can examine all information related to this group.
View tasks related to a group	Click View tasks. You will be taken to the Tasks view with pre-filtered tasks related to the selected backup group.
View log related to a group	Click View log.

	This opens Log view with pre-filtered log entries of the selected group.
Remove machine from a group.	Click Remove. The centralized plans, which were deployed to the parent group, will no longer affect this machine.

Hosted virtual machines

The tab displays a list of the machines hosted on the selected virtualization server or managed by the specified virtual appliance.

You can create a dynamic group based on the list of the hosted virtual machines. To do this, click **Create a dynamic group**. The created group will be accessible in the Virtual machines view (p. 328).

Inheritance order

The **Inheritance order** window lets you examine where the policy applied to the machine came from.

The policy that was directly applied to the machine is displayed as follows:

Machine name

The policy that is applied on the machine through inheritance is displayed as in the following example:

Group1 in the root contains *Group2* to which the policy is applied directly. *Group2*, in turn, contains child *Group3* that inherits the policy from the parent and applies the policy to *Machine1* respectively.

The machine (or group) to which the policy was applied directly is boldfaced and marked with an icon.

All items are interactive, i.e. when you click on a machine or a group, its parent group view will be opened.

Filtering and sorting machines

То	Do
Sort machines by any column	Click the column's header to sort the machines in ascending order.
	Click it once again to sort the machines in descending order.
Filter machines by name.	Type a machine's name in the field below the corresponding column's header. As a result you will see the list of machines, whose names fully or just partly coincide with the entered value.
Filter machines by status, last connect, last backup, availability.	In a field below the corresponding column's header, select the required value from the list.

Configuring the machines table

By default, the table has five 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.

7.1.3.2 Actions on groups

Actions are available when you select the Physical machines view in the Navigation tree, and then click on a group.

The following is a guideline for you to perform actions on selected groups.

То	Do
Create a custom static or a dynamic group	Click Create group.
	In the Create group (p. 324) window, specify the required parameters of the group.
	Custom groups can be created in the root folder (Physical machines), or in other custom groups.
Apply a backup policy to a	Click 3 Apply backup policy.
group	In the Policy selection window, specify the backup policy you need to apply to the selected group. If there are child groups in the selected group, the backup policy will be applied to them as well.
View detailed information on	Click View details.
a group	In the Group details (p. 326) window, examine information on the selected group.
Rename a custom	Click Rename.
group/subgroup	In the Name column, type a new name for the selected group.
	Built-in groups cannot be renamed.
Edit a custom group	Click Edit .
	In the Edit group (p. 326) window, change the required parameters of the group.
Move one custom group to	Click Move to.
another	In the Move to group (p. 326) window, specify a group that will be a new parent of the selected group.
Delete a custom group	Click X Delete.
	Deletion of a parent group will delete its child groups as well. Backup policies applied to the parent group and inherited by its child groups will be revoked from all members of the deleted groups. The the policies that are directly applied to the members will remain.

Refresh a list of groups	Click C Refresh.
	The management console will update the list of groups from the management server with the most recent information. Though the list of groups is refreshed automatically based on events, the data may not be retrieved immediately from the management server due to some latency. Manual Refresh guarantees that the most recent data is displayed.

Creating a custom static or dynamic group

To create a group

- 1. In the **Group name** field, enter a name for the group being created.
- 2. Choose the type of group:
 - a. Static to create a group that will contain machines added manually.
 - b. **Dynamic** to create a group that will contain machines added automatically according to the specified criteria.

Click the **Add criteria** and select the criterion pattern.

Operating system

All the machines running the selected operating system will be members of the dynamic group.

Organizational unit (p. 325)

All the machines belonging to the specified organizational unit (OU) will be members of the dynamic group.

IP address range

All the machines whose IP addresses are within the specified IP range will be members of the dynamic group.

Listed in txt/csv file (p. 325)

All the machines that are listed in the specified .txt or .csv file will be members of the dynamic group.

- 3. In the **Comments** field, enter a description of the created group.
- 4. Click OK.

Adding multiple criteria

Adding multiple criteria forms a condition according to the following rules:

a) All the entries of the same criteria are combined by logical addition (OR).

For example, the following set of criteria

Operating system: Windows Server 2008 Operating system: Windows Server 2003

will add to the same group all the machines whose operating system is Windows 2000 OR Windows 2003.

b) Entries of different criteria are combined by logical multiplication (AND)

For example, the following set of criteria

Operating system: Windows Server 2008 Operating system: Windows Server 2003

Organizational unit: SERVERS

IP range: 192.168.17.0 - 192.168.17.55

will add to the same group all the machines whose operating system is Windows 2000 or Windows 2003 and belong to the SERVERS organizational unit and whose IP addresses are within the range 192.168.17.0 - 192.168.17.55.

How long does a dynamic group member remain in the group?

A dynamic group member remains in the group as long as the member meets the criteria. The member is removed from the group automatically as soon as

- the member changes so that it no longer meets the criteria
- the administrator changes the criteria so that the member no longer meets the criteria

There is no way to remove a machine from a dynamic group manually except for deleting the machine from the management server.

Organizational unit criterion

Organizational unit criterion is specified for the domain the management server is currently in, as follows: *OU=OU1*

Select an organizational unit from the Active Directory tree by clicking **Browse**, or typing it manually. If the domain access credentials were not specified in the management server options, the program will ask you to provide them. The credentials will be saved in the Domain access credentials (p. 97) option.

For example, suppose that the domain *us.corp.example.com* has OU1 (which is in the root), OU1 has OU2, and OU2 has OU3. And you need to add the machines of OU3. So, the criterion will be: *OU=OU3*, *OU=OU1*

If OU3 has child containers and you also need to add the machines of those containers to the group, select the **Include child containers** check box.

Listed in txt/csv file criterion

When you use this criterion, the dynamic group will include machines from the list given in the specified .txt or .csv file.

If you later modify the file, the contents of the group will change accordingly. The file is checked every 15 minutes.

If you later delete the file or if it becomes unavailable, the contents of the group will correspond to the list that was last stored in the file.

Text file requirements

The file should contain machine names or IP addresses, one machine per line.

Example:

```
Machine_name_1
Machine_name_2
192.168.1.14
192.168.1.15
```

A registered machine has to be specified by its registration address, that is, you need to provide exactly the same host name, fully qualified domain name (FQDN), or IP address as was specified

when the machine was initially added to the management server. Otherwise, the machine will not be added to the group. The registration address of each machine can be found in the **Registration** address column in any management server view that contains the machine (the column is hidden by default).

Move one group to another

To move the selected group to another group or to the root

- 1. In the groups tree, click the group to move the selected group to. You can move any type of custom group (either static, or dynamic) to another custom group of any type, or to the root folder.
 - The root folder of the machines tree contains *groups of the first level*. Groups that include other groups are called *parent groups*. Groups that are in parent groups are called *child groups*. All the backup policies applied to the parent group will be applied to its child groups as well.
- 2. Click OK.

Editing custom groups

Editing a custom group is performed in the same way as creating (p. 324) one.

Changing the type of group will result in its conversion. Any custom group can be converted to a dynamic group if it was static, and vice versa.

- When converting a static group to dynamic, provide grouping criteria. All the members that exist in the static group that do not match the provided criteria will be removed from the dynamic group.
- When converting a dynamic group to static, two options are available either to leave the current content of the group or to empty the group.

Group details

Aggregates in two tabs all information on the selected group. Allows performing operations with the policies applied to the group.

This information is also duplicated in the **Information** panel.

Group

Displays the following information on the group:

- Name name of the selected group
- Parent group (for subgroups only) name of the parent group
- Machines number of machines in the group
- Type type of the group (static, or dynamic)
- Criteria (for dynamic groups only) grouping criteria
- Comments the group description (if specified)

Backup policies

Displays a list of backup policies related to the group and allows performing the following operations:

То	Do	

View details of a policy	Click View details.
	In the Policy details (p. 309) window, examine all information related to the selected backup policy.
View tasks of a policy	Click View tasks.
	The Tasks (p. 339) view will display a list of the tasks related to the selected backup policy.
View log of a policy	Click View log.
	The Log (p. 342) view will display a list of the log entries related to the selected backup policy.
Revoke a policy from the	Click Revoke.
group.	The management server revokes the policy from the group. While the changes are being transferred to the machines and the agents are deleting the backup plans, the policy state of the group is Revoking . The policy itself remains on the management server.
Examine where the policy	Click Explore inheritance .
applied to the group came from	The Inheritance order (p. 327) window will display the inheritance order of the policy applied to the group.

Filtering and sorting

Filtering and sorting of the backup policies is performed in the same way as for the Backup policies view. See the Filtering and sorting backup policies (p. 308) section for details.

Inheritance order

The Inheritance order window lets you examine where the policy applied to the group came from.

The policy that is directly applied to the group is displayed as follows:

Group name

The following example illustrates how the policy that is applied on the group through inheritance is displayed.

Group1 in the root contains *Group2* to which the policy is applied directly. *Group2*, in turn, contains child *Group3* that inherits the policy from the parent.

The group to which the policy was applied directly is boldfaced and marked with an icon.

All items are interactive, i.e. when you click on a group, its parent group view will be opened.

7.1.4 Virtual machines

You can centrally manage virtual machines using either of the following methods or both:

Adding a virtual machine as a physical machine

Install Acronis Backup & Recovery 10 Agent for Windows or Agent for Linux on the virtual machine and register (p. 314) it on the management server. The machine will be treated as a physical one. You will be able to apply any backup policy to the machine, including policies that back up files.

This approach comes in handy when:

- the machine is not hosted on a virtualization server
- the virtualization product installed on the host server is not supported by Acronis Backup & Recovery 10 Advanced Server Virtual Edition
- you want to use pre/post backup or pre/post data capture commands on the machine
- you want to apply file backup policies to the machine.

Adding a virtual machine as a virtual machine

On Acronis Backup & Recovery 10 Management Server, a machine is considered virtual if it can be backed up from the virtualization host without installing an agent on the machine. This is possible when using Acronis Backup & Recovery 10 Advanced Server Virtual Edition. 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.

Adding Hyper-V virtual machines

- 1. Integration services (p. 53) have to be installed in the guest systems.
- 2. Install Acronis Backup & Recovery 10 Agent for Hyper-V on the Hyper-V host. The agent is installed as an add-on to Acronis Backup & Recovery 10 Agent for Windows.
- 3. Register (p. 314) the Hyper-V host on the management server. If the machine is already registered, skip this step.
- 4. The virtual machines hosted on the Hyper-V server appear in the All virtual machines group.

Adding ESX/ESXi virtual machines

- 1. VMware Tools (p. 53) have to be installed in the guest systems.
- 2. Acronis Backup & Recovery 10 Agent for ESX/ESXi is delivered as a virtual appliance. Do either of the following:
 - Deploy the agent (p. 330) to the ESX/ESXi server

or

- Install and configure the agent manually as described in "Installing ESX/ESXi virtual appliance"
- Add (p. 314) the virtual appliance to the management server as an ordinary physical machine

The virtual machines hosted on the ESX/ESXi server (except for the virtual appliance with the agent) appear in the **All virtual machines** group.

Virtual machines added to the management server as virtual machines are present under the **Virtual machines** in the **Navigation** tree. This section describes available operations with these machines.

7.1.4.1 Virtual machines on a management server

Availability of virtual machines

Virtual machines are displayed as available when both the agent is available for the management server and the machines are available for the agent. The list of virtual machines is refreshed dynamically every time the management server synchronizes with the agents.

When the virtualization server or the virtual appliance becomes unavailable or is withdrawn, the virtual machines are grayed out.

When virtual machines become unavailable for the agent (this happens when machines are removed from the virtualization server inventory, deleted from the disk, or the server's storage is down or disconnected), the machines disappear from the **All virtual machines** groups and other groups they are included in. Tasks that back up these virtual machines will fail with an appropriate log record; as a result, the generative policy will have the Error status.

The online or offline state of a virtual machine does not affect its backup since virtual machines can be backed up in both states.

Policies for virtual machines

Any policy that backs up disks and volumes can be applied to virtual machines as well as to physical machines. Policies that perform file-level backup cannot be applied to virtual machines. For more information about backup and recovery of virtual machines, supported guest operating systems and disk configurations, see "Backing up virtual machines (p. 51)".

What happens when a policy is applied to a group of virtual machines

Each machine will be backed up by a separate task to a separate archive. The default archive name will include the virtual machine name and the policy name. It is advisable to keep default archive naming so that you can easily find each machine's backups in the storage vault.

Grouping of virtual machines

The **Virtual machines** section of the navigation tree contains one built-in group called **All virtual machines**. You cannot modify this group manually, delete or move it. You can apply policies that back up disks or volumes to this group.

You can create both static and dynamic groups of virtual machines. Any virtual machine that is currently available can be added to a static group. You cannot create groups that contain both physical and virtual machines.

The dynamic membership criteria for virtual machines are as follows:

Virtualization server type (Hyper-V, ESX/ESXi).

Using this criterion, you can create a dynamic group of virtual machines hosted on all registered Hyper-V (or ESX/ESXi, respectively) servers. Any machine added to the servers will appear in this group. Any machine deleted from the servers will disappear from this group.

Host/VA

Using this criterion, you can create a dynamic group of virtual machines hosted on a specified virtualization server or managed by the specified virtual appliance.

7.1.4.2 VMware vCenter integration

If you are using VMware vSphere, it is recommended that you integrate the management server with your vCenter Server.

To integrate the management server with a VMware vCenter Server:

- 1. In the Navigation tree, right click Virtual machines and select VMware vCenter Integration
- 2. Click Configure integration
- 3. Select the Enable VMware vCenter integration check box
- 4. Specify the vCenter Server's IP address or name and provide access credentials for the server

5. Click OK

As a result, a group that has the same name as the vCenter Server appears on the management server under **Virtual machines**. For more information, please refer to "VMware vCenter integration (p. 98)."

To remove integration with a VMware vCenter Server:

- 1. In the Navigation tree, right click Virtual machines and select VMware vCenter Integration
- 2. Click Configure integration
- 3. Clear the Enable VMware vCenter integration check box
- 4. Click OK

The group that has the same name as the vCenter Server will be removed and the policies applied to this group or its child groups will be revoked.

Virtual machines remain in the **All virtual machines** group and in other groups if their host is managed by Agent for ESX/ESXi. Policies applied to these groups or directly to the machines continue functioning on the machines. This way, by removing the integration you remove only the machines that are not manageable.

7.1.4.3 Deploying and updating Agent for ESX/ESXi

Acronis Backup & Recovery 10 Management Server provides an easy way to deploy Agent for ESX/ESXi to every VMware ESX or ESXi server whose virtual machines you want to back up.

A virtual appliance with an agent will be created on every ESX/ESXi server you specify and registered on the management server. Virtual machines, dynamically grouped by their hosts, will appear on the management server and you will be able to apply backup policies to the virtual machines or back up each machine individually.

The update of already installed agents is performed in the same way as deployment. Upon selecting a host or cluster where the agent is installed, you will be suggested to update the agent on that host.

If you are using VMware vSphere, it is recommended that you integrate (p. 329) the management server with your vCenter Server before starting the agent deployment. You will not have to specify every host manually in this case.

To deploy Agent for ESX/ESXi to VMware ESX/ESXi servers:

- 1. In the **Navigation** tree, right click **Virtual machines** or right click the group that has the same name as the vCenter Server.
- 2. Click Deploy ESX agent.
- 3. ESX/ESXi hosts

For a vCenter Server, a list of ESX/ESXi hosts and clusters obtained from the vCenter Server will be displayed. Select the hosts and clusters to deploy the agent to, or check the **Select all** check box.

In a vCenter cluster, a single Agent for ESX/ESXi backs up virtual machines hosted on all the cluster's hosts. For more information, please see "Support for vCenter clusters (p. 332)".

You can add a single host to the list by specifying its IP address or name. Provide a user name and password for every host you add to the list. A vCenter Server cannot be specified in this window.

When you select a host or cluster where the agent is already installed, the right panel of the **ESX Agent Deployment** window displays: **Update ESX agent on this host**. Other settings are not available. If you only need the update, proceed directly to step 6.

4. [Optional] The agents' settings

You can deploy Agents for ESX/ESXi with default settings or specify custom settings for any agent. The settings are as follows:

Datastore: This is the datastore on the ESX/ESXi host where the virtual appliance will be stored. When deploying the agent onto a vCenter cluster, this is the datastore shared by all the servers included into the cluster. For more information, please see "Support for vCenter clusters (p. 332)".

Network interface: This is the host's internal network the virtual appliance will be included in. If there are multiple networks on the host, the program selects the one that is more suitable for the agent operation and specifies this network as **default**. Only those networks that have a connection to the host's Service Console (or Management Network, in terms of the VMware Infrastructure) are available for selection. This is critical to the operation of the agent.

The next setting appears differently, depending on how you are going to deploy the agent.

When deploying through the vCenter server - The account that will be used for agent connection to the vCenter server.

When deploying directly to the ESX/ESXi server - The account that will be used for agent connection to the ESX server.

The management server will use this account to establish a trusted relationship with the agent during registration. Centralized backup plans and recovery tasks, originating from the management server, will run under this account by default. This means the account must have the necessary privileges (p. 333) on the vCenter Server.

By default, the software will use the account that you have already specified, either when configuring integration with the vCenter, or when getting access to the ESX/ESXi server. You have the option to specify credentials for a different account if need be.

The virtual appliance's **time zone** will be set automatically according to the management server's time zone. You can change the time zone directly in the virtual appliance GUI as described in "Installing ESX/ESXi Virtual Appliance." Changing the account or network settings is also possible but is not recommended, unless it is absolutely necessary.

5. Licenses

Click Provide license.

When installing the trial product version, select **Use the following trial license key** and enter the trial license key. Deduplication is always enabled in the trial version.

When installing the purchased product, select **Use a license from the following Acronis License Server** and specify the license server that has the appropriate number of licenses for Acronis Backup & Recovery 10 Advanced Server Virtual Edition. You need one license for every host you selected.

To be able to deduplicate backups, an agent needs a separately sold license for deduplication. If you have imported such licenses into the license server, you can select the **Enable deduplication...** check box to let the agents acquire these licenses.

When installing the product for online backup *only*, select **Online backup only (license key is not required)**. This option presumes that you have or will obtain a subscription to the Acronis Backup & Recovery 10 Online service by the time of the first backup.

6. Click Deploy ESX agent.

Monitoring the deployment progress and result

Creating or updating virtual appliances may take some time. Watch the progress of the operations at the bottom of the virtual machines' views underneath the **Information** bar. After a virtual appliance is created and registered, a corresponding group of virtual machines appears on the management server.

If the deployment completed but the group of virtual machines is missing

Access the virtual appliance console using the vSphere/VMware Infrastructure client and check the agent configuration. Configure the agent manually, if required, as described in "Installing ESX/ESXi virtual appliance." Add the virtual appliance to the management server manually as described in "Adding a machine to the management server (p. 314)."

7.1.4.4 Support for vCenter clusters

In a vCenter cluster, a single Agent for ESX/ESXi backs up virtual machines hosted on all the cluster's hosts.

Deploying Agent for ESX/ESXi to a cluster

When configuring the agent deployment from a management server, you can select a cluster as a regular ESX host. The agent virtual appliance (VA) is deployed to a storage shared by all the cluster's hosts. Normally, this is an NFS share or a SAN-LUN attached to each of the hosts.

Let's assume that the cluster contains three servers.

- Server 1 uses storages A, B, C, D
- Server 2 uses storages C, D, E
- Server 3 uses storages B, C, D

The VA can be deployed to either C or D. If there is no storage shared by all the servers, you can import the VA manually into any of the hosts. This will work, but backup performance will be far from optimal.

After deployment, the agent virtual appliance can appear on any of the hosts included in the cluster, depending on how the load balancing is configured.

Moving the agent VA around the cluster

The agent's work is not affected when the Distributed Resource Scheduler (DRS) migrates the virtual appliance to another host.

Creating a cluster of servers that already have agents

It is recommended that you remove Agents for ESX/ESXi from all but one of the servers. Retain the agent whose VA resides on the shared storage. Restart the VA so that it becomes aware of the cluster.

7.1.4.5 Support for VM migration

This section informs what you can expect when migrating virtual machines within a datacenter using vCenter Server migration options. Performance considerations apply to both "hot" and "cold" migration.

VMotion

VMotion moves a virtual machine's state and configuration to another host while the machine's disks remain in the same location on shared storage. VMotion is fully supported for both Agent for ESX/ESXi Virtual Appliance and the virtual machines being backed up by the agent. Migration of either the virtual appliance or a machine can take place during backup.

Storage VMotion

Storage VMotion moves a virtual machine disks from one datastore to another. Migration of Agent for ESX/ESXi Virtual Appliance with Storage VMotion is possible unless a backup or recovery is in progress. During migration, the agent postpones any backup that has to start. It starts the backup after the migration has been completed.

Migration of a virtual machine with Storage VMotion during backup is possible, but the backup may fail or succeed with warnings. The agent will not be able to delete the snapshot taken before migration because the machine is gone. To avoid this situation, do not migrate a virtual machine until its backup is completed.

Performance considerations

It is critical to understand that backup performance degrades when Agent for ESX/ESXi does not have direct access to the storage where the backed up disks are. In this case, the agent cannot attach the disks. Instead, it obtains data from these disks through LAN. This process is much slower than obtaining data from directly attached disks.

So the best practice is that Agent for ESX/ESXi Virtual Appliance be hosted on a host for which all shared storages of the cluster are accessible. In this case, backup performance remains optimal, wherever (within the shared storages) a virtual machine or the virtual appliance migrates. Once a machine migrates to a local storage of a different host, its backups will run slower.

7.1.4.6 Privileges for VM backup and recovery

Once Agent for ESX(i) is deployed to a vCenter's host or cluster, any user of the vCenter Server can connect a management console to the agent. The scope of available operations depends on the privileges a user has on the vCenter Server. Only those actions are available that the user has permission to perform. The below tables contain the privileges required for backup and recovery of ESX virtual machines and, additionally, for virtual appliance deployment.

If the agent was deployed directly to an ESX/ESXi host or manually imported to the host, and you want the vCenter users to be able to connect to the agent and the below privileges to take effect, connect the agent to the vCenter Server rather than to the ESX/ESXi host. To change the connection, access the virtual appliance GUI using the vSphere Client and specify access credentials for the vCenter Server in the ESX(i)/vCenter setting.

Privileges on vCenter Server or ESX/ESXi host

Outlined in the below table are the privileges a vCenter Server user must have to perform operations on all the vCenter hosts and clusters.

To enable a user to operate on a specific ESX host only, assign the user the same privileges on the host. In addition, the **Global** > **Licenses** privilege is required to be able to back up virtual machines of a specific ESX host.

Object	Privilege	Back up a VM	Back up a VM's disk	Recover to a new VM	Recover to an existing VM	VA deployment
Datastore	Allocate space			+	+	+
	Browse datastore					+
	Low level file operations					+
Global	Licenses	+ (required on ESX host only)	+ (required on ESX host only)	+	+	
Network	Assign network			+	+	+
Resource	Assign VM to resource pool			+	+	+
Virtual machine > Configuration	Add existing disk	+	+	+		
	Add new disk			+	+	+
	Add or remove device			+		+
	Change CPU count			+		
	Memory			+		
	Remove disk	+	+	+	+	
	Rename			+		
	Settings				+	
Virtual machine > Interaction	Configure CD media			+		
	Console interaction					+
	Power off				+	+
	Power on			+	+	+
Virtual machine > Inventory	Create from existing			+	+	
	Create new			+	+	+
	Remove			+	+	+
Virtual machine > Provisioning	Allow disk access			+	+	

Virtual machine > State	Create snapshot	+	+	+	+
	Remove snapshot	+	+	+	+

Privileges for a folder

To enable a user to operate within a specific vCenter folder, assign the user the following privileges on the folder.

		Operation		
Object	Privilege	Back up a VM	Back up a VM's disk	Recover to an existing VM
Datastore	Allocate space			+
Global	Licenses	+	+	+
Network	Assign network			+
Resource	Assign VM to resource pool			+
Virtual machine > Configuration	Add existing disk	+	+	
	Add new disk			+
	Remove disk	+	+	+
	Settings			+
Virtual machine > Interaction	Power off			+
	Power on			+
Virtual machine > Inventory	Create from existing			+
	Create new			+
	Remove			+
Virtual machine > Provisioning	Allow disk access			+
Virtual machine > State	Create snapshot	+	+	+
	Remove snapshot	+	+	+

7.1.4.7 Removing Agent for ESX/ESXi

You can remove Agent for ESX/ESXi from an ESX/ESXi server by deleting the corresponding virtual appliance. When integration with vCenter is enabled, Agent for ESX/ESXi can be removed automatically from the hosts managed by the vCenter Server.

To remove Agent for ESX/ESXi automatically:

- 1. In the **Navigation** tree, right click the group that has the same name as the vCenter Server.
- 2. Click Remove ESX agents.
- 3. A list of ESX/ESXi hosts obtained from the vCenter Server will be displayed. Select the hosts to remove the agents from, or check the **Select all** check box.
- 4. Click Remove ESX agents.

What happens when you remove an agent

The virtual appliance that contains the agent will be deleted from the server's disk. The virtual machines hosted on this ESX/ESXi server will disappear from the management server or become unavailable for backup and restore, if integration with the vCenter is still enabled.

The Virtual Edition license on the license server will not be freed up automatically. Revoke it manually from this host using the **Manage licenses** tool if you need to free up the license.

7.1.5 Storage nodes

Acronis Backup & Recovery 10 Storage Node helps you to optimize usage of various resources required for the enterprise data protection. This goal is achieved through organizing managed vaults (p. 417) that serve as dedicated storages of the enterprise backup archives.

Storage node enables you to:

- relieve managed machines of unnecessary CPU load by using the storage node-side cleanup (p. 420) and storage node-side validation (p. 420).
- drastically reduce backup traffic and storage space taken by the archives by using deduplication (p. 76).
- prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted vaults (p. 416).

To learn more about Acronis Backup & Recovery 10 Storage Node, see the Acronis Backup & Recovery 10 Storage Node (p. 21) section.

The key elements of the "Storage nodes" view

Storage nodes list with toolbar

The toolbar lets you perform operations (p. 337) with the selected storage node. The list of storage nodes displays online and offline storage nodes added to the management server. It also informs you about the total number of backups and archives on the storage node.

Information panel

Contains the detailed information about the selected storage node and lets you manage the compacting task. The panel is collapsed by default. To expand the panel, click the chevron. The content of the pane is also duplicated in the **Storage node details** (p. 338) window.

Way of working with storage nodes (typical workflow)

1. Install the Acronis Backup & Recovery 10 Storage Node.

2. Create a user account for each user whom you want to allow to access the storage node.

Note: You can skip this step if both the storage node and the users' machines are in the same Active Directory domain.

For information about user rights on a storage node and in its managed vaults, see User rights on a storage node (p. 85).

- 3. Add (p. 338) the storage node to the Acronis Backup & Recovery 10 Management Server.
- 4. Create a managed vault (p. 139): specify the path to the vault, select the storage node that will manage the vault and select the management operations such as deduplication or encryption.
- 5. Create a backup policy (p. 367) or a backup plan that will use the managed vault.

7.1.5.1 Actions on storage nodes

All the operations described here, are performed by clicking the corresponding buttons on the toolbar. The operations can be also accessed from the **Storage nodes** bar (on the **Actions and tools** pane) and from the **Storage nodes** item of the main menu.

To perform an operation with a storage node added to the management server, first select the storage node.

The following is a guideline for you to perform operations with storage nodes.

То	Do
Add a storage node to	Click Add.
the management server	In the Add storage node (p. 338) window, specify the machine the storage node is installed on.
	Adding a storage node establishes a trusted relationship between the management server and the storage node, in the same way as when you add machines to the server. Once the storage node is added to the management server, you will be able to create managed vaults on the node.
Remove a storage	Click X Remove.
node from the management server	Once the storage node is removed from the management server, the vaults being managed by the storage node disappear from the vault list (p. 135) and become unavailable for performing operations. All the plans and tasks that use these vaults will fail. All the databases and vaults of this storage node remain untouched.
	It is possible to add the previously removed storage node to the management server again. As a result, all the vaults managed by the storage node will appear in the vault list and become available once again for all the plans and tasks that used these vaults.
Create a centralized	Click Create vault.
managed vault on the selected storage node	The Create managed vault page (p. 139) will be opened with the pre-selected storage node. Perform the remaining steps to create the vault.

Change the compacting task schedule	After deleting backups from deduplicating vaults, either manually or during cleanup, unreferenced data may appear in the deduplicating vaults and their databases. The compacting procedure deletes such data in order to free up more storage space. Only one compacting task is available per storage node.		
	Click Reschedule compacting.		
	In the Schedule window, set up the schedule for the compacting procedure. Only the time events (daily (p. 174), weekly (p. 176), and monthly (p. 178) schedules) are available for setting up.		
	The preset is: Start the task every 1 week on Sunday at 03:00:00 AM . Repeat once .		
View details of the	Click View details.		
storage node	In the Storage node details (p. 338) window (the content of which is duplicated on the Information panel), examine information about the storage node and the vaults managed by this node. You can also manage the compacting task: manually start and stop the task.		
Refresh the list of	Click C Refresh.		
storage nodes	The management console will update the list of storage nodes from the management server with the most recent information. Though the list of storage nodes is refreshed automatically based on events, the data may not be retrieved immediately from the management server due to some latency. Manual refresh guarantees that the most recent data is displayed.		

Adding a storage node

To add a storage node

- 1. In the **IP/Name** field, enter the name or the IP address of the machine the storage node resides on, or click **Browse...** and browse the network for the machine.
 - Use the fully-qualified domain name (FQDN) of the storage node, that is, a completely specified domain name ending in a top-level domain. Do not enter "127.0.0.1" or "localhost" as the storage node IP/name. These settings are no good even if the management server and the storage node are on the same machine, because, after the policy using the storage node is deployed, each agent will try to access the storage node as if it were installed on the agent's host.
- 2. To provide a valid user account for the machine, click **Options>>**, and specify:
 - User name. When entering a name of an Active Directory user account, be sure to also specify the domain name (DOMAIN\Username or Username@domain). The user account has to be a member of the Administrators group on the machine.
 - Password. The password for the account.

Select the **Save password** check box to store the password for the account.

3. Click OK.

Because registration requires the storage node's participation, it cannot take place when the machine is offline.

7.1.5.2 Storage node details

The **Storage node details** window accumulates in four tabs all information on the selected Acronis Backup & Recovery 10 Storage Node. This information is also duplicated on the **Information** pane.

Storage node properties

The tab displays the following information about the selected storage node:

- Name the name of the machine where the storage node is installed
- IP the IP address of the machine where the storage node is installed
- Availability:
 - Unknown this status is displayed until the first connection between the management server and the storage node is established after adding the storage node or starting the management server's service.
 - Online the storage node is available for the management server. This means that the last management server's connection to the node was successful. Connection is established every 2 minutes.
 - Offline the storage node is unavailable.
 - **Withdrawn** the storage node was registered on another management server. As a result, it is not possible to control the node from the current management server.
- Archives the total number of archives stored in all the vaults managed by the storage node
- Backups the total number of backups stored within the archives in all the vaults managed by the storage node.

Vaults

This tab displays a list of the vaults, managed by the storage node.

To open a managed vault for detailed examination and to perform operations on it, select the vault, then click View vault (on the tab's toolbar). In the Centralized vault (p. 136) view, perform the required actions.

Services

This tab displays the compacting task scheduling parameters.

Service tasks

This tab lets the management server administrator manage the compacting task and review its parameters. Only one compacting task can exist on a storage node.

7.1.6 Tasks

The **Tasks** view lets you monitor and manage tasks existing on the registered machines. You can view tasks' details, their states and execution results, as well as run, stop and delete tasks.

To find out what a task is currently doing on a machine, check the task execution state. The status of a task helps you to estimate whether the task is successfully accomplished.

To learn more about task states and statuses, see the Task states (p. 192) and Task statuses (p. 193) sections.

Way of working with tasks

- Use the filtering and sorting (p. 341) capabilities to display the desired tasks in the table.
- Select a task to take an action on it.

7.1.6.1 Actions on tasks

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

То	Do
Create a new backup plan,	Click New, and select one of the following:
or a task on a registered machine	■ Backup plan (p. 203)
	■ Recovery task
	■ Validation task (p. 251)
	Then, you have to specify the registered machine on which the selected task, or the backup plan will run.
View details of a task	Click View details. In the Tasks details (p. 198) window, examine all information related to the selected task.
View a task's log	Click View log.
	The Log (p. 342) view will display a list of the log entries related to the selected task.
Run a task	Click Run.
	The task will be executed immediately in spite of its schedule.
Stop a task	Click Stop.
	What will happen if I stop the task?
	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.
	■ recovery task (from the disk backup): The target volume will be deleted and its space unallocated – the same result you will get if the recovery is unsuccessful. To recover the "lost" volume, 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 period when you stopped the task. To recover all the files, you will have to run the task once again.
Edit a task	Click Edit .
	Why can't I edit the task?
	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 plan. Only the management server administrator can do this.
	Do not have the appropriate privilege
	Without Administrator privileges on the machine, a user cannot modify
840	Convright © Acronic Inc. 2000-2011

	tasks owned by other users
Delete a task	Click Delete. Why can't I delete the task? 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 Administrator privileges on the machine, a user cannot delete tasks owned by other users. This is a built-in compacting task Each storage node has a built-in service task called a compacting task. This task cannot be deleted.
Refresh tasks table	Click Refresh. The management console will update the list of tasks existing on the machines with the most recent information. Though the list of tasks 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.

7.1.6.2 Filtering and sorting tasks

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

То	Do
Set a number of tasks to display Select Options > Console options > Number of tasks (p. 94) and set the value. The maximum number of tasks that can be displayed is 500. If the tasks exceeds the specified value, use filters to display the tasks that a the scope.	
Sort tasks by column	Click the column's header to sort the tasks in ascending order. Click it once again to sort the tasks in descending order.
Filter tasks by name, owner, or backup plan.	Type the task's name (owner name, or the backup plan name) in the field below the corresponding column header.
	As a result you will see the list of tasks, whose names (owner names, or backup plan names) fully or just partly coincide with the entered value.

Filter tasks by ty	pe, execution	In a field below the corresponding header, select the required value from the list.
state, status, typ	e, origin, last	
result, schedule.		

Configuring tasks table

By default, the table has eight 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.

7.1.7 Log

The Acronis Backup & Recovery 10 log stores the history of actions the software does on a machine or a user does on a machine using the software. For example, when a user edits a task, an entry is added to the log. When the software executes a task, it adds multiple entries saying what it is currently doing.

Local and centralized logging in Acronis Backup & Recovery 10

Acronis Backup & Recovery 10 has local and centralized logs of events.

Local event log

A local event log holds information about Acronis Backup & Recovery 10 operations on a managed machine. For example, creating a backup plan, executing a backup plan, managing archives in personal vaults, executing a recovery task, will generate events logged in the local event log. Physically, a local event log is a collection of XML files stored on the machine. The managed machine local event log is accessible when the console is connected to the machine. Local event logging cannot be disabled.

Operations performed using bootable media are logged as well, but the log's lifetime is limited to a current session. Rebooting eliminates the log, but you can save the log to a file while the machine is booted with the media.

Acronis Backup & Recovery 10 Storage Node has its own local event log. This log's events are accessible through the centralized log only.

Centralized event log

Way of working with log entries

- The maximum number of entries stored in the centralized log is 50000. The maximum number of entries that can be displayed is 10000. In case the number of log entries is greater than 10000, use filtering and sorting capabilities to display the desired log entries in the table. You can also hide the unneeded columns and show the hidden ones. See the Filtering and sorting log entries (p. 344) section for details.
- Select the log entry (or log entries) to take action on it (them). See the Actions on log entries (p. 343) section for details.

■ Use the **Information** panel to review the 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. 345) window.

Ways to open the "Log" view with the pre-filtered log entries

Having selected items in other administration views (Dashboard, Machines, Backup policies, Tasks), you can open the Log view with already filtered log entries for the item in question. Thus, you do not have to configure filters in the log table by yourself.

View	Action
Dashboard	In the calendar, right-click on any highlighted date, and then select View log . The Log view will appear with the list of the log entries already filtered by the date in question.
Machines	Select a machine or a group of machines, then click View log . The Log view will display a list of the log entries related to the selected machine or group.
Backup policies	Select a backup policy, then click View log . The Log view will display a list of the log entries related to the selected policy.
Tasks	Select a task, and then click View log . The Log view appears with the log entries belonging to the selected task.

7.1.7.1 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.

То	Do
Select a single log entry	Click on it.
Select multiple log entries	■ non-contiguous: hold down CTRL and click the log entries one by one
	 contiguous: 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	1. Select a log entry.
	2. Do one of the following
	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	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	1. Make sure, that the filters are not set.
to a file	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	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	Click 🖺 Clear Log.
	All the log entries will be deleted from the log, and a new log entry will be created. It will contain information about who deleted the entries and when.
Set up the logging level	Click Configure logging level.
	In the Logging level (p. 95) window, specify whether to collect log events from the registered machines to the centralized log.

7.1.7.2 Filtering and sorting log entries

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

То	Do
Display log entries for a given time period	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	Press or release the following toolbar buttons:
	S to filter error messages
	⚠ to filter warning messages
	🕠 to filter information messages
Filter log entries by the original backup plan or managed entity type	Under the Backup plan (or Managed entity type) column header, select the backup plan or the type of managed entity from the list.
Filter log entries by task, managed entity, machine, code, owner	Type the required value (task name, machine name, owner name, etc.) in the field below the respective column header.
	As a result you will see that the list of log entries fully or just partly coincide with the entered value.
Sort log entries by date and time	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.

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.

7.1.7.3 Centralized 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 centralized log entry contains the following data fields:

- Type Type of event (Error; Warning; Information)
- Date Date and time when the event took place
- Policy The backup policy the event relates to (if any)
- Task The task the event relates to (if any)
- Managed entity type Type of managed entity where the event has occurred (if any)
- Managed entity The name of the managed entity where the event has occurred (if any)
- Machine The name of the machine where the event has occurred (if any)
- Code Blank or the program error code if the event type is error. Error code is an integer number that may be used by Acronis support service to solve the problem.
- Module Blank or the number of program module where an error was occurred. It is an integer number that may be used by Acronis support service to solve the problem.
- Owner User name of the policy/backup plan owner (p. 34)
- Message The event text description.

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

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

7.1.8 Reporting

Reporting provides the management server administrator with detailed and well-structured information concerning the enterprise data protection operations. Reports can be used as an instrument for profound analysis of the whole backup infrastructure within a corporate network.

The management server generates reports using statistics and logs which are collected from registered machines and are stored in the dedicated databases.

Reports are generated based on report templates. The templates define the information to be included in the report and the way the information is represented.

Acronis Backup & Recovery 10 Management Server offers report templates for:

- Registered machines
- Backup policies existing on the management server
- Local and centralized backup plans existing on the registered machines
- Local and centralized tasks existing on the registered machines
- Archives and backups stored in the centralized managed vaults
- Statistics about centralized managed vaults
- Task activities history

Reports about machines, backup policies, backup plans, tasks, and archives and backups contain information as of the current time.

Reports about vaults' statistics and task activities are interval-based and provide historical information for the specified time interval that can last from days to years, depending on the amount of data kept in the databases.

Generating reports

To start generating a report, select a report template in the **Reports** view, and then click **Generate** on the toolbar.

There are two types of report templates: customizable and predefined. In a customizable report template, you can specify which entries to include in the report, by using filters. A predefined report template is preset so that you can generate a report with one click.

The report will contain the information selected, grouped and sorted according to the template settings. The report appears in a separate interactive window that enables expanding and collapsing the tables. You can export the report to an XML file and open it later using Microsoft Excel or Microsoft Access.

7.1.8.1 Report about the machines

In this view, you can generate a report about the machines that are registered on the management server. This report consists of one or more tables.

Filters

Under **Filters**, choose which machines to include in the report. Only the machines that meet all filter criteria are included.

- Machines: The list of machines. Select either physical machines or virtual machines.
- Status: The machine statuses—OK, Warning, and/or Error.
- Last connection (physical machines only): The period within which the last connection between the machines and the management server occurred.
- Last successful backup: The period within which the last successful backup finished on each of the machines.
- Next backup: The period within which the next scheduled backup will start on each of the machines.
- **Operating system**: The operating systems that the machines run.
- IP address (physical machines only): The range for the latest-known IP addresses of the machines.

Availability (physical machines only): The types of the machines' availability—Online or Offline.

With the default filter settings, the report includes all physical machines.

Report view

Under **Report view**, choose how the report will look:

- Select whether to show all items in a single table or to group them by a particular column.
- Specify which table columns to show, and in which order.
- Specify how to sort the table.

7.1.8.2 Report about the backup policies

In this view, you can generate a report about the backup policies existing on the management server. This report consists of one or more tables.

Filters

Under **Filters**, choose which backup policies to include in the report. Only the backup policies that meet all filter criteria are included.

- Backup policies: The list of backup policies.
- Source type: The type of data backed up under the backup policies—Disks/volumes and/or Files.
- **Deployment state**: The deployment states of the backup policies—for example, **Deployed**.
- **Status**: The statuses of the backup policies—**OK**, **Warning**, and/or **Error**.
- Schedule: The types of the backup policies' schedules—Manual and/or Scheduled. Manual schedule means that the corresponding centralized backup plan runs only when you start it manually.
- Owner: The list of users who created the backup policies.

With the default filter settings, the report includes all backup policies.

Report view

Under **Report view**, choose how the report will look:

- Select whether to show all items in a single table or to group them by a particular column.
- Specify which table columns to show, and in which order.
- Specify how to sort the table.

7.1.8.3 Report about the backup plans

In this view, you can generate a report about backup plans existing on registered machines. This report consists of one or more tables.

Filters

Under **Filters**, choose which backup plans to include in the report. Only the backup plans that meet all filter criteria are included.

- Origin: The types of origin of the backup plans—Local and/or Centralized.
- Backup policies (available only for centralized backup plans): The backup policies on which the centralized backup plans are based.
- Machines: The list of machines on which the backup plans exist.

- Execution state: The execution states of the backup plans—for example, Running.
- Status: The statuses of the backup plans—OK, Warning, and/or Error.
- Last finish time: The period within which the last backup finished under each of the backup plans.
- Schedule: The types of the backup plans' schedules—Manual and/or Scheduled. Manual schedule means that a backup plan runs only when you start it manually.
- Owner: The list of users who created the backup plans.

With the default filter settings, the report includes all backup plans from all machines.

Report view

Under **Report view**, choose how the report will look:

- Select whether to show all items in a single table or to group them by a particular column.
- Specify which table columns to show, and in which order.
- Specify how to sort the table.

7.1.8.4 Report about the tasks

In this view, you can generate a report about the tasks that run on registered machines. This report consists of one or more tables.

Filters

Under **Filters**, choose which tasks to include in the report. Only the tasks that meet all filter criteria are included.

- Origin: The types of origin of the tasks—Centralized, Local, and/or Local without backup plan. A
 centralized task belongs to a centralized backup plan. A local task might not belong to a backup
 plan (for example, a recovery task).
- Backup policies (centralized tasks only): The backup policies on which the tasks are based.
- Machines: The list of machines on which the tasks exist.
- Type: The task types—for example, disk backup tasks.
- Execution state: The execution states of the tasks—for example, Running.
- Last result: The last results of the tasks—Succeeded, Succeeded with warnings, and/or Failed.
- **Schedule**: The types of the tasks' schedules—**Manual** or **Scheduled**. Manual schedule means that a task runs only when you start it manually.
- Owner: The list of users who created the tasks.
- **Duration**: The limits for the amount of time within which each of the tasks last ran.

With the default filter settings, the report includes all tasks from all machines.

Report view

Under **Report view**, choose how the report will look:

- Select whether to show all items in a single table or to group them by a particular column.
- Specify which table columns to show, and in which order.
- Specify how to sort the table.

7.1.8.5 Report about the archives and backups

In this view, you can generate a report about the archives that are stored in managed centralized vaults. This report consists of one or more tables.

Filters

Under **Filters**, choose which archives to include in the report. Only the archives that meet all filter criteria are included.

- Vaults: The list of centralized managed vaults that store the archives.
- Machines: The list of registered machines from which the archives were created.
- **Type**: The archive types—disk-level archives and/or file-level archives.
- Owner: The list of users who created the archives.
- Creation time: The period within which the newest backup was created in each of the archives.
- Occupied space: The limits for the space occupied by each of the archives.
- **Data backed up**: The limits for the total size of data that is currently stored in each of the archives. This size may differ from the occupied space because of compression or deduplication.
- Number of backups: The limits for the number of backups that each of the archives contains.

With the default filter settings, the report includes all archives that are stored in the centralized managed vaults.

Report view

Under **Report view**, choose how the report will look:

- Select whether to show all items in a single table or to group them by a particular column.
- Specify which table columns to show, and in which order.
- Specify how to sort the table.

7.1.8.6 Report about the vaults' statistics

In this view, you can generate a report about the use of the centralized managed vaults that are currently added to the management server. This report consists of one or more tables and diagrams.

Report coverage

Under **Report coverage**, choose the time interval for which you want to generate the report. The report will show the state of the selected vaults at the specified time on each day in the report period.

Filters

Under **Filters**, select which centralized managed vaults to include in the report, and whether to include information about the combined total of all the selected vaults.

A combined total is the total free and occupied space, total amount of backed up data, total number of archives and backups, and average ratios across the selected vaults.

With the default filter settings, the report includes information about all centralized managed vaults plus the combined total.

Report view

Under **Report view**, choose how the report will look:

- Specify which table columns to show, and in which order.
- Select which diagrams to include in the report. The diagrams show space usage in the vaults.

7.1.8.7 Report about the task activities

In this view, you can generate a report about the tasks that existed on registered machines within a chosen period. This report consists of one or more diagrams, one diagram per machine.

The diagrams show how many times each task finished on a particular day with each of these results: "Succeeded", "Succeeded with warnings", and "Failed".

Report coverage

Under **Report coverage**, choose the time interval for which you want to generate the report.

Filters

Under **Filters**, choose which tasks to include in the report. Only the tasks that meet all filter criteria are included.

- Origin: The types of origin of the tasks—Centralized, Local, and/or Local without backup plan. A
 centralized task belongs to a centralized backup plan. A local task might not belong to a backup
 plan (for example, a recovery task).
- Backup policies (centralized tasks only): The backup policies on which the tasks are based. The default setting means all backup policies that ever existed during the report period.
- Machines: The list of machines on which the tasks exist.
- Type: The task types—for example, disk backup tasks.
- Owner: The list of users who created the tasks.

With the default filter settings, the report includes all tasks that existed on the registered machines any time during the report period.

7.1.8.8 Column selection

In the **Column Selection** window, you can choose which table columns to include in the report and in which order.

The tables in the report will contain columns, from left to right, as listed in **Display in report**. The topmost column in the list will be the leftmost column in the report.

When choosing the columns to display, use the left arrow and right arrow buttons to include or exclude columns, and the up arrow and down arrow buttons to change the order of columns.

Some columns—such as **Machine name** in a report about machines—cannot be excluded from the list, or moved up or down in it.

7.1.8.9 Report view

In order for your web browser to correctly display dates and other information in generated reports, enable active content (JavaScript). You can allow active content to run temporarily for the currently displayed webpage, or enable it permanently. To allow active content to run temporarily in Internet

Explorer, click the Information bar that appears at the top of the webpage by default, and then click **Allow blocked content**.

To allow active content permanently

in Internet Explorer

- 1. On the **Tools** menu, click **Internet Options**, and then click the **Advanced** tab.
- 2. Select the Allow active content to run files on My Computer check box under Security.
- 3. Click OK.

in Mozilla Firefox

- 1. On the Options menu, click Content.
- 2. Make sure, that the **Enable JavaScript** check box is selected.
- 3. Click OK.

7.2 Configuring Acronis Backup & Recovery 10 components

There are three ways to configure various parameters of Acronis Backup & Recovery 10 components in Windows:

- By using Acronis Administrative Template
- By using the graphical user interface (GUI)
- By modifying the Windows registry

In Linux, instead of using the administrative template and modifying the registry, parameters are configured by editing the corresponding configuration files.

If the values of any of these parameters set through the administrative template differ from those set through the graphical user interface, the template-based parameters take precedence and are effective immediately; the parameters shown in the GUI will be changed accordingly.

The following subtopics describe each way of configuration and the parameters that can be configured through it.

7.2.1 Parameters set through administrative template

The following are the parameters of Acronis Backup & Recovery 10 components that can be set by using Acronis Administrative Template. For information on how to apply the administrative template, see How to load Acronis Administrative Template (p. 351).

The administrative template contains the configuration parameters of Acronis Backup & Recovery 10 Agent, Acronis Backup & Recovery 10 Management Server, and Acronis Backup & Recovery 10 Storage Node, as described in the correspondent subtopics of this topic.

7.2.1.1 How to load Acronis Administrative Template

The Administrative Template, provided by Acronis, enables the fine-tuning of some security related features, including encrypted communication settings. Through the Microsoft Group Policy mechanism, the template policy settings can be applied to a single computer as well as to a domain.

To load the Acronis Administrative Template

- Run Windows Group Policy Objects Editor (%windir%\system32\gpedit.msc.)
- 2. Open the Group Policy object (GPO) you want to edit.
- 3. Expand Computer Configuration.
- 4. Right click Administrative Templates.
- 5. Click Add/Remove Templates.
- 6. Click Add.
- Browse to Acronis Administrative Template (\Program files\Common Files\Acronis\Agent\Acronis_agent.adm or \Program files\Acronis\BackupAndRecoveryConsole\Acronis agent.adm), and click Open.


Once the template is loaded, you can open it and edit the desired settings. After loading the template or editing its settings, you should restart the configured component(s) or some of their services.

For detailed information about Windows GPO Editor please see:

http://msdn2.microsoft.com/en-us/library/aa374163.aspx

For detailed information about Group Policies please see:

http://msdn2.microsoft.com/en-us/library/aa374177.aspx



7.2.1.2 Acronis Backup & Recovery 10 Storage Node

The following are the parameters of Acronis Backup & Recovery 10 Storage Node that can be set by using Acronis Administrative Template.

Client Connection Limit

Description: Specifies the maximum number of simultaneous connections to the storage node by the agents that perform backup or recovery.

Possible values: Any integer number between 1 and 2147483647

Default value: 10

Acronis Backup & Recovery 10 agents connect to the storage node to access its managed vaults during backup or recovery. The **Client Connection Limit** parameter determines the maximum number of such connections that the storage node can handle simultaneously.

When this limit is reached, the storage node will use the backup queue (see the next parameter) for the agents that are awaiting connection.

Backup Queue Limit

Description: Specifies the maximum number of Acronis Backup & Recovery 10 components in the storage node's backup queue.

Possible values: Any integer number between 1 and 2147483647

Default value: 50

The backup queue is a list of Acronis Backup & Recovery 10 components that are awaiting connection to the storage node or are currently connected to it (see the previous parameter).

When the number of components in the backup queue is equal to the value in

Backup Queue Limit, and another component tries to establish a connection, the storage node does not put the component in the queue.

In this case, the component's connection to the storage node will fail. If the component is an Acronis Backup & Recovery 10 Agent, the corresponding backup or recovery task will stop with the **Failed** status.

Vault Warnings and Limits

Specifies the amount of free space in a vault (both as an absolute value and as a percentage) below which a warning or error is recorded in the log.

This parameter contains the following settings:

Vault Free Space Warning Limit

Description: Specifies the amount of free space in a managed vault, in megabytes, below which a warning is recorded in the storage node's log.

Possible values: Any integer number between 0 and 2147483647

Default value: 200

A vault's free space is the amount of free space on the medium—such as a disk volume—that stores the vault.

When the amount of free space in a vault is equal to the value in

Vault Free Space Warning Limit or less, a warning is recorded in the storage node's log, indicating the vault in question. You can view storage node warnings in the Dashboard.

Vault Free Space Warning Percentage

Description: Specifies the amount of free space in a managed vault, as a percentage of its total size, below which a warning is recorded in the storage node's log.

Possible values: Any integer number between 0 and 100

Default value: 10

The total size of a vault is the vault's free space plus the size of all archives that are contained in the vault.

For example, suppose that two vaults, Vault A and Vault B, are both stored on a disk volume. Suppose further that the size of the archives in Vault A is 20 GB and the size of the archives in Vault B is 45 GB.

If the volume has 5 GB of free space, then the total size of Vault A is 20 GB + 5 GB = 25 GB, and that of Vault B is 45 GB + 5 GB = 50 GB, regardless of the size of the volume.

The percentage of free space in a vault is the vault's free space divided by the vault's total size. In the previous example, Vault A has 5 GB / 25 GB = 20% of free space, and Vault B has 5 GB / 50 GB = 10% of free space.

When the percentage of free space in a vault is equal to the value in

Vault Free Space Warning Percentage or less, a warning is recorded in the storage node's log, indicating the vault in question. You can view storage node warnings in the Dashboard.

Note: The parameters **Vault Free Space Warning Limit** and **Vault Free Space Warning Percentage** are independent of each other: a warning will be recorded every time that either of the thresholds is reached.

Vault Free Space Error Limit

Description: Specifies the amount of free space in a managed vault, in megabytes, below which an error is recorded in the storage node's log and any backup to the vault becomes prohibited.

Possible values: Any integer number between 0 and 2147483647

Default value: 50

When the amount of free space in a vault is equal to the value in

Vault Free Space Error Limit or less, an error is recorded in the storage node's log. Backups performed to the vault will keep failing until the vault's free space is above the limit.

Vault Database Free Space Warning Limit

Description: Specifies the amount of free space, in megabytes, on the volume containing a managed vault's database, below which a warning is recorded in the storage node's log.

Possible values: Any integer number between 0 and 2147483647

Default value: 20

If the amount of free space on the volume containing a managed vault's database is less than the value in **Vault Database Free Space Warning Limit**, a warning is recorded in the storage node's log, indicating the vault in question. You can view storage node warnings in the Dashboard.

The database is stored on the storage node in a local folder whose name is specified in **Database path** when creating the vault.

Vault Database FreeSpace Error Limit

Description: Specifies the amount of free space on the volume containing a managed vault's database, in megabytes, below which an error is recorded in the storage node's log and any backup to the vault becomes prohibited.

Possible values: Any integer number between 0 and 2147483647

Default value: 10

If the amount of free space on the disk containing a managed vault's database is less than the value of **Vault Database Free Space Error Limit**, an error is recorded in the storage node's log. Backups performed to the vault will keep failing until the amount of free space is above the limit.

You can view storage node errors in the Dashboard.

The database is stored on the storage node in a local folder whose name is specified in **Database path** when creating the vault.

7.2.1.3 Acronis Backup & Recovery 10 Management Server

The following are the parameters of Acronis Backup & Recovery 10 Management Server that can be set by using Acronis Administrative Template.

Collecting Logs

Specifies when to collect log entries from machines managed by Acronis Backup & Recovery 10 Management Server.

This parameter contains two settings:

Trace State

Description: Specifies whether to collect the log entries about the components' events from the registered machines.

Possible values: True or False

Default value: True

Trace Level

Description: Specifies the minimum level of severity of collected entries. Only entries of levels greater than or equal to the value in **Trace Level** will be collected.

Possible values: 0 (Internal event), 1 (Debugging information), 2 (Information), 3 (Warning),

4 (Error), or 5 (Critical error)

Default value: 0 (all entries will be collected)

Log Cleanup Rules

Specifies how to clean up the centralized event log stored in the management server's reporting database.

This parameter has the following settings:

Max Size

Description: Specifies the maximum size of the centralized event log, in kilobytes.

Possible values: Any integer number between 0 and 2147483647

Default value: 1048576 (that is, 1 GB)

Percentage to Keep

Description: Specifies the percentage of the maximum log size to keep on cleanup

Possible values: Any integer number between 0 and 100

Default value: **95**

For details on how the centralized event log is cleaned up, see Log cleanup rules (p. 95).

Windows Event Log

Specifies when to record Acronis Backup & Recovery 10 Management Server's events into the Application event log in Windows.

This parameter has two settings:

Trace State

Description: Specifies whether to record Acronis Backup & Recovery 10 Management Server's events into the event log.

Possible values: True or False

Default value: False

Trace Level

Description: Specifies the minimum level of severity of events to be recorded into the event log. Only events of levels greater than or equal to the value in **Trace Level** will be recorded.

Possible values: 0 (Internal event), 1 (Debugging information), 2 (Information), 3 (Warning),

4 (Error), or 5 (Critical error)

Default value: 4 (only errors and critical errors will be recorded—if Trace State is set to True)

SNMP

Specifies the types of the management server's events to send notifications about by means of Simple Network Management Protocol (SNMP).

This parameter contains the following settings:

Trace State

Description: Specifies whether to send the SNMP notifications.

Possible values: True or False

Default value: False

Trace Level

Description: Specifies the minimum level of severity of events for sending SNMP notifications about them. Only notifications about events of levels greater than or equal to **Trace Level** will be sent.

Possible values: **0** (Internal event), **1** (Debugging information), **2** (Information), **3** (Warning), **4** (Error), or **5** (Critical error)

Default value: 4 (only errors and critical errors will be sent—if **Trace State** is set to **True**)

SNMP Address

Description: Specifies the network name or IP address of the SNMP server.

Possible values: Any string 0 to 32765 characters long

Default value: Empty string

SNMP Community

Description: Specifies the community name for the SNMP notifications.

Possible values: Any string 0 to 32765 characters long

Default value: public

Synchronization

Specifies how Acronis Backup & Recovery 10 Management Server connects to registered machines for deployment of centralized policies, retrieval of logs and backup plan states, and similar actions—collectively called synchronization.

This parameter has the following settings:

Maximum Connections

Description: Specifies the maximum number of simultaneous synchronization connections to keep

Possible values: Any integer number between 1 and 500

Default value: 200

If the total number of online registered machines does not exceed the value in **Maximum Connections**, connections to those machines are always kept, and the management server periodically performs synchronization with each machine.

Otherwise, it connects to a number of registered machines depending on the allotted number of simultaneous connections. After synchronization for a machine is complete, the management server may disconnect from that machine and use the free connection for synchronization with another machine, and so on.

(Note: Connections to machines with high synchronization priority—see **Period-High Priority** later in this topic—are likely to be always kept.)

Synchronization connections are unrelated to connections such as those between Acronis Backup & Recovery 10 Management Server and Acronis Backup & Recovery 10 Management Console.

Maximum Workers

Description: Specifies the maximum number of threads to use for synchronization.

Possible values: Any integer number between 1 and 100

Default value: 30

The management server's process uses special threads—called worker threads or workers—to perform synchronization for a registered machine which is connected for synchronization.

Each worker performs synchronization for exactly one machine at a time.

A connected machine to be synchronized waits for an available worker. For this reason, the actual number of workers will never exceed the maximum number of connections (see **Maximum Connections** described previously).

Period (in seconds)

Description: Specifies how often, in seconds, to perform synchronization for machines that have a normal synchronization priority—typically, the machines without currently running centralized backup tasks.

Possible values: Any integer number between 120 and 2147483647

Default value: 120

Acronis Backup & Recovery 10 Management Server tries to perform synchronization for each normal-priority machine once in the number of seconds given by **Period**, by using an available worker thread (see **Maximum Workers** described previously).

If there are fewer worker threads than normal-priority machines, the actual interval between synchronizations may be longer than the value of this parameter.

Period-High Priority (in seconds)

Description: Specifies how often, in seconds, to perform synchronization for machines that have a high synchronization priority—typically, the machines with currently running centralized backup tasks.

Possible values: Any integer number between 15 and 2147483647

Default value: 15

This parameter is analogous to the **Period** parameter described previously.

Real-Time Monitoring

Description: Specifies whether to perform real-time monitoring of registered machines instead of using a polling mechanism.

Possible values: True or False

Default value: False

By default, Acronis Backup & Recovery 10 Management Server connects to registered machines to perform synchronization—in particular, to retrieve data such as backup logs. This approach is known as a polling mechanism.

If **Real Time Monitoring** is set to **True**, the management server instead sends requests to machines to provide new data whenever it will appear, and then enters a listening mode. This approach is called real-time monitoring.

Real-time monitoring may reduce network traffic—for example, when centralized backup tasks run infrequently. However, it is effective only when there are relatively few registered machines.

Avoid enabling real-time monitoring if the number of registered machines exceeds the maximum number of simultaneous connections (see **Maximum Connections** earlier in this topic).

Second Connection Attempt

Description: Specifies whether to try to connect to a registered machine by using its last-known IP address after an attempt to connect to it by using its host name has failed.

Possible values: True or False

Default value: False

When connecting to a registered machine, Acronis Backup & Recovery 10 Management Server first uses the machine's network name—provided that the machine was added to the management server by name.

If **Second Connection Attempt** is set to **True** and a connection to the machine by using its network name has failed, the management server performs a second connection attempt, this time using the latest IP address which was associated with that network name.

We recommend setting **Second Connection Attempt** to **True** only in networks which often experience problems with their DNS servers, and provided that the machines' IP addresses change infrequently—as in cases of fixed IP addresses or long DHCP lease times.

This setting has no effect on machines that were added to the management server by IP address.

Offline Period Threshold (in seconds)

Description: Specifies the maximum interval, in seconds, between attempts to connect to a registered machine which appears to be offline.

Possible values: Any integer number between 120 and 2147483647

Default value: 1800

Normally, the management server connects to each registered machine with a certain time interval (see **Period** and **Period-High Priority** earlier in this section). When the management server discovers that the machine is offline, it doubles this interval; it keeps doubling this interval on each further attempt until reaching the value specified in

Offline Period Threshold. If the machine comes back online, the time interval becomes normal again.

This approach aims at efficient use of the management server resources and reducing the network load.

Backup

Specifies the location and initial size of the snapshot storage—a temporary file that is used when backing up data by taking a snapshot. This file is deleted as soon as the backup is complete.

With the default settings, the snapshot storage is created in a Windows' temporary files folder and occupies 50 percent of the space available on the volume containing that folder. This size may then grow if more space is needed for the snapshot.

You may want to increase the initial size of the snapshot storage—or to place it on a different volume—when experiencing problems with backing up data that changes extensively during backup.

This parameter is used when creating a backup policy and applies to all centralized backup plans that will be based on this policy. Changes to this parameter do not affect already existing backup policies (and their centralized backup plans).

This parameter has the following settings:

Snapshot Storage Path

Description: Specifies the folder in which to place the snapshot storage.

Possible values: Any string 0 to 32765 characters long

Default value: Empty string

An empty string means a temporary files folder, which is typically given by the TMP or TEMP environment variable.

You can specify a local folder on any volume, including a volume you are backing up.

Snapshot Storage Absolute Size

Description: Specifies the initial size of the snapshot storage, in megabytes.

Possible values: Any integer number between 0 and 2147483647

Default value: 0

If this setting is 0, the management server uses the Snapshot Storage Relative Size setting.

The initial size will not exceed the available space minus 50 MB.

Snapshot Storage Relative Size

This setting is effective only when the **Snapshot Storage Absolute Size** setting is **0**.

Description: Specifies the initial size of the snapshot storage as a percentage of the disk space that is available at the time of starting the backup.

Possible values: Any integer number between **0** and **100**

Default value: 50

If this setting is **0**, the snapshot storage will not be created.

The initial size will not exceed the available space minus 50 MB.

Without the snapshot storage, taking snapshots is still possible.

The size of the snapshot storage does not affect the size of the backup.

7.2.1.4 Acronis Backup & Recovery 10 Agent for Windows

The following are the parameters of Acronis Backup & Recovery 10 Agent that can be set by using Acronis Administrative Template.

Licensing

Specifies how often the agent checks its license on the license server, and how long it can work without a license server.

License Check Interval (in days)

Description: Specifies how often, in days, to check for license availability on Acronis License

Possible values: any integer number between 0 and 5

Default value: 1

Acronis Backup & Recovery 10 Agent periodically checks whether its license key is present on the license server. The first check is performed every time that Acronis Backup & Recovery 10 Agent starts and subsequent checks are performed once in the number of days given by **License Check Interval**.

When the agent cannot connect to the license server, a warning is recorded into the agent's log. You can view this warning in the Dashboard.

If the value is **0**, no license check will be performed; without a license, Acronis Backup & Recovery 10's functionality will be disabled after the number of days given by

Maximum Time Without License Server (see the next parameter).

See also License Server Connection Retry Interval later in this topic.

Maximum Time Without License Server (in days)

Description: Specifies how long, in days, Acronis Backup & Recovery 10 will work as normal until its functionality is disabled.

Possible values: any integer number between 0 and 60

Default value: 30

If Acronis License Server is unavailable, Acronis Backup & Recovery 10 will continue working with full functionality for the number of days specified in

Maximum Time Without License Server, as counted from the moment of installation or from the last successful check.

License Server Connection Retry Interval (in hours)

Description: Specifies the interval, in hours, between connection attempts when Acronis License Server is unavailable.

Possible values: any integer number between 0 and 24

Default value: 1

If, during a check for the license key (see **License Check Interval** earlier in this topic), Acronis Backup & Recovery 10 Agent could not connect to the license server, it will try to reconnect once in the number of hours given by **License Server Connection Retry Interval**.

If the value is **0**, no reconnection attempts will be performed; the agent will only check for the license as determined by **License Check Interval**.

License Server Address

Description: Specifies the network name or IP address of Acronis License Server.

Possible values: Any string 0 to 32765 characters long

Default value: Empty string

Log Cleanup Rules

Specifies how to clean up the agent log.

This parameter has the following settings:

Max Size

Description: Specifies the maximum size of the agent log folder, in kilobytes.

Possible values: Any integer number between 0 and 2147483647

Default value: 1048576 (that is, 1 GB)

Percentage To Keep

Description: Specifies the percentage of the maximum log size to keep on cleanup.

Possible values: Any integer number between 0 and 100

Default value: 95

For details on how the agent log is cleaned up, see Log cleanup rules (p. 102).

Windows Event Log

Specifies when to record Acronis Backup & Recovery 10 Agent's events into the Application event log in Windows.

This parameter has two settings:

Trace State

Description: Specifies whether to record the agent's events into the event log.

Possible values: True or False

Default value: False

Trace Level

Description: Specifies the minimum level of severity of events to be recorded into the event log. Only events of levels greater than or equal to the value in **Trace Level** will be recorded.

Possible values: **0** (Internal event), **1** (Debugging information), **2** (Information), **3** (Warning), **4** (Error), or **5** (Critical error)

(Littor), or 3 (Critical error)

Default value: 4 (only errors and critical errors will be recorded—if Trace State is set to True)

SNMP

Specifies the types of the agent's events to send notifications about by means of Simple Network Management Protocol (SNMP).

This parameter has the following settings:

Trace State

Description: Specifies whether to send the SNMP notifications.

Possible values: True or False

Default value: False

Trace Level

Description: Specifies the minimum level of severity of events for sending SNMP notifications about them. Only notifications about events of levels greater than or equal to **Trace Level** will be sent.

Possible values: **0** (Internal event), **1** (Debugging information), **2** (Information), **3** (Warning), **4** (Error), or **5** (Critical error)

Default value: 4 (only errors and critical errors will be recorded—if Trace State is set to True)

SNMP Address

Description: Specifies the network name or IP address of the SNMP server.

Possible values: Any string 0 to 32765 characters long

Default value: Empty string

SNMP Community

Description: Specifies the community name for the SNMP notifications.

Possible values: Any string 0 to 32765 characters long

Default value: public

Backup

Specifies the location and initial size of the snapshot storage—a temporary file that is used when backing up data by taking a snapshot. This file is deleted as soon as the backup is complete.

With the default settings, the snapshot storage is created in a Windows' temporary files folder and initially occupies 50 percent of the space available on the volume containing that folder. This size may then grow if more space is needed for the snapshot.

You may want to increase the initial size of the snapshot storage—or to place it on a different volume—when experiencing problems with backing up data that changes extensively during backup.

This parameter is used when creating a backup plan. Changes to this parameter do not affect already existing backup plans.

This parameter has the following settings:

Snapshot Storage Path

Description: Specifies the folder in which to create the snapshot storage.

Possible values: Any string 0 to 32765 characters long

Default value: Empty string

An empty string means a temporary files folder, which is typically given by the TMP or TEMP

environment variable.

You can specify a local folder on any volume, including a volume you are backing up.

Snapshot Storage Absolute Size

Description: Specifies the initial size of the snapshot storage, in megabytes.

Possible values: Any integer number between 0 and 2147483647

Default value: 0

If this setting is **0**, the management server uses the **Snapshot Storage Relative Size** setting.

The initial size will not exceed the available space minus 50 MB.

Snapshot Storage Relative Size

This setting is effective only when the **Snapshot Storage Absolute Size** setting is **0**.

Description: Specifies the initial size of the snapshot storage as a percentage of the disk space that is available at the time of starting the backup.

Possible values: Any integer number between **0** and **100**

Default value: 50

If this setting is **0**, the snapshot storage will not be created.

The initial size will not exceed the available space minus 50 MB.

Without the snapshot storage, taking snapshots is still possible.

The size of the snapshot storage does not affect the size of the backup.

7.2.1.5 Acronis Backup & Recovery 10

This section of the administrative template specifies the connection parameters and event tracing parameters for the following Acronis Backup & Recovery 10 components:

- Acronis Backup & Recovery 10 Management Server
- Acronis Backup & Recovery 10 Agent
- Acronis Backup & Recovery 10 Storage Node

Connection parameters

Remote Agent ports

Specifies the port that the component will use for incoming and outgoing communication with other Acronis components.

Select one of the following:

Not Configured

The component will use the default TCP port number 9876.

Enabled

The component will use the specified port; type the port number in the **Server TCP Port** box.

Disabled

The same as Not configured.

Client Encryption options

Specifies whether to encrypt the transferred data when the component acts as a client application, and whether to trust self-signed SSL certificates.

Select one of the following:

Not Configured

The component will use the default settings, which is to use encryption if possible and to trust self-signed SSL certificates (see the following option).

Enabled

Encryption is enabled. In **Encryption**, select one of the following:

Enabled

Data transfer will be encrypted if encryption is enabled on the server application, otherwise it will be unencrypted.

Disabled

Encryption is disabled; any connection to a server application which requires encryption will not be established.

Required

Data transfer will be performed only if encryption is enabled on the server application (see "Server Encryption options"); it will be encrypted.

Authentication parameters

Selecting the **Trust self-signed certificates** check box allows the client to connect to the server applications that use self-signed SSL certificates such as certificates created during the installation of Acronis Backup & Recovery 10 components—see SSL certificates (p. 91).

You should keep this check box selected, unless you have a Public Key Infrastructure (PKI) in your environment.

In Use Agent Certificate Authentication, select one of the following:

Do not use

The use of SSL certificates is disabled. Any connection to a server application which requires the use of SSL certificates will not be established.

Use if possible

The use of SSL certificates is enabled. The client will use SSL certificates if their use is enabled on the server application, and will not use them otherwise.

Always use

The use of SSL certificates is enabled. The connection will be established only if the use of SSL certificates is enabled on the server application.

Disabled

The same as Not configured.

Server Encryption options

Specifies whether to encrypt the transferred data when the component acts as a server application.

Select one of the following:

Not Configured

The component will use the default setting, which is to use encryption if possible (see the following option).

Enabled

Encryption is enabled. In **Encryption**, select one of the following:

Enabled

Data transfer will be encrypted if encryption is enabled on the client application, otherwise it will be unencrypted.

Disabled

Encryption is disabled; any connection to a client application which requires encryption will not be established.

Required

Data transfer will be performed only if encryption is enabled on the client application (see "Client Encryption options"); it will be encrypted.

Authentication parameters

In Use Agent Certificate Authentication, select one of the following:

Do not use

The use of SSL certificates is disabled. Any connection to a client application which requires the use of SSL certificates will not be established.

Use if possible

The use of SSL certificates is enabled. The server will use SSL certificates if their use is enabled on the client application, and will not use them otherwise.

Always use

The use of SSL certificates is enabled. The connection will be established only if the use of SSL certificates is enabled on the client application.

Disabled

The same as **Not configured**.

Event tracing parameters

In Windows, the events occurring in Acronis Backup & Recovery 10 can be recorded into the event log, a file, or both.

Each event has a level from zero to five based on the event's severity, as shown in the following table:

Level	Name	Description
0	Unknown	Event whose level of severity is unknown or not applicable

1	Debug	Event used for debug purposes
2	Information	Informational event, such as one about the successful completion of an operation or startup of a service
3	Warning	Event which is a possible impending problem, such as low free space in a vault
4	Error	Event that resulted in a loss of data or functionality
5	Critical	Event that resulted in the termination of a process such as the agent's process

Event tracing parameters are specified as the following settings in the administrative template:

File Trace Minimal Level

Description: Specifies the minimum severity level of events to be recorded in the file. Only events of levels greater than or equal to **File Trace Minimal Level** will be recorded.

Possible values: Any severity level from **Unknown** through **Critical**, or **Blocked** to not record any events

Default value: 2 (meaning that events with severity levels two through five will be recorded)

The log files are located inside the folder **%ALLUSERSPROFILE%\Application Data\Acronis**, in the **Logs** subfolder for the particular component.

Win32 Trace Minimal Level

Description: Specifies the minimum severity level of events to be recorded in the System event log. Only events of levels greater than or equal to **Win32 Trace Minimal Level** will be recorded.

Possible values: Any severity level from **Unknown** through **Critical**, or **Blocked** to not record any events

Default value: 4 (meaning that events about errors and critical errors will be recorded)

Customer Experience Program

Specifies whether the machine where the Acronis Backup & Recovery 10 component is installed will participate in the Customer Experience Program.

Select one of the following:

Not Configured

By default, the machine does not participate in the Customer Experience Program.

Enabled

In **Enable sending reports to Acronis**, select one of the following:

Enable

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. The terms of participation can be found on the Acronis Web site.

Disable

The information will not be sent.

Disabled

The same as **Not configured**.

7.2.2 Parameters set through GUI

The following parameters can be set through the graphical user interface (GUI):

- For Acronis Backup & Recovery 10 Management Server: Collecting Logs, Windows Event Log,
 SNMP, SNMP Address, and SNMP Community
- For Acronis Backup & Recovery 10 Agent: Windows Event Log, SNMP, SNMP Address,
 SNMP Community and Customer Experience Program

You will find the description of these parameters in the correspondent topic about configuration through the administrative template.

7.2.3 Parameters set through Windows registry

The following are the parameters of Acronis Backup & Recovery 10 Storage Node that can be set only by editing the registry.

Parameter related to deduplication

CompactingTriggerThreshold

Description: Specifies the percentage of used items in the data stores below which compacting occurs.

Possible values: Any integer number between 0 and 100

Default value: **80**Registry key:

 $\label{thm:local_machine} HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\ASN\Configuration\Storage\Node\Compacting\Trigge\ rThreshold$

As backups are deleted from a deduplicating vault, its deduplication data stores (p. 77) may contain unused items: files or disk blocks that are no longer referred to from any backup. The storage node processes both data stores in turn to delete the unused items. This operation is called compacting.

Since compacting is a resource-consuming operation, it should occur only when the number of unused items is significant.

The **CompactingTriggerThreshold** parameter enables you to set up a balance between the extra space required to store unused items and the compacting frequency. The larger the value of this parameter, the fewer unused items are allowed in the data stores, but compacting will likely be more frequent.

This parameter applies separately to disk-level and file-level backups. So, compacting may be performed for one data store and skipped for the other.

Parameters related to vault databases

The following two parameters determine paths to Acronis Backup & Recovery 10 Storage Node's internal databases, which contain information about managed vaults.

The database that is located in the folder specified by the **DatabasePath** parameter is typically small. However, the database that is located in the folder specified by the **TapeDatabasePath** parameter

(called the tape database), may be large if the tape library contains thousands of archives. In this case, you may want to store the tape database on a different volume.

Important: We do not recommend modifying these parameters. If you do need to modify either of them, you should do this before creating any corresponding (tape or non-tape) managed vaults. Otherwise, the storage node will lose access to those vaults until you re-attach them, and re-attaching a vault—especially a deduplicating one—may take a considerable amount of time.

DatabasePath

Description: Specifies the folder where Acronis Backup & Recovery 10 Storage Node stores its non-tape vaults database.

This database contains a list of vaults that are managed by the storage node, other than tape vaults (see the next parameter). Its typical size does not exceed a few kilobytes.

Possible values: Any string 0 to 32765 characters long

Default value: C:\Program Files\Acronis\StorageNode

Registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\ASN\Configuration\StorageNode\DatabasePath

TapesDatabasePath

Description: Specifies the folder where Acronis Backup & Recovery 10 Storage Node stores its tape vaults database.

This database contains a list of tape vaults that are managed by the storage node. Its size depends on the number of archives stored in the tape libraries, and approximately equals 10 MB per hundred archives.

Possible values: Any string 0 to 32765 characters long

Default value: C:\Documents and Settings\All Users\Application

Data\Acronis\BackupAndRecovery\TapeLocation\

Registry key:

 $\label{local_MACHINE} IN EVALUATION The SOFTWARE Acronis Backup And Recovery \end{array} Tape Location \end{array} Tape$

7.3 Creating a backup policy

A backup policy can be applied to both Windows and Linux machines.

To create a backup policy, perform the following steps.

General

Policy name

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

Source type

Select the type of items to back up: **Disk/volumes** or **Files**.

Policy credentials (p. 370)

[Optional] You can change the policy account credentials if necessary. To access this option, select the **Advanced view** check box.

Policy comments

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

Label (p. 206)

[Optional] Type a text label for the machine(s) you are going to back up. The label can be used to identify the machine or group of machines in various scenarios. To access this option, select the **Advanced view** check box.

What to back up

Items to back up (p. 370)

Specify which data items to back up on each machine the policy will be deployed to. On each of the machines, the agent will find the data items using the rules you specify. For example, if the selection rule is [All volumes], the entire machine will be backed up.

Access credentials (p. 375)

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

Exclusions (p. 375)

[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. 377)

Specify the 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 location must be available at the time when the management server starts to deploy the policy.

Access credentials (p. 378)

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

Archive comments

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

How to back up

Backup scheme (p. 378)

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 a 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

[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. 103) as set in the management server, 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 Policy** page.

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

During the backup operation, the registered machines' default backup options are ignored.

Convert to VM

Applies to: Disk/volume backup

Not effective for 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 any machine that is registered on the management server and has Acronis Backup & Recovery 10 Agent with the corresponding functionality. The archive has to be stored in a shared location, such as a network folder or a managed vault, so that the selected machine can access the archive.

When to convert (p. 229)

[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. 229)

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. 229)

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

Storage (p. 229)

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

Resultant VMs

Specify a name for the virtual machines to be created. The default name consists of variables that reflect the policy name and the name of the machine that will be backed up. You can add suffixes to the name but never delete variables, since each virtual machine has to have a distinct and unique name.

Folder on VMware vCenter

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

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

7.3.1 Policy credentials

Provide the credentials under which the centralized tasks will run on the machines.

To specify credentials

- 1. Select one of the following:
 - Use Acronis service credentials

The tasks will run under the Acronis service account, whether started manually or executed on schedule.

Use the following credentials

The tasks will 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 Acronis service credentials, see the Rights for Acronis services (p. 86) section.

To learn more about operations available depending on the user privileges, see the User privileges on a managed machine (p. 33) section.

7.3.2 Items to back up

Specify selection rules for backing up items, selected in the **Source type** field of the General section.

Volumes to back up selection rules (p. 370)

Files to back up selection rules (p. 374)

7.3.2.1 Volumes to back up selection rules

Define volume selection rules, according to which the volumes will be backed up on the machines the policy will be applied to.

To define volume selection rules

In the first line, select the rule from the list, or type it manually. To add another rule, click the next empty line, and select the rule from the list, or type it manually. The program remembers the rules typed manually, and the next time you open the window, these rules will be available for selection in the list.

The following table explains the pre-defined rules that can be selected from the list.

To include	In the Volumes column:	Comments
Windows and Linux volumes		
All volumes	Type or select: [All Volumes]	Refers to all volumes on machines running Windows, and all mounted volumes on machines running Linux.
Windows volumes		

Volume C:	Type C:\ or select it from the list	
System volume	Type or select: [SYSTEM]	The system volume contains the hardware- specific files that are needed to start Windows, such as Ntldr, Boot.ini, and Ntdetect.com.
		There is only one system volume even if multiple Windows operating systems are installed on the computer.
		For more details, see "Note on Windows machines" below.
Boot volume	Type or select: [BOOT]	Refers to the registered machine's boot volume.
		The boot volume contains the Windows folder and the supporting files for the Windows operating system (typically located in the Windows\System32 folder). It may or may not be the same as the system volume.
		If multiple operating systems are installed on the computer, this is the boot volume of the operating system in which the agent is working.
		For more details, see "Note on Windows machines" below.
All fixed volumes	Type or select: [Fixed Volumes]	Refers to all volumes other than removable media. Fixed volumes include volumes on SCSI, ATAPI, ATA, SSA, SAS and SATA devices, and on RAID arrays.
	Linux volum	es
First partition on the first IDE hard disk of a Linux machine	Type or select: /dev/hda1	hda1 is the standard device name for the first partition of the first IDE hard disk drive. For more details, see "Note on Linux machines" below.
First partition on the first SCSI hard disk of a Linux machine	Type or select: /dev/sda1	sda1 is the standard device name for the first partition of the first SCSI hard disk drive. For more details, see "Note on Linux machines" below.
First partition on the first software RAID hard disk of a Linux machine	Type or select: /dev/md1	md1 is the standard device name for the first partition of the first software RAID drive. For more details, see "Note on Linux machines" below.

The names of templates are case-sensitive.

What does a disk or volume backup store?

For supported file systems, 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.

Volumes with unsupported file systems will be backed up sector-by-sector.

Note on Windows machines

Windows operating systems prior to Windows 7 and Windows Server 2008 R2 keep system files and the loader on the same volume, unless a different volume has been specified during the system installation. If Windows files and the loader are on the same volume, selecting either [SYSTEM] or [BOOT] is enough to back up the entire operating system. Otherwise select both [SYSTEM] and [BOOT].

Operating systems starting from Windows 7 and Windows Server 2008 R2 create a dedicated system volume called **System Reserved**. If you select **[SYSTEM]**, only this dedicated volume will be backed up. Always select both **[SYSTEM]** and **[BOOT]** when backing up machines running these operating systems.

Since backup policies are commonly applied to multiple machines with various operating systems, Acronis recommends that you always select both the system and the boot volumes for backup, to ensure the integrity of every operating system.

Note on Linux machines

You can include both Windows and Linux volumes (partitions) in one centralized backup policy.

For instance, it is possible to set up a policy to back up volume **C**: on Windows machines and partition /dev/hda1 on Linux machines.

Unlike Windows, there is no clear distinction between a volume (partition) and a folder (directory) in Linux. Linux has the root partition (denoted as /), to which elements of various types—including hard disks, directories, and system devices—are attached (mounted), forming a tree similar to the file and folder structure in Windows.

For example, let a Linux machine contain a hard disk which is split into three volumes, or partitions: the first, second, and third partitions. These partitions are available in the tree as /dev/hda1, /dev/hda2, and /dev/hda3, respectively. To perform a disk backup of the, say, third partition, one can type /dev/hda3 in the row of the **Volumes to back up selection rules** dialog box.

Furthermore, a Linux partition can be mounted anywhere inside the tree. Say, /dev/hda3, can be mounted as a "subdirectory" inside the tree, such as /home/usr/docs. In this case, one can type

either /dev/hda3 or /home/usr/docs in the Volume field to perform a disk backup of the third partition.

In general, when setting up a centralized policy to perform volume backups of Linux machines, make sure that the paths entered in the Volume field correspond to partitions (such as /dev/hda2 or /home/usr/docs in the previous example), and not to directories.

Standard names for Linux partitions

Names such as /dev/hda1 reflect the standard way of naming IDE hard disk partitions in Linux. The prefix hd signifies the disk type (IDE); a means that this is the first IDE hard disk on the system, and 1 denotes the first partition on the disk.

In general, the standard name for a Linux partition consists of three components:

- Disk type; hd for IDE drives, sd for SCSI drives, md for software RAID drives (for example, dynamic volumes);
- Disk number; a for the first disk, b for the second disk, etc.;
- Partition number on the disk; 1 for the first partition, 2 for the second partition, etc.

To guarantee backing up selected disks regardless of their type, consider including three entries in the **Volumes to back up selection rules** dialog box, one for each possible type. For example, to back up the first hard disk of each Linux machine under a centralized policy, you may want to type the following lines in the Volume field:

/dev/hda1

/dev/sda1

/dev/mda1

Names for logical volumes

To back up logical volumes, also known as LVM volumes, specify their full names in the selection rules. The full name of a logical volume includes the volume group to which the volume belongs.

For example, to back up two logical volumes, **lv_root** and **lv_bin**, both of which belong to the volume group **vg_mymachine**, specify the following selection rules:

```
/dev/vg_mymachine/lv_root
/dev/vg_mymachine/lv_bin
```

To see the list of logical volumes on a machine, run the **lvdisplay** utility. In our example, the output would be similar to the following:

```
--- Logical volume ---
LV Name /dev/vg_mymachine/lv_root
VG Name vg_mymachine
...
--- Logical volume ---
LV Name /dev/vg_mymachine/lv_bin
VG Name vg_mymachine
...
```

Tip: To be able to automatically create the volume structure information during recovery, make sure that the volume with the **/etc/Acronis** directory of each machine is selected for backup. For more details, see "Saving the volume structure information" (p. 49).

7.3.2.2 Files to back up selection rules

Define file selection rules, according to which the files and (or) folders will be backed up on the machines the policy will be applied to.

To define file selection rules

In the first line, select the rule from the list, or type it manually. To add another rule, click the next empty line, and select the rule from the list, or type it manually.

The program remembers the rules typed manually, and the next time you open the window, these rules will be available for selection in the list along with the default ones.

Windows

Full path

Point to the folders and files to be backed up. If you specified a path to a file or folder explicitly, the policy will back up this item on each machine where this exact path will be found.

To include	In the Files and folders column, type or select:
File Text.doc in folder D:\Work	D:\Work\Text.doc
Folder C:\Windows	C:\Windows

Environment variables

Some environment variables point to Windows folders. Using such variables instead of full folder and file paths ensures that proper Windows folders are backed up regardless of where Windows is located on a particular machine.

To include	In the Files and folders column, type or select	Comments
Program Files folder	%PROGRAMFILES%	Points to the Program Files folder (for example, C:\Program Files)
Windows folder	%WINDIR%	Points to the folder where Windows is located (for example, C:\Windows)
Common data for all user profiles	%ALLUSERSPROFILE%	Points to the folder where the common data of all user profiles is located (typically, C:\Documents and Settings\All Users in Windows XP and C:\ProgramData in Windows Vista)

You can use other environment variables or a combination of environment variables and text. For example, to refer to the Acronis folder in the machines' Program Files folder, type: **%PROGRAMFILES%\Acronis**

Templates

Templates are similar to environment variables, but are already pre-customized.

To include	In the Files and folders column, type or select:	Comments

All files on all volumes on a machine	[All Files]	Points to all files on all volumes of the machine.
All user profiles existing on a machine	[All Profiles Folder]	Points to the folder where all user profiles are located (typically, C:\Documents and Settings in Windows XP, and C:\Users in Windows Vista).

Linux

To include	In the Files and folders column, type or select:
Text file file.txt on the volume /dev/hda3 mounted on /home/usr/docs	/dev/hda3/file.txt or /home/usr/docs/file.txt
Home directory of the common users	/home
The root user's home directory	/root
Directory for all user- related programs	/usr
Directory for system configuration files	/etc

7.3.3 Access credentials for source

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

To specify credentials

- 1. Select one of the following:
 - Use the policy credentials

The program will access the source data using the credentials of the backup policy 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 policy credentials do 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.

7.3.4 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

7.3.5 Archive

Specify where to store the archives and define names for the new backup archives.

1. Selecting the archives destination

Choose where to store machines' archives:

- Store all machines' archives in a single location
 - 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. 401) to the online backup service and activate (p. 402) the subscription on the machine(s) you want to back up. Online backup is not available in Linux.

Acronis Backup & Recovery 10 Online might be unavailable in your region. To find more information, click here: http://www.acronis.com/my/backup-recovery-online/

- To store archives in a centralized vault, expand the Centralized group and click the vault.
- To store archives 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.
- To store archives on an FTP or SFTP server, expand the corresponding group and reach the appropriate server, then select the folder that will be used for storing archives.

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.

- Store each machine's archive in the specified folder on the machine Enter the full path to the folder in the Path field. This path will be created on each machine the policy will be applied to.
- Store each machine's archive in the machine's Acronis Secure Zone Acronis Secure Zone has to be created on each machine the policy will be applied to. For information on how to create Acronis Secure Zone, see the Creating Acronis Secure Zone (p. 265) section.

2. Naming the archives

Data from each machine will be backed up to a separate archive. Specify names for the archives.

The program generates a common name for the new archives and displays it in the Name field. The name looks like [PolicyName]_[MachineName]_Archive1. If you are not satisfied with the automatically generated name, construct another name.

If you selected Store all machines' archives in a single location, you have to use variables in order to provide the unique archive names within the location.

1. Click Add variables, then select

- [Machine name] substitution for the machine's name
- [Policy name] substitution for the backup policy's name

As a result, in the Name field the following rule will appear: [Machine name]_[Policy name]_Archive1

So, if the backup policy named, say SYSTEM_BACKUP will be applied to three machines (say, FINDEPT1, FINDEPT3, FINDEPT3), the following three archives will be created in the location:

FINDEPT1_SYSTEM_BACKUP_Archive1

FINDEPT2 SYSTEM BACKUP Archive1

FINDEPT3 SYSTEM BACKUP Archive1

2. Click OK.

The name looks like ArchiveN, where N is a sequence number. If the program finds that the archive Archive1 is already stored in the location, it will automatically suggest the name Archive2.

7.3.6 Access credentials for location

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

To specify credentials

- 1. Select one of the following:
 - Use the policy credentials

The program will access the location using the credentials of the backup policy 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 policy credentials do not have access permissions to the location. You might need to provide special credentials for a network share or a storage node.

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.

7.3.7 Backup scheme selection

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.

7.3.7.1 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. 35).

7.3.7.2 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.

7.3.7.3 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. 173) section.

Retention rule	With the simple scheme, only one retention rule (p. 43) is available. Set the retention period for the backups.

7.3.7.4 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.

	Мо	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:	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. The default values for each backup type are as follows.
	Daily: 7 days (recommended minimum) Weekly: 4 weeks Monthly: indefinitely
	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. We recommend setting a retention period of at least one week for daily backups.
Advanced settings:	To specify Advanced scheduling settings (p. 182), click Change in the Advanced settings area.

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 weekWeekly: 1 monthMonthly: 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 PMBack up on: WorkdaysWeekly/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 monthsWeekly: 6 monthsMonthly: 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.

7.3.7.5 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. 174), weekly (p. 176), or monthly (p. 178) 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

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.

					6		
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
ĺ	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.

7.3.7.6 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	Specifies on what schedule and under which conditions to perform an incremental backup.
	If the archive contains no backups at the time of the task run, a full backup is created instead of the incremental backup.
Differential	Specifies on what schedule and under which conditions to perform a differential backup.
	If the archive contains no full backups at the time of the task run, a full backup is created instead of the differential backup.
Clean up archive	Specifies how to get rid of old backups: either to apply retention rules (p. 43) regularly or clean up the archive during a backup when the destination location runs out of space.
	By default, the retention rules are not specified, which means older backups will not be deleted automatically.
	Using retention rules
	Specify the retention rules and when to apply them.
	This setting is recommended for backup destinations such as shared folders or centralized vaults.
	When there is insufficient space while backing up

	The archive will be cleaned up only during backup and only if there is not enough space to create a new backup. In this case, the program will act as follows:
	 Delete the oldest full backup with all dependent incremental/differential backups
	 If there is only one full backup left and a full backup is in progress, then delete the last full backup with all dependent incremental/differential backups
	 If there is only one full backup left, and an incremental or differential backup is in progress, an error occurs saying there is a lack of available space
	This setting is recommended when backing up to a USB drive or Acronis Secure Zone. This setting is not applicable to managed vaults.
	This setting enables deletion of the last backup in the archive, in case your storage device cannot accommodate more than one backup. However, you might end up with no backups if the program is not able to create the new backup for some reason.
Apply the rules	Specifies when to apply the retention rules (p. 43).
(only if the retention rules are set)	For example, the cleanup procedure can be set up to run after each backup, and also on schedule.
	This option is available only if you have set at least one retention rule in Retention rules .
Cleanup schedule	Specifies a schedule for archive cleanup.
(only if On schedule is selected)	For example, the cleanup can be scheduled to start on the last day of each month.
	This option is available only if you selected On schedule in Apply the rules .

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. 43).

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. 43).

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.

7.3.7.7 Initial seeding

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

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

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

Please refer to the "Initial Seeding FAQ (p. 392)" section for more details.

7.3.8 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. 173) section.

8 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.com/my/backup-recovery-online/

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

Creating a backup plan (p. 203)

Creating a backup policy (p. 367)

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.

8.1 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.

8.1.1 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.

Virtual machines hosted on a VMware ESX(i) or Microsoft Hyper-V host can be backed up using a single subscription for virtual machines.

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.

8.1.2 What data can I back up and recover?

You can back up files, volumes, disks, 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.

For information about backing up virtual machines see "How to back up virtual machines to the online storage? (p. 390)"

8.1.3 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.

8.1.4 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.

8.1.5 How to back up virtual machines to the online storage?

Use either of the following methods or both.

Install Acronis software onto the virtualization host

This approach comes in handy when the virtualization product installed on the host server *is* VMware ESX(i) or Windows Server 2008 with Hyper-V or Microsoft Hyper-V Server.

You will be able to back up entire virtual machines or their volumes using a single subscription for virtual machines. The subscription applies to all machines hosted on the host or to the entire vCenter cluster. File-level backup and recovery is not possible for virtual machines, but is possible for a Windows host.

Installing the software, backing up, and recovery are described in the Quick Start Guide for Acronis Backup & Recovery 10 Virtual Edition. When installing Acronis Backup & Recovery 10 for online backup only, you do not need to enter a license key during installation.

Host-based backup is available only for paid licenses of VMware ESXi. Choose the below approach if your ESXi Server uses a free license.

Install Acronis software into the guest system

The machine will be treated as a physical one. You will need a separate server or workstation subscription for this machine. This approach comes in handy when:

- the machine is not hosted on a virtualization server.
- the virtualization product installed on the host server is not VMware ESX(i) or Windows Server 2008 with Hyper-V or Microsoft Hyper-V Server
- you want to use pre/post backup or pre/post data capture commands on the machine
- you want to perform file-level backup and recovery
- you want to back up an independent disk or an RDM disk attached in the physical compatibility mode on a running ESX(i) machine.

Installing the software, backing up, and recovery are the same as with a physical machine.

8.1.6 Backup and recovery FAQ

This section answers questions related to backup and recovery processes.

8.1.6.1 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. 392) service.

8.1.6.2 Is the online storage available under Acronis bootable media?

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

8.1.6.3 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.

8.1.6.4 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.

8.1.6.5 What happens if I run out of space?

When a machine's backups are about to exceed the storage space allowed by its subscription, you receive an e-mail notification with an alert. In addition, you can see this alert on the account management 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. 390) so that an overflow does not occur. Once the occupied space reaches the limit, the backups will cease to run.

8.1.6.6 What is the cleanup task for?

Any backup plan where the retention rule is set contains a cleanup task in addition to a backup task. The cleanup task checks the archive created by the backup plan for backups that have outlived their lifetime. If such backups are found, the task makes the 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. 390)"

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.

8.1.6.7 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. 403) 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.

8.1.7 Initial Seeding FAQ

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

8.1.7.1 What is Initial Seeding?

Initial Seeding is an extra service that lets you save an initial full backup locally and then send it to Acronis on a hard disk drive.

Acronis uploads the backup to the 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. 397) option.

8.1.7.2 Why would I want to use Initial Seeding?

This service helps you save time and network traffic during the initial full backup. It is useful when backing up very large volumes of data or entire machines to the online storage.

8.1.7.3 Is Initial Seeding a paid service?

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

8.1.7.4 What types of hard drive can I use for Initial Seeding?

Acronis accepts hard disk drives of the following interface types: IDE, ATA, SATA, USB connected drives. SCSI drives are not accepted.

You can back up directly to the device or back up to a local or network folder and then copy the backup to the device. Make sure that the device has only one volume and that the file system on that volume is NTFS or FAT32.

8.1.7.5 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.

8.1.7.6 Can I send backups taken from a number of machines on a single hard drive?

Yes. However, the number of required licenses is still one per machine.

8.1.7.7 How to buy an Initial Seeding license?

You can buy an Initial Seeding license from an Acronis partner or in Acronis online store. Follow the link www.acronis.com/my/backup-recovery-online/#buy to locate a partner or to buy online.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Click **Enter new registration code** on the same web page and register the license. The license becomes available on the **Initial Seeding / Recovery** tab.

A license purchased in Acronis online store becomes available immediately after the payment is processed.

8.1.7.8 How to perform initial seeding?

1. Decide on the media (p. 393) 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.
- Package (p. 394) 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. 396) 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.

8.1.7.9 How to package a hard drive for shipment?

It is very important that your hard drive be packaged carefully. Careful packaging will protect your drive from any damage during shipment.

Hard drive types

Acronis accepts hard disk drives of the following interface types: IDE, ATA, SATA, USB connected drives.

SCSI drives are not accepted.

Packaging

If possible, use the original packaging. Otherwise, packaging materials can be obtained at any shipping outlet or stationary store.

The following are instructions about how to package your hard disk drive.

Step 1Delicately 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:

- 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.



8.1.7.10 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. 394) 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.

8.1.8 Large Scale Recovery FAQ

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

8.1.8.1 What is Large Scale Recovery?

Large Scale Recovery is an extra service that enables you to obtain a copy of the backups you have in the 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.

8.1.8.2 Why would I use Large Scale Recovery?

In the event of a disaster or the need to recover large volumes of data or the entire machines quickly, this service helps you save time and network traffic. Recovering hundreds of gigabytes over the Internet may take days. This process will deliver a faster recovery.

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

No, these services are independent.

8.1.8.4 Is Large Scale Recovery a paid service?

Yes, you need to buy one Large Scale Recovery license per machine. The license enables you to get a disk with all of the currently available backups of this machine. To obtain backups that will be created in the future, you will need a new Large Scale Recovery license.

8.1.8.5 Can I perform large scale recovery on a different machine?

Yes. You can recover the data an unlimited number of times on any machine you wish. Acronis Universal Restore is included to help you recover an operating system to dissimilar hardware.

8.1.8.6 Can I obtain backups taken from a number of machines on a single hard drive?

No. A separate hard drive is required for each machine.

8.1.8.7 How to buy a Large Scale Recovery license?

You can buy a Large Scale Recovery license from an Acronis partner or in Acronis online store. Follow the link www.acronis.com/my/backup-recovery-online/#buy to locate a partner or to buy online.

Having purchased a license from an Acronis partner, you receive a confirmation e-mail with a registration code. Click **Enter new registration code** on the same web page and register the license. The license becomes available on the **Initial Seeding / Recovery** tab.

A license purchased in Acronis online store becomes available immediately after the payment is processed.

8.1.8.8 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.

8.1.8.9 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.

8.1.9 Subscription lifecycle FAQ

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

8.1.9.1 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).

8.1.9.2 Where do I find the subscriptions that I purchased?

If you purchased your subscriptions from an Acronis partner, you should have received an e-mail confirming the registration codes for each subscription. 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.

8.1.9.3 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.

8.1.9.4 What happens when my subscription expires?

A month before the subscription expiration date you receive an e-mail notification with an alert. In addition, you can see this alert on the account management web page near the machine. This means you need to renew (p. 400) 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.

8.1.9.5 How do I renew a subscription?

Buy another subscription and specify it as the next subscription of the same machine. The new subscription will be activated as soon as the current subscription expires.

An expired subscription can be renewed within five days after expiration. In such cases, the new subscription will be activated immediately.

Renewing a single subscription

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.

8.1.9.6 What is the "Group" column for?

So you can apply actions, such as **Renew all** or **Auto-renew all**, to your selection of the subscriptions. Specify the desired group name, for example, SalesDept, near each of the subscriptions you want to group. Click the **Group** column header to sort the subscriptions and then apply the desired actions to the group.

8.1.9.7 Can I revoke a subscription from a machine?

You cannot return an activated subscription to the list of available subscriptions, but you can reassign (p. 403) it to a different machine in Acronis Backup & Recovery 10 GUI.

8.1.9.8 Can I cancel my subscription?

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

8.2 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. 402) each subscription to a machine. The subscriptions become activated. After that, you can start backing up to Acronis Online Backup Storage.

8.3 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.

8.4 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.

8.4.1 Activating subscriptions

To begin with, make sure that the machines whose subscriptions you want to activate are registered on the management server and available (turned on).

To activate subscriptions

- 1. Connect the console to the management server.
- 2. In the Actions pane, click Activate online backup subscriptions.
- 3. Specify the credentials to log in to the online storage.
- 4. Select the machine and then click **Select subscription**.
- 5. From Available subscriptions, select the subscription that you want to activate for the machine.
- 6. Click Activate now.
- 7. Perform the previous three steps for each machine for which you want to activate a subscription.

Alternatively, you can activate a subscription when the console is connected to a machine instead of the management server.

8.4.2 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.

8.5 Configuring proxy settings

If one or more machines where you installed agents connect to the Internet through a proxy server, you need to configure each of the agents to use the proxy server.

The management server connects to the Internet to retrieve information about online backup subscriptions. The proxy settings for the management server also need to be configured.

Proxy settings for the agent and the management server are configured separately, even if both are installed on the same machine.

To configure proxy settings for an agent

- 1. Connect the console to the machine for which you want to configure proxy settings.
- 2. In the **Options** menu, located on the toolbar, click **Machine options**.
- 3. Click Online backup proxy.
- 4. Enter the proxy server settings. For detailed information (p. 102) about the settings please refer to the context-sensitive help.
- 5. Repeat steps 2–5 for all machines that connect to the Internet through a proxy server.

To configure proxy settings for the management server

- 1. Connect the console to the management server.
- 2. In the **Options** menu, located on the toolbar, click **Management server options**.
- 3. Click Online backup proxy.
- 4. Enter the proxy server settings. For detailed information (p. 102) about the settings please refer to the context-sensitive help.

8.6 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. 117)

Dual destination (p. 120)

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.

8.7 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.

9 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. 414) 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. 422) image using Bootable Media Builder. The resulting bootable media (p. 411) can be used to boot any PC-compatible machine and perform, with certain limitations, most of the direct management (p. 413) 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. 413).

Acronis Secure Zone

A secure volume for storing backup archives (p. 409) within a managed machine (p. 417). 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. 415) or a disk using the GPT partitioning style.

Acronis Secure Zone is considered as a personal vault (p. 418).

Acronis Startup Recovery Manager (ASRM)

A modification of the bootable agent (p. 411), 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. 417), 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. 412) performed by an agent (p. 409) according to the backup plan (p. 410) that produces the archive (p. 409). Agent-side cleanup is performed in unmanaged vaults (p. 421).

Agent-side validation

Validation (p. 421) performed by an agent (p. 409) according to the backup plan (p. 410) that produces the archive (p. 409). Agent-side validation is performed in unmanaged vaults (p. 421).

Archive

See Backup archive (p. 409).



Backup

A backup is the result of a single backup operation (p. 409). 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. 416) or consolidation (p. 413) are also called backups.

Backup archive (Archive)

A set of backups (p. 409) created and managed by a backup plan (p. 410). An archive can contain multiple full backups (p. 416) as well as incremental (p. 417) and differential backups (p. 413). 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. 421), 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. 417) 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. 409), 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. 410).

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. 409) (the backup archive name and location)
- the backup scheme (p. 411), that includes the backup schedule and [optionally] the retention rules
- [optionally] the archive validation rules (p. 421)
- the backup options (p. 409).

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. 420) configured for execution on a managed machine (p. 417).

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. 410) deployment (centralized plan (p. 412)).

Backup policy (Policy)

A backup plan template created by the management server (p. 418) 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. 419), 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. 417).
- 3. The management server deploys the policy to the machines.

- 4. On each machine, the agent (p. 409) 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. 410) using other rules specified by the policy. Such backup plan is called a centralized plan (p. 412).
- 6. On each machine, the agent installed on the machine creates a set of centralized tasks (p. 412) that will carry out the plan.

Backup scheme

A part of the backup plan (p. 410) that includes the backup schedule and [optionally] the retention rules and the cleanup (p. 412) schedule. For example, perform a full backup (p. 416) monthly on the last day of the month at 10:00AM and an incremental backup (p. 417) 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. 409). Bootable agent is based on Linux kernel. A machine (p. 417) can be booted into a bootable agent using either bootable media (p. 411) or Acronis PXE Server. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 413).

Bootable media

A physical media (CD, DVD, USB flash drive or other media supported by a machine (p. 417) BIOS as a boot device) that contains the bootable agent (p. 411) or Windows Preinstallation Environment (WinPE) (p. 422) with the Acronis Plug-in for WinPE (p. 408). 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. 416) 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. 418).

A management server has two built-in groups that contain all machines of each type: All physical machines (p. 418), All virtual machines (p. 421).

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. 410) can be applied to a built-in group.



Centralized backup plan

A backup plan (p. 410) that appears on the managed machine (p. 417) as a result of deploying a backup policy (p. 410) from the management server (p. 418). 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. 418). The centralized management operations include:

- creating, applying and managing backup policies (p. 410)
- creating and managing static (p. 419) and dynamic groups (p. 415) of machines (p. 417)
- managing the tasks (p. 420) existing on the machines
- creating and managing centralized vaults (p. 412) for storing archives
- managing storage nodes (p. 419)
- monitoring activities of the Acronis Backup & Recovery 10 components, viewing the centralized log and more.

Centralized task

A task (p. 420) belonging to a centralized backup plan (p. 412). Such task appears on the managed machine (p. 417) as a result of deploying a backup policy (p. 410) from the management server (p. 418) and can be modified only by editing the backup policy.

Centralized vault

A networked location allotted by the management server (p. 418) administrator to serve as storage for the backup archives (p. 409). A centralized vault can be managed by a storage node (p. 419) 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. 419) on the server. The shortcut to the vault appears on the machines in the Centralized vaults list. Any backup plan (p. 410) 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. 409) from a backup archive (p. 409) 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. 410) 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. 43).

Console (Acronis Backup & Recovery 10 Management Console)

A tool for remote or local access to Acronis agents (p. 409) and Acronis Backup & Recovery 10 Management Server (p. 418).

Having connected the console to the management server, the administrator sets up and manages backup policies (p. 410) and accesses other management server functionality, that is, performs centralized management (p. 412). Using the direct console-agent connection, the administrator performs direct management (p. 413).

Consolidation

Combining two or more subsequent backups (p. 409) belonging to the same archive (p. 409) into a single backup.

Consolidation might be needed when deleting backups, either manually or during cleanup (p. 412). For example, the retention rules require to delete a full backup (p. 416) that has expired but retain the next incremental (p. 417) 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.



Deduplicating vault

A managed vault (p. 417) in which deduplication (p. 413) 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. 409) stored on storage nodes (p. 419). 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. 416). 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. 417) using the direct console (p. 413)-agent (p. 409) connection (as opposed to centralized management (p. 412) when the operations are configured on the management server (p. 418) and propagated by the server to the managed machines).

The direct management operations include:

- creating and managing local backup plans (p. 417)
- creating and managing local tasks (p. 417), such as recovery tasks
- creating and managing personal vaults (p. 418) and archives stored there
- viewing the state, progress and properties of the centralized tasks (p. 412) 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. 411). 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. 409) 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. 415) 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. 417) are members of the same disk group.

As soon as the first dynamic disk is created by the LDM or another disk management tool, the disk group name can be found in the registry key

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\dmio\Boot Info\Primary Disk Group\Name.

The next created or imported disks are added to the same disk group. The group exists until at least one of its members exists. Once the last dynamic disk is disconnected or converted to basic, the group is discontinued, though its name is kept in the above registry key. In case a dynamic disk is created or connected again, a disk group with an incremental name is created.

When moved to another machine, a disk group is considered as 'foreign' and cannot be used until imported into the existing disk group. The import updates the configuration data on both the local and the foreign disks so that they form a single entity. A foreign group is imported as is (will have the original name) if no disk group exists on the machine.


For more information about disk groups please refer to the following Microsoft knowledge base article:

222189 Description of Disk Groups in Windows Disk Management http://support.microsoft.com/kb/222189/EN-US/

Dynamic disk

A hard disk managed by Logical Disk Manager (LDM) that is available in Windows starting with Windows 2000. LDM helps flexibly allocate volumes on a storage device for better fault tolerance, better performance or larger volume size.

A dynamic disk can use either the master boot record (MBR) or GUID partition table (GPT) partition style. In addition to MBR or GPT, each dynamic disk has a hidden database where the LDM stores the dynamic volumes' configuration. Each dynamic disk holds the complete information about all dynamic volumes existing in the disk group which makes for better storage reliability. The database occupies the last 1MB of an MBR disk. On a GPT disk, Windows creates the dedicated LDM Metadata partition, taking space from the Microsoft Reserved Partition (MSR.)



Dynamic disks organized on MBR (Disk 1) and GPT (Disk 2) disks.

For more information about dynamic disks please refer to the following Microsoft knowledge base articles:

Disk Management (Windows XP Professional Resource Kit) http://technet.microsoft.com/en-us/library/bb457110.aspx

816307 Best practices for using dynamic disks on Windows Server 2003-based computers http://support.microsoft.com/kb/816307

Dynamic group

A group of machines (p. 417) which is populated automatically by the management server (p. 418) 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. 415), or more precisely, on a disk group (p. 414). 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.)



Encrypted archive

A backup archive (p. 409) 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. 409), each backup belonging to the archive is encrypted by the agent (p. 409) 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. 417) to which anything written is encrypted and anything read is decrypted transparently by the storage node (p. 419), 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. 416) will be encrypted over the encryption performed by the agent (p. 409).

Export

An operation that creates a copy of an archive (p. 409) 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. 409) or to your choice of backups belonging to the same archive. An entire vault (p. 421) can be exported by using the command line interface.

F

Full backup

A self-sufficient backup (p. 409) 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. 411) aimed to maintain the optimal balance between a backup archive (p. 409) size and the number of recovery points (p. 419) 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. 37).

Ι

Image

The same as Disk backup (p. 414).

Incremental backup

A backup (p. 409) that stores changes to the data against the latest backup. You need access to other backups from the same archive (p. 409) to restore data from an incremental backup.

Local backup plan

A backup plan (p. 410) created on a managed machine (p. 417) using direct management (p. 413).

Local task

A task (p. 420) belonging to a local backup plan (p. 417) 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. 417), either physical or virtual, where at least one Acronis Backup & Recovery 10 Agent (p. 409) is installed.

Managed vault

A centralized vault (p. 412) managed by a storage node (p. 419). Archives (p. 409) 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. 420) and storage node-side validation (p. 420) for each archive stored in the managed vault. An administrator can specify additional operations that the storage node will perform (deduplication (p. 413), 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. 417) using backup policies (p. 410) and grouping
- enterprise-wide monitoring functionality
- the ability to create centralized vaults (p. 412) for storing enterprise backup archives (p. 409)
- the ability to manage storage nodes (p. 419).

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. 411).

P

Personal vault

A local or networked vault (p. 421) created using direct management (p. 413). 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. 419). 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. 410).

Policy

See Backup policy (p. 410).

R

Recovery point

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

Registered machine

A machine (p. 417) managed by a management server (p. 418). A machine can be registered on only one management server at a time. A machine becomes registered as a result of the registration (p. 419) procedure.

Registration

A procedure that adds a managed machine (p. 417) to a management server (p. 418).

Registration sets up a trust relationship between the agent (p. 409) 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. 410). Enables the management server (p. 418) administrator to select the data to back up within a machine.

Static group

A group of machines which a management server (p. 418) 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. 417). Storage node enables the administrator to:

- relieve managed machines (p. 417) of unnecessary CPU load by using the storage node-side cleanup (p. 420) and storage node-side validation (p. 420)
- drastically reduce backup traffic and storage space taken by the archives (p. 409) by using deduplication (p. 413)

 prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted vaults (p. 416).

Storage node-side cleanup

Cleanup (p. 412) performed by a storage node (p. 419) according to the backup plans (p. 410) that produce the archives (p. 409) stored in a managed vault (p. 417). Being an alternative to the agent-side cleanup (p. 409), the cleanup on the storage node side relieves the production servers of unnecessary CPU load.

Since the cleanup schedule exists on the machine (p. 417) the agent (p. 409) 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. 421) performed by a storage node (p. 419) according to the backup plans (p. 410) that produce the archives (p. 409) stored in a managed location (p. 417). Being an alternative to the agent-side validation (p. 409), the validation on the storage node side relieves the production servers of unnecessary CPU load.



Task

In Acronis Backup & Recovery 10, a task is a set of sequential actions to be performed on a managed machine (p. 417) 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. 411) aimed to maintain the optimal balance between a backup archive (p. 409) size and the number of recovery points (p. 419) available from the archive. Unlike the GFS (p. 417) 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. 40)".



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. 408) (using F11) or
- the image being recovered is located in Acronis Secure Zone (p. 408) or
- when using Acronis Active Restore (p. 408),

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. 421) that is not a managed vault (p. 417).



Validation

An operation that checks the possibility of data recovery from a backup (p. 409).

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. 410). Rules that define when and how often to perform validation (p. 421) and whether to validate the entire archive (p. 409) or the latest backup in the archive.

Vault

A place for storing backup archives (p. 409). 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. 412), 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. 417) is considered virtual if it can be backed up from the virtualization host without installing an agent (p. 409) 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. 408) enables running the Acronis Backup & Recovery 10 Agent (p. 409) in the preinstallation environment.