Acronis
COMPUTE WITH CONFIDENCE
WWW.ACRONIS.COM

Acronis Backup & Recovery 10

User's Guide

End User License Agreement (EULA)

Acronis, INC

BEFORE INSTALLING AND USING THE SOFTWARE PRODUCT WHICH EITHER YOU HAVE DOWNLOADED OR IS CONTAINED ON THESE DISKS ("SOFTWARE") YOU SHOULD CAREFULLY READ THE FOLLOWING LICENSE AGREEMENT ("AGREEMENT") THAT APPLIES TO THE SOFTWARE. CLICK "ACCEPT" IF YOU FULLY ACCEPT AND AGREE TO ALL OF THE PROVISIONS OF THIS AGREEMENT. OTHERWISE, CLICK "DO NOT ACCEPT." CLICKING "ACCEPT" OR OTHERWISE DOWNLOADING, INSTALLING AND OR USING THE SOFTWARE ESTABLISHES A BINDING AGREEMENT BETWEEN YOU AS THE PERSON LICENSING THE SOFTWARE (THE "LICENSEE") AND Acronis, INC. LOCATED AT: Acronis INTERNATIONAL GMBH VERWALTUNG EURO HAUS RHEINWEG 5 SCHAFFHAUSEN, SWITZERLAND CH-8200, ("LICENSOR"). IF YOU DO NOT ACCEPT ALL OF THE TERMS OF THIS AGREEMENT, YOU SHALL HAVE NOT RIGHT TO DOWNLOAD, INSTALL AND/OR USE THE SOFTWARE AND MUST DELETE THE SOFTWARE AND ASSOCIATED FILES IMMEDIATELY.

This Agreement applies to the Software, whether licensed under a Software License and/or an Evaluation License, each as defined and described below:

Purchased License of Software. Subject to the terms and conditions of this Agreement, upon purchase of a license to the Software, LICENSOR grants and LICENSEE accepts a nonexclusive, nontransferable, nonassignable license to use Software only for LICENSEE's own internal use solely on the specific number of Hardware (as defined below) licensed owned, leased or otherwise controlled by LICENSEE. LICENSEE may make one copy of Software only for archival purposes, only in machine readable form, provided that such archival copy is only used for archival purposes and never in a production environment and is marked with every notice on the original Installation of Software is LICENSEE's responsibility. The license described in this section shall be referred to as a "Software License").

Evaluation License of Software: The LICENSEE has the right to evaluate the Software for a period of time not to exceed fifteen (15) days (the "Evaluation Period') unless extended by LICENSOR. Software licensed under this Evaluation License may not be used in a production environment. There will be no charge to the LICENSEE for said evaluation of the Software under this Evaluation License. At the conclusion of the Evaluation Period, unless a Software License to the Software is purchased, the LICENSEE will delete the Software from its systems and have no further license or other rights with respect to the Software except as to the rights and responsibilities in this Agreement. LICENSOR SHALL NOT BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, PUNITIVE, OR CONSEQUENTIAL DAMAGES RESULTING FROM USE OF SOFTWARE UNDER THE EVALUATION LICENSE UNDER ANY CONTRACT, NEGLIGENCE, STRICT LIABILITY OR OTHER THEORY. THESE LIMITATIONS SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY. The following sections of this Agreement also apply to Evaluation License(s) of the Software: Limitations, Confidentiality, Disclaimer of Warranties, LICENSEE Indemnity, Law, Export Restrictions, and Miscellaneous. The license described in this section shall be referred to as an "Evaluation License").

Use Rights:

Assigning the License. Before you run any instance of the Software under a Software License, you must assign that license to one of your PCs or Servers (depending upon the license you have purchased or are evaluating and such purchase or evaluation is based upon the operating system on which that hardware operates, such PCs or Servers shall be referred to as the "Hardware"). That Hardware is the licensed Hardware for that particular Software license. You may assign other Software Licenses to the same Hardware, but you may not assign the same Software License to more than one Hardware except as identified herein.

You may reassign a Software License if you retire the licensed Hardware due to permanent Hardware failure. If you reassign a Software License, the Hardware to which you reassign the license becomes the new licensed Hardware for that particular Software License.

Running Instances of the Hardware Software. You have the rights to run the Software on one piece of Hardware as defined above. Every Hardware creating an image and every Hardware to which an image is either deployed to or restored from must have a valid license. In the case of Universal Deploy or Universal Restore, every time an image is deployed or restored (as appropriate to either Universal Restore or Universal Deploy), to Hardware that is dissimilar to the Hardware from which the image was originally created, a valid license of the Universal Deploy or Universal Restore is required.

You have the rights below for each Software License you assign:

Standard Software license. A standard Software License is the general license that is available to LICENSEE. Unless the Virtual Edition Software is purchased, you have purchased the standard Software License and may run on the licensed Hardware, at any one time as follows:

- One (1) instance of the Hardware installed Software in one physical operating system environment; and
- Up to four (4) instances of the Hardware installed Software in virtual operating system environments (only one (1) instance per virtual operating system environment).

If you run all five (5) permitted instances at the same time, the instance of the Hardware installed Software running in the physical operating system environment may be used only to run hardware virtualization software and to manage and service operating system environments on the licensed Hardware.

Virtual Edition. Virtual Edition (as identified by the product name (for example, True Image Virtual Edition)) is licensed by physical Hardware. You may run on the licensed Hardware, at any one time as follows:

- One instance of the Hardware software in one physical operating system environment; and
- As many instances of the Hardware software in virtual operating system environments.

Support. If LICENSEE is under a current support contract with LICENSOR with respect to the Software and is current in paying all amounts due thereunder, LICENSOR shall make available to LICENSEE the support described in this paragraph (the "Support") on a twenty four (24) hours a day, seven (7) days per week basis. Support shall consist of: (i) supplying telephone or other electronic support, as determined by LICENSOR in its sole discretion, to LICENSEE in order to help LICENSEE locate and, on its own, correct problems with the Software and (ii) supplying all extensions, enhancements and other changes that LICENSOR, at its sole discretion, makes or adds to the Software and which LICENSOR makes generally available, without additional charge, to other licensees of the Software that are enrolled in Support. Upon mutual written agreement by both parties, LICENSOR may, but shall not be required to: (i) supply code corrections to LICENSEE to correct Software malfunctions in order to bring such Software into substantial conformity with the published operating specifications for the most current version of the Software unless LICENSEE's unauthorized modifications prohibit or hamper such corrections or cause the malfunction; or (ii) supply code corrections to correct insubstantial problems at the next general release of the Software.

Limitations. Notwithstanding any references to "purchase" the Software is licensed and not sold pursuant to this Agreement. This Agreement confers a limited license to the Software and does not constitute a transfer of title to or sale of all or a portion of the Software, and LICENSOR retains ownership of all copies of the Software. LICENSEE acknowledges that the Software contain trade secrets of LICENSOR, its suppliers or licensors, including but not limited to the specific internal design and structure of individual programs and associated interface information. Accordingly, except as otherwise expressly provided under this Agreement, LICENSEE shall have no right, and LICENSEE specifically agrees not to: (i) transfer, assign or sublicense its license rights to any other person or entity, or use the Software on ay equipment other than the PC, and LICENSEE acknowledges that any attempted transfer, assignment, sublicense or use shall be void; (ii) make error corrections to or otherwise modify or adapt the Software or create derivative works based upon the Software, or permit third parties to do the same; (iii) reverse engineer or decompile, decrypt, disassemble or otherwise reduce the Software to human-readable form, except to the extent otherwise expressly permitted under applicable law notwithstanding this restriction; (iv) use or permit the Software to be used to perform services for third parties, whether on a service bureau or time sharing basis or otherwise, without the express written authorization of LICENSEE; or (v) disclose, provide, or otherwise make available trade secrets contained within the Software in any form to any third party without the prior written consent of LICENSOR.

Confidentiality. The Software is a trade secret of LICENSOR and is proprietary to LICENSOR. LICENSEE shall maintain Software in confidence and prevent disclosure of Software using at least the same degree of care it uses for its own similar proprietary information, but in no event less than a reasonable degree of care. LICENSEE shall not disclose Software or any part thereof to anyone for any purpose, other than to employees for the purpose of exercising the rights expressly granted under this Agreement. License shall not, and shall not allow any third party to, decompile, disassemble or otherwise, reverse engineer or attempt to reconstruct or discover any source code or underlying ideas, algorithms, file formats or programming or interoperability interfaces of Software or of any files contained or generated using Software by any means whatsoever. The obligations under this paragraph shall survive any termination of the Agreement.

Disclaimer of Warranties. THE SOFTWARE IS PROVIDED "AS IS" AND LICENSOR DISCLAIMS ALL WARRANTIES, EXPRESSED OR IMPLIED WITH RESPECT TO SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY, TITLE, NON-INFRINGEMENT OF THIRD PARTIES' RIGHTS, AND FITNESS FOR A PARTICULAR USE. WITHOUT LIMITING THE FOREGOING, LICENSOR DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN SOFTWARE WILL OPERATE IN THE COMBINATION LICENSEE SELECTS, THAT OPERATION OF SOFTWARE WILL BE UNINTERRUPTED OR ERROR-FREE AND/OR THAT DEFECTS IN THE SOFTWARE WILL BE CORRECTED. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY LICENSEE. FURTHERMORE, LICENSOR DOES NOT WARRANT OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE OR RELATED DOCUMENTATION IN TERMS OF THEIR CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS, OR OTHERWISE. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY LICENSOR SHALL CREATE A WARRANTY OR IN ANY WAY INCREASE THE SCOPE OF THIS WARRANTY.

Liability Limitations. LICENSOR SHALL NOT BE LIABLE FOR ANY SPECIAL, INDIRECT, PUNITIVE, OR CONSEQUENTIAL DAMAGES RESULTING FROM USE OF SOFTWARE UNDER ANY CONTRACT, NEGLIGENCE, STRICT LIABILITY OR OTHER THEORY. LICENSOR'S CUMULATIVE LIABILITY FOR DAMAGES HEREUNDER, WHETHER IN AN ACTION IN CONTRACT, WARRANTY, TORT, NEGLIGENCE, STRICT LIABILITY, INDEMNITY, OR OTHERWISE, SHALL IN NO EVENT EXCEED THE AMOUNT OF LICENSE FEES PAID BY THE LICENSEE FOR THE SOFTWARE LICENSED UNDER THIS AGREEMENT. THESE LIMITATIONS SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

LICENSEE Indemnity. LICENSEE agrees to indemnify and defend LICENSOR, and hold it harmless from all costs, including attorney's fees, arising from any claim that may be made against LICENSOR by any third party as a direct or indirect result of any use by LICENSEE of the Software,

Termination. This Agreement and the license may be terminated without fee reduction (i) by LICENSEE without cause on thirty (30) days notice; (ii) by LICENSOR, in addition to other remedies, if LICENSEE is in default and fails to cure within ten (10) days following notice; (iii) on notice by either party hereto if the other party ceases to do business in the normal course, becomes insolvent, or becomes subject to any bankruptcy, insolvency, or equivalent proceedings. Upon termination for any reason, LICENSEE shall immediately return Software and all copies to LICENSOR and delete all Software and all copies from the Hardware.

Law. This Agreement shall be governed by the laws of the Commonwealth of Massachusetts, exclusive of its conflicts of laws provisions and without regard to the United Nations Convention on Contracts for the International Sale of Goods, and any suit under this Agreement shall exclusively be brought in a federal or state court in Massachusetts. Any action against LICENSOR under this Agreement must be commenced within one year after such cause of action accrues.

Government End Users. This provision applies to all Software acquired directly or indirectly by or on behalf of the United States Government. The Software is a commercial product, licensed on the open market at market prices, and was developed entirely at private expense and without the use of any U.S. Government funds. If the Software is supplied to the Department of Defense, the U.S. Government acquires only the license rights customarily provided to the public and specified in this Agreement. If the Software is supplied to any unit or agency of the U.S. Government other than the Department of Defense, the license to the U.S. Government is granted only with restricted rights. Use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c) of the Commercial Computer Software Restricted Rights clause of FAR 52.227-19.

Export Restriction. LICENSEE will not remove or export from the United States or the country originally shipped to by LICENSOR (or re-export from anywhere) any part of the Software or any direct product thereof except in compliance with applicable export laws and regulations, including without limitation, those of the U.S. Department of Commerce.

Miscellaneous. This Agreement contains the entire understanding of the parties and supersedes all other agreements, oral or written, including purchase orders submitted by LICENSEE, with respect to the subject matter covered in this Agreement The delay or failure of either party to exercise any right provided in the Agreement shall not be deemed a waiver. All notices must be in writing and shall be delivered by hand (effective when received) or mailed by registered or certified mail (effective on the third day following the date of mailing). The notices addressed to LICENSOR shall be sent to its address set out above. If any provision is held invalid, all others shall remain in force. LICENSEE may not assign, pledge, or otherwise transfer this agreement, nor any rights or obligations hereunder in whole or in part to any entity. Paragraph headings are for convenience and shall have no effect on interpretation. In the event that it is necessary to undertake legal action to collect any amounts payable or to protect or to defend against the unauthorized use, disclosure, distribution, of the Software hereunder and/or other violation of this Agreement, LICENSOR shall be entitled to recover its costs and expenses including, without limitation, reasonable attorneys' fees.

Table of Contents

1. Introducing Acronis Backup & Recovery 10

1.1. Getting started

Direct management

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

   Windows

   Start the console by selecting it from the start menu.

   Linux

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

   ```
   trueimageremote
   ```

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. 14).

For how to enable the non-root users to start the console under Linux see Privileges for local connection (p. 78).

For how to enable the remote connection to a machine running Linux see Privileges for remote connection in Linux (p. 79).

Centralized management

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

To start with the centralized management:

1. Install Acronis Backup & Recovery 10 Management Server (p. 24).
2. 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 server administrator 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 Windows console flavor if you have a choice between Windows and Linux. Install Acronis Bootable Media Builder along with the console.
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. 315) 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. 317). 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 in the **Navigation** tree **Machines - All machines** 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, 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 locations**

If you opt for storing all backup archives in a single or a few networked locations, create shortcuts to these locations. After a shortcut is created, you can view and administer the location content under **Archive locations - Centralized locations** - Location name in the **Navigation** tree. The shortcut will be deployed to all the registered machines. The shortcut 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 the Acronis Backup & Recovery Storage Node (p. 24).
  2. Create the deduplicating managed location on the storage node.
  3. Ensure that the backup plans you create use the managed location 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 the 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. 63).

The comprehensive example of the advanced centralized management is provided in the section Setting up the centralized data protection in a heterogeneous network (p. 61).

1.1.1. Using the management console

As soon as the console connects to a managed machine (p. 324) or to a management server (p. 324), the respective items appear across the console's workspace (in the menu, in the main area with the **Welcome** view, the **Navigation** pane, the **Actions and tools** pane) enabling you to perform the agent-specific or server-specific operations.

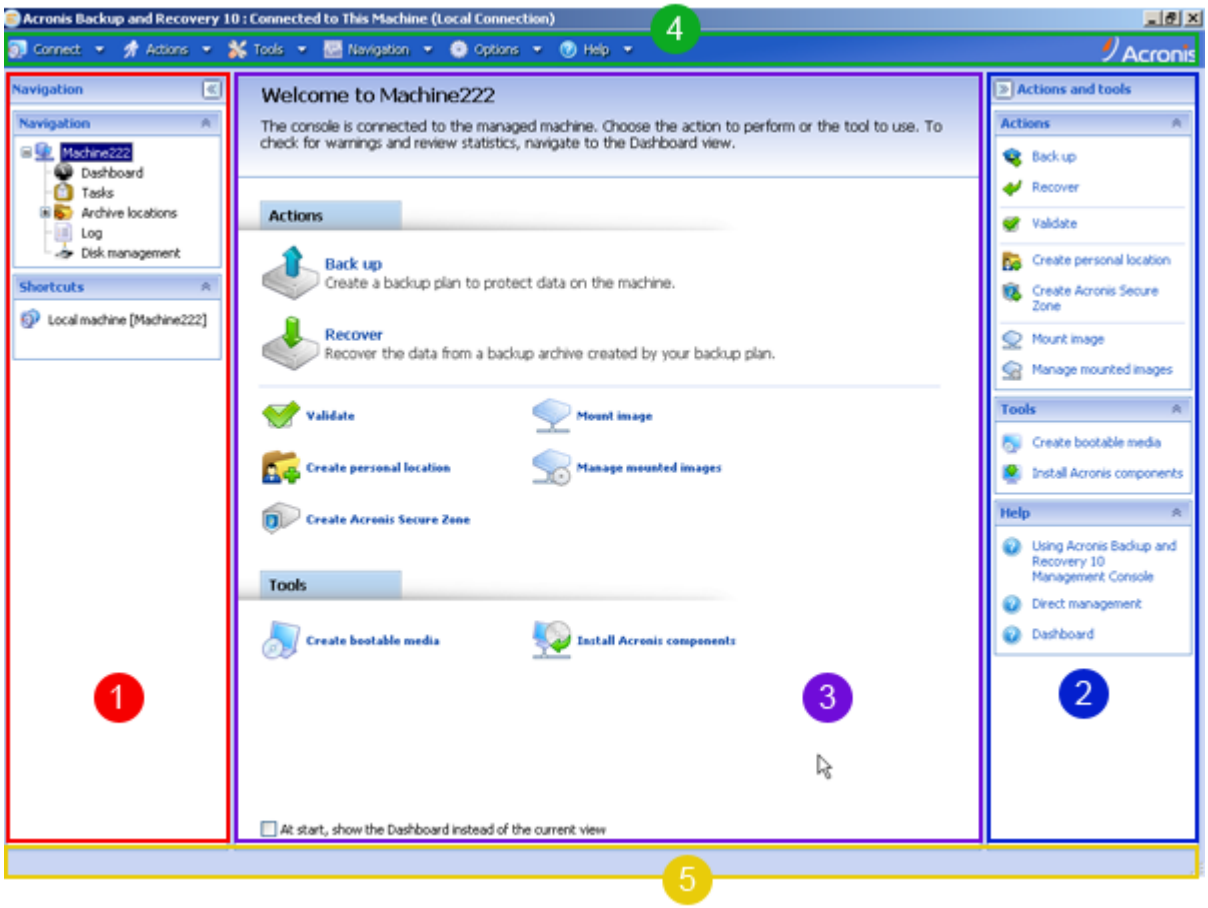

Acronis Backup & Recovery 10 Management Console - Welcome view

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. 15) section.)
2	**Actions and tools** pane	Contains bars with set of actions and tools that can be performed (see the Actions and Tools pane (p. 16) 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. 18) according 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 Some menu items are hidden until an item is not selected in the **Navigation** tree or in the view.
5	**Status** bar	Appears at the bottom of the window. The status bar is divided into two parts. The left side briefly describes the selected operation; the right side indicates operation progress and results. If you double-click on the operation results, you will see the Log.

1.1.1.1. "Navigation" pane

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 to quick access 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.
  - **Archive locations**. Use this view to manage locations and archives stored in there, add new locations, rename and delete the existing ones, validate locations, explore backup content, mount backups as virtual drives, convert backups to virtual disks, 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 to quick access the main operations, available with 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.
  - **Backup policies**. Use this view to manage backup policies existing on the management server.
  - **Machines**. Use this view to manage machines: add new machines, arrange them into groups in order to apply policies.
  - **Tasks**. Use this view to manage tasks, run, edit, stop and delete tasks, monitor their states, examine task history.
  - **Archive locations**. Use this view to manage locations and archives stored in there, create new locations, rename and delete the existing ones, validate locations and explore backup content.
  - **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 location; 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 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 a 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]**.

1.1.1.2. "Actions and tools" pane

The **Actions and tools** pane enables you to easy 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 bars 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. 19). Items of different navigation views have their own set of actions. The bar's name changes in accordance with item you select. For example, if you select in the **Backup plans and tasks** view the backup plan named *System backup*, 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 navigation views.

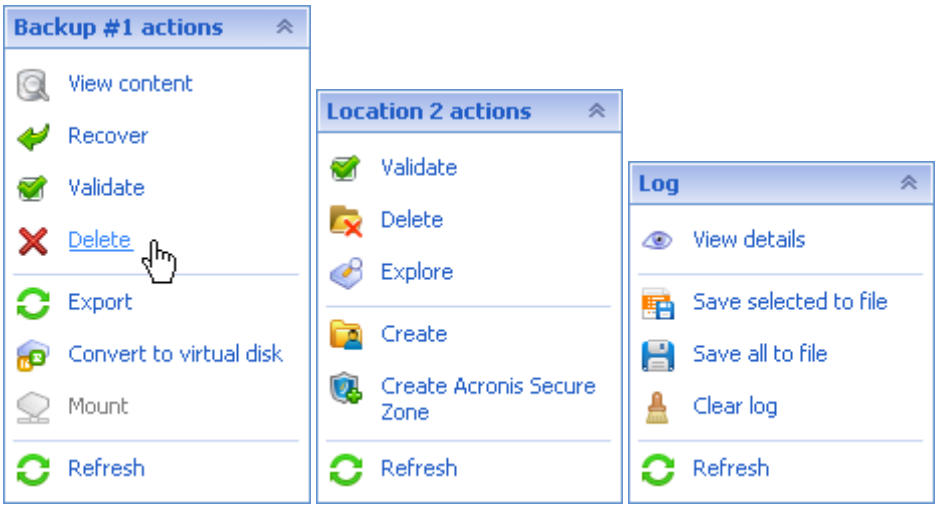

Examples of "'Item name' actions" bars

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. 19) section.)

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

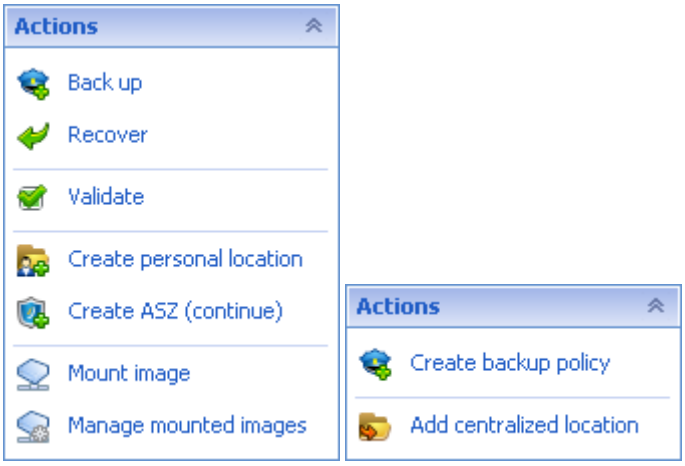

"Actions" bar on a managed machine and on a management server

Tools

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

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

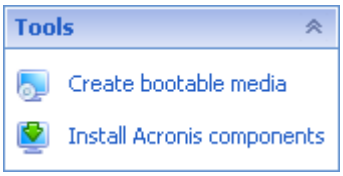

"Tools" bar

Help

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

1.1.1.3. Operations with panes

How to expand/minimize panes

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

How to change the panes' borders

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

1.1.1.4. Main area, views and action pages

Main area is a basic place of your working 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. 15).

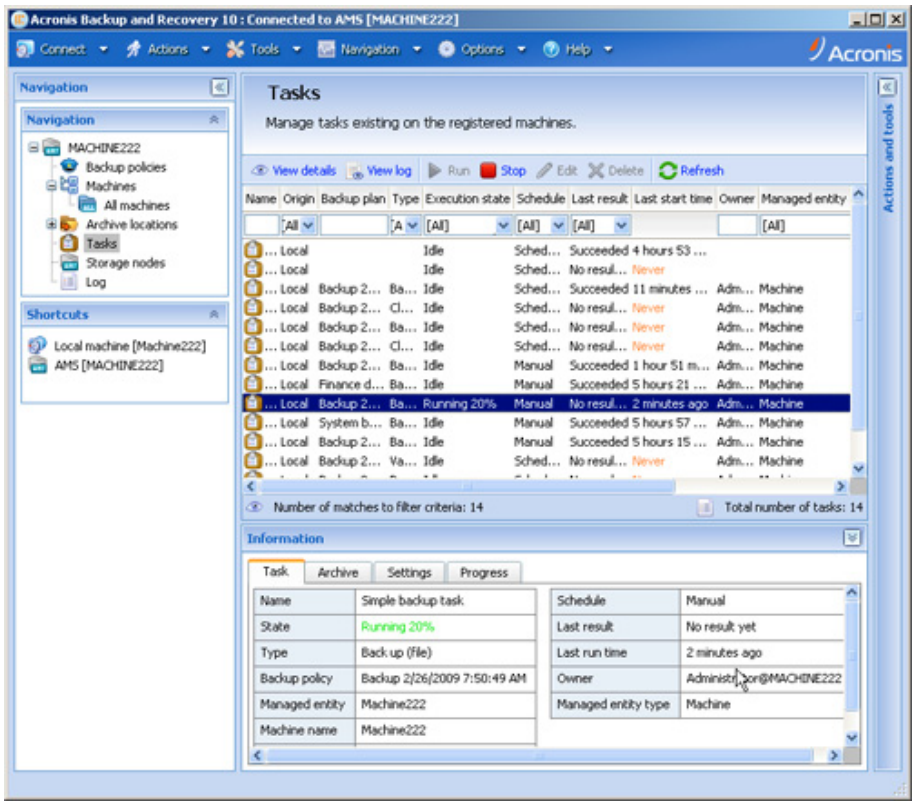

"Tasks" view

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 items'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 on 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.

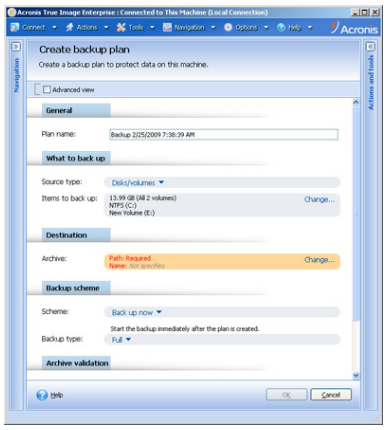

Action page - Create backup plan

Using controls and specifying settings

The action pages 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 check box for **Advanced view** at the top of the action page.

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

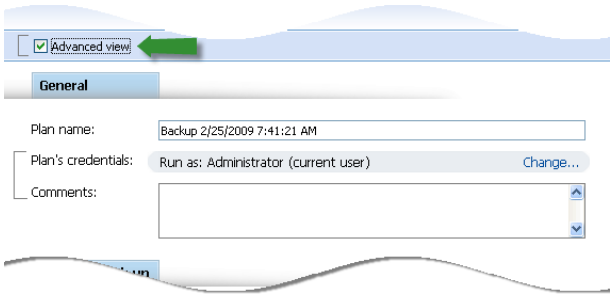

Action page - Controls

The Acronis Backup & Recovery 10 remembers changes you made on 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, on the **Actions** bar (as well as in the **Actions** menu) you will see the **Create backup plan (Continue)** item. Clicking this item brings you back to the page where you started the plan creation. Thus, you can perform the steps remained and accomplish the backup plan or the task creation.

"Actions" bar - Continue operation

1.2. Acronis Backup & Recovery 10 components

1.2.1. Agent for Windows

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

Disk backup

The disk-level data protection is based on backing up either the disk or 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. You can recover disks or volumes as a whole from such backup, as well as individual folders or files.

File backup

The file-level data protection is based on backing up the user's selection of files and folders. You can recover all files and folders that were backed up or select which of them to recover.

1.2.1.1. Bootable components

The bootable components of the agent enable operations with reboot, such as recovery of the volume containing the currently active operating system. Once the operations are completed, the machine boots into the operating system again. You can choose not to install the bootable components and perform operations that require reboot using bootable media.

1.2.1.2. Acronis Disk Director Lite

Acronis Disk Director Lite enables disk management operations such as cloning disks; converting disks; creating, extending, formatting and deleting volumes; removing or breaking a volume mirror; and some additional operations like changing a disk partitioning style between MBR and GPT or changing a disk label. The disk management operations can be performed either in the operating system or using bootable media.

This application is supplied free with the Acronis Backup & Recovery 10 Agent for Windows.

1.2.1.3. Agent for Hyper-V

Acronis Backup & Recovery 10 Agent for Hyper-V protects virtual machines residing on virtualization servers.

The agent installs on Windows 2008 Server (any edition) or Microsoft Hyper-V Server 2008 as an add-on to the Acronis Backup & Recovery Agent for Windows (p. 21).

Integration services have to be installed on the guest systems.

Where to go from here

Backing up virtual machines (p. 50)

Creating a backup plan (p. 167)

1.2.2. Agent for Linux

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

Disk backup

The disk-level data protection is based on backing up either the disk or 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. You can recover disks or volumes as a whole from such backup, as well as individual directories or files.

File backup

The file-level data protection is based on backing up the user's selection of files and directories. You can recover all files and directories that were backed up or select which of them to recover.

1.2.2.1. Bootable components

The bootable components of the agent enable operations with reboot, such as recovery of the volume containing the currently active operating system. Once the operations are completed, the machine boots into the operating system again. You can choose not to install the bootable components and perform operations that require reboot using bootable media.

1.2.2.2. Agent for ESX

Acronis Backup & Recovery 10 Agent for ESX protects virtual machines residing on virtualization servers.

The agent installs on VMware ESX Infrastructure 3.5 Update 2 as an add-on to the Acronis Backup & Recovery Agent for Linux (p. 21).

VMware Tools have to be installed on the guest systems.

In the future, the Agent for ESX will operate on the ESX server as virtual appliance and so it will be possible to install it without the Agent for Linux.

Where to go from here

Backing up virtual machines (p. 50)

Creating a backup plan (p. 167)

1.2.3. 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.

Using the direct console-agent connection, the administrator sets up and manages local backup plans, recovers data, creates and explores personal archive locations and performs other direct management operations.

Having connected the console to the management server, the administrator sets up and manages backup policies and accesses other functionality of the management server, that is, performs centralized management.

The console installation package includes **Acronis Bootable Media Builder.**

Flavors

The console has two flavors for installation on Windows and installation on Linux. Both flavors enable connection to any Acronis Backup & Recovery 10 Agent and Acronis Backup & Recovery 10 Management Server. The difference is as follows:

- the Active Directory-related features, such as browsing the AD, are not available when the console is running on Linux
- the Windows console flavor includes **Acronis WinPE ISO Media Builder** while the Linux console flavor does not.

1.2.3.1. Acronis Bootable Media Builder

Acronis Bootable Media Builder is a dedicated tool for creating bootable media (p. 317) with bootable Acronis components based on the Linux kernel.

Flavors

The media builder of the **Windows console flavor** creates bootable media that represents volumes and network in the Windows-like style (C:, D:, \\server\share.)

The media builder of the **Linux console flavor** creates bootable media that represents volumes and network in the Linux-like style (hda1, sdb2, smb://server/share.)

1.2.3.2. Acronis WinPE ISO Builder

Acronis WinPE ISO Builder is a dedicated tool for creating bootable media (p. 317) based on Windows preinstallation environment.

1.2.4. 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.

1.2.4.1. Acronis Backup & Recovery 10 Management Server

Acronis Backup & Recovery 10 Management Server is the central server that manages data protection within the local network. The management server provides the administrator with:

- a single entry point to the Acronis Backup & Recovery 10 infrastructure
- easy way to protect data on numerous machines (p. 323) using backup policies (see "Backup policy (Policy)" on page 316) and grouping
- enterprise-wide monitoring functionality
- ability to create centralized locations (p. 317) for storing enterprise backup archives (p. 315)
- ability to manage storage nodes (p. 325).

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

The management server's databases

The management server uses two Microsoft SQL instances for storing the configuration information and the centralized log. When installing a management server, you can choose between:

1. Using 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 several databases will be created on the machine.
2. Using Microsoft SQL Server 2008 (any edition) previously installed on any machine.
3. Using Microsoft SQL Server 2005 (any edition) previously installed on any machine.

This provides flexibility to store, say, the configuration on a locally installed server instance, while the logging information is placed to another database configured on a different machine.

1.2.4.2. Acronis Backup & Recovery 10 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 locations.)

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 node being able to manage up to 20 locations. The administrator controls the storage nodes centrally from the Acronis Backup & Recovery 10 Management Server (p. 324). The 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. 325)) and create centralized locations (p. 317). When creating a centralized location, specify the path to the location, the storage node that will manage the location, and the management operations to be performed on the location.

Managed locations 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.)

The management operations are as follows.

Storage node-side cleanup and validation

Archives, stored in unmanaged locations, are maintained by the agents (p. 314) 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. 315). 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 and validation according to the schedule. To do so, the agent must be online. The further processing is performed by the storage node.

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

Deduplication

A managed location can be configured as deduplicating location. This means that identical data will be backed up to this location only once to minimize the network usage during backup and storage space taken by the archives. For more information, please see Deduplication (p. 319).

Encryption

A managed location can be configured so that every write to the location is encrypted and every read is decrypted transparently by the storage node, using a location-specific encryption key stored on the node server. In case the storage medium is stolen or accessed by unauthorized person, the malefactor will not be able to decrypt the location 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.

1.2.4.3. Acronis 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. 317) into the system that has to be booted
- during group operations, reduces time required for booting multiple machines as compared to using physical bootable media.

Where to go from here

Acronis PXE Server installation (p. 229)

Bootable Media Builder (p. 222)

Configuring Acronis PXE Server (p. 230)

1.2.4.4. Acronis License Server

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

1.3. Supported operating systems

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

- Windows Server 2000/Advanced Server 2000/Server 2003/ SBS/ Server 2008
- Windows XP Professional x64 Edition, Windows Server 2003/2008 x64 Editions
- Windows Vista all Editions
- Windows Professional 2000 SP4/XP Professional SP2

Acronis Backup & Recovery 10 Management Console

- Windows Server 2000/Advanced Server 2000/Server 2003/ SBS/ Server 2008
- Windows XP Professional x64 Edition, Windows Server 2003/2008 x64 Editions
- Windows Professional 2000 SP4/XP Professional SP2
- Windows Vista all Editions (except for Vista Home Basic and Vista Home Premium) (except for installation of Acronis components on remote machines running Vista and some specific features such as tape devices support)

Acronis Backup & Recovery 10 Agent for Windows

- Windows Server 2000/Advanced Server 2000/Server 2003/ Server 2008
- Windows XP Professional x64 Edition, Windows Server 2003/2008 x64 Editions
- Windows Professional 2000 SP4/XP Professional SP2
- Windows Vista all Editions

Acronis Backup & Recovery 10 Agent for Linux

- Linux with kernel 2.4.18 or later (including 2.6.x kernels) and glibc 2.3.2.

- Various Linux distributions (see your man for the exact list) among which support for the following Linux distributions is especially considered: SuSE 9.3, SuSE 10.0, SLES 9.0, RedHat 9.0, Fedora Core 4, Fedora Core 5, Enterprise Server 3.0, Mandrake 10.x, Slackware 10, Debian (Sarge), ASPLinux 10, Gentoo, Ubuntu 4.10.
- x64 versions of the above Linux distributions and other Linux distributions are also supported.

  The agent for Linux is in fact a 32-bit executable. For the 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 Scientific Linux, 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.
- VMware ESX Infrastructure 3.5 Update 2.

Acronis Backup & Recovery 10 Agent for Hyper-V

- Windows Server 2008 with Hyper-V (x64)

Acronis Backup & Recovery 10 Agent for ESX

- VMware Infrastructure

Acronis Universal Restore

- Windows Server 2000/Advanced Server 2000/Server 2003/ SBS/ Server 2008
- Windows XP Professional x64 Edition, Windows Server 2003/2008 x64 Editions
- Windows Professional 2000 SP4/XP Professional SP2
- Windows Vista all Editions

Bootable version of Acronis Backup & Recovery 10 enables disk backup and recovery on a machine running any PC-based operating system.

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
- ReiserFS3 - particular files cannot be recovered from volume backups
- XFS - volume recovery without the volume resize ability, particular files cannot be recovered from volume backups
- JFS - particular files cannot be recovered from volume backups
- Linux SWAP

Acronis Backup & Recovery 10 can back up and recover corrupted or non-supported file systems, such as ReizerFS4, using the sector-by-sector approach.

1.5. Technical support

As part of a purchased annual Support charge you are entitled to Technical Support as follows: to the extent that electronic services are available, you may electronically access at no additional charge, Support services for the Software, which Acronis shall endeavor to make available twenty four (24) hours a day, seven (7) days per week. Such electronic services may include, but are not limited to: user forums; software-specific information; hints and tips; bug fix retrieval via the internet; software maintenance and demonstration code retrieval via a WAN-accessible FTP server; and access to a problem resolution database via Acronis customer support system.

Support shall consist of supplying telephone or other electronic support to you in order to help you locate and, on its own, correct problems with the Software and supplying patches, updates and other changes that Acronis, at its sole discretion, makes or adds to the Software and which Acronis makes generally available, without additional charge, to other licensees of the Software that are enrolled in Support.

Upon mutual agreement by both parties, Acronis shall:

(i) supply code corrections to you to correct Software malfunctions in order to bring such Software into substantial conformity with the published operating specifications for the most current version of the Software unless your unauthorized modifications prohibit or hamper such corrections or cause the malfunction;

or (ii) supply code corrections to correct insubstantial problems at the next general release of the Software.

More information about contacting Acronis Technical Support is available at the following link: http://www.acronis.com/enterprise/support/ http://www.acronis.com/enterprise/support/

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. 314) on the machine which becomes a managed machine (on page 324) from this point on.
2. To be able to manage the machine using Graphical User Interface, install Acronis Backup & Recovery 10 Management Console (see "Console (Acronis Backup & Recovery 10 Management Console)" on page 318) 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 (on page 317).
4. Connect the console to the managed machine.
5. Create a backup plan.

   To do so, you have to specify, at the very least, the data to be protected and the location (see "Archive location" on page 314) where the backup archive (see "Backup archive (Archive)" on page 315) will be stored. This will create a minimal backup plan consisting of one task (on page 326) that will create a full backup (on page 314) 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. 32), perform archive maintenance operations such as backups validation (on page 327) or deleting outdated backups (archive cleanup (on page 318).) You can customize backup operations using various backup options, such as pre/post backup commands, network bandwidth throttling, error handling or notification options.

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

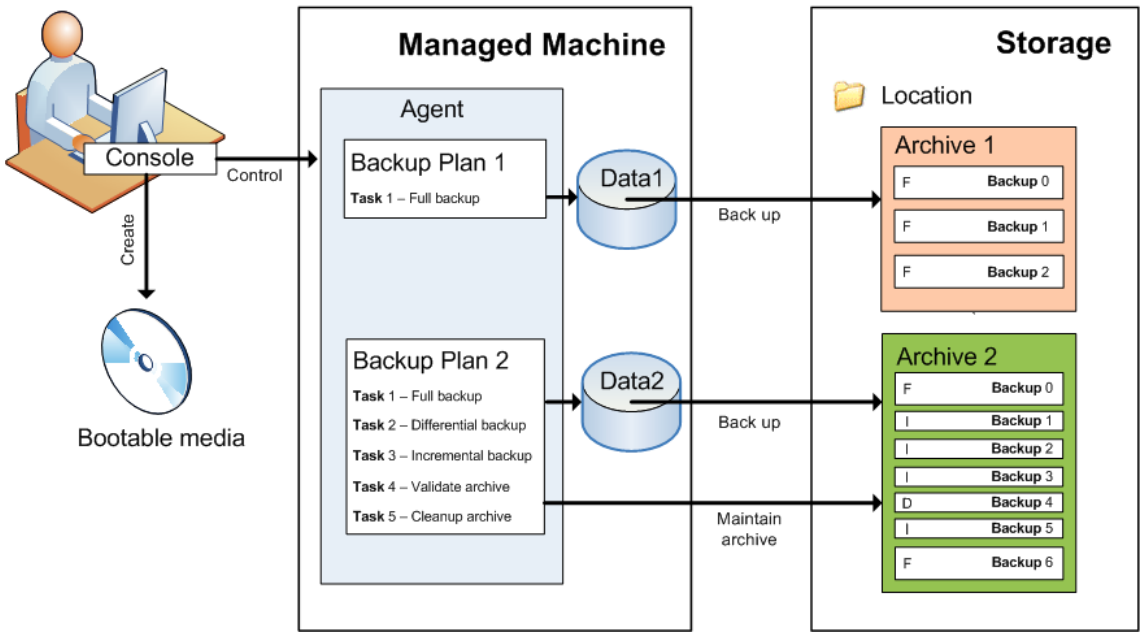

Backup using bootable media

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.

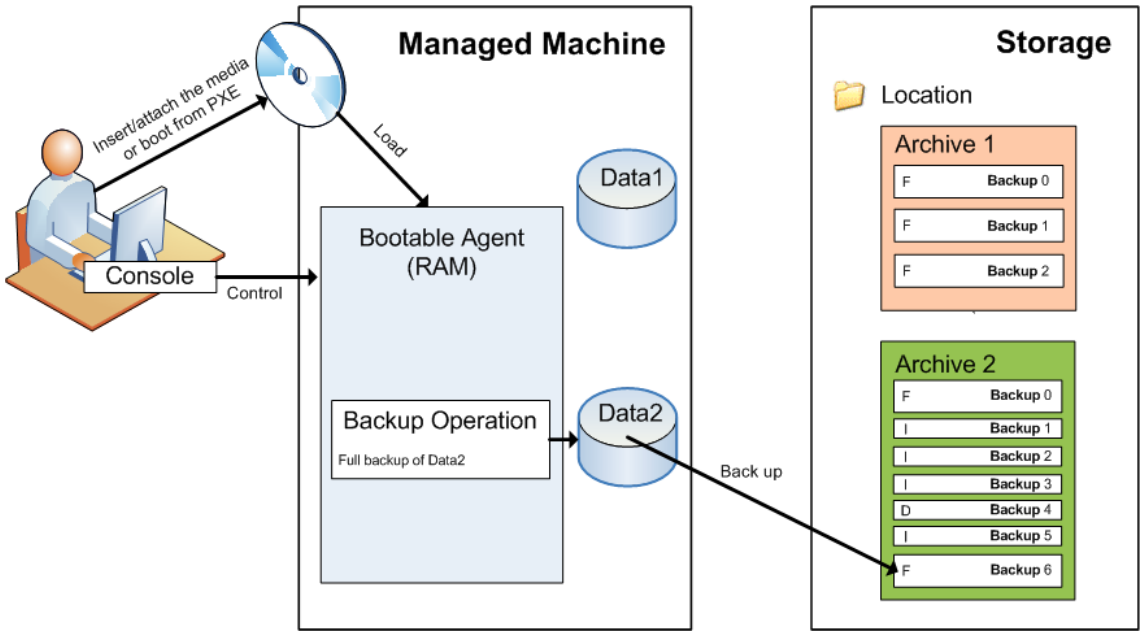

Recovery under operating system

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

Examples of the exceptions from this rule:

1. 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)
2. 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 takes place. But you can connect the console to another machine and configure a recovery operation on that machine, if need be. This ability (the remote parallel recovery) first appeared in Acronis Backup & Recovery 10, the previous product versions do not provide it.

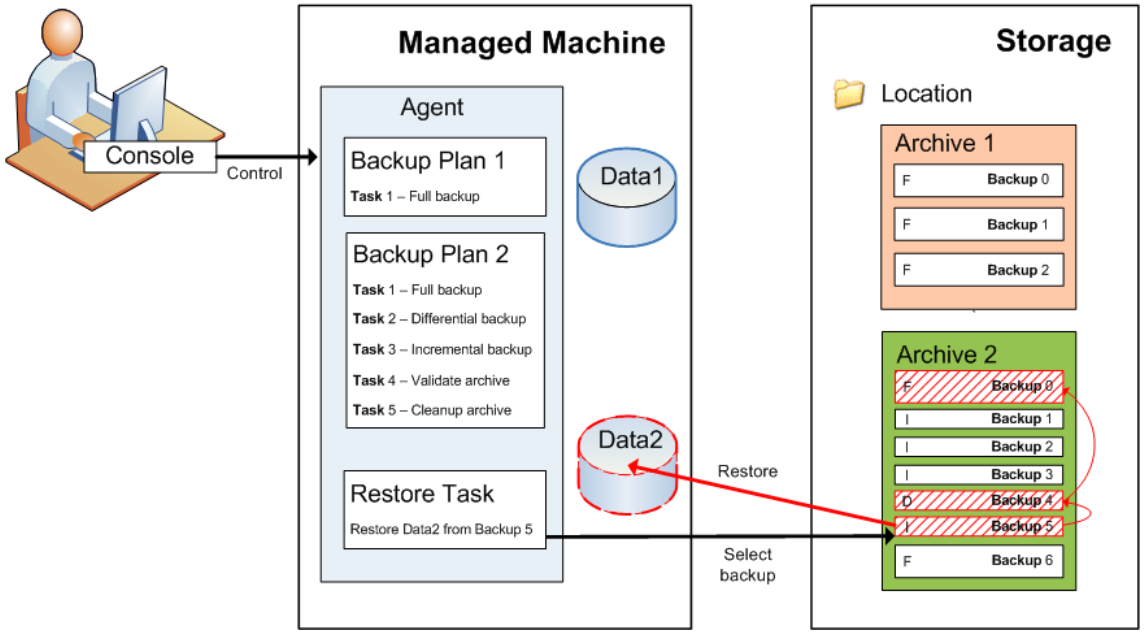

Recovery using bootable media

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

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

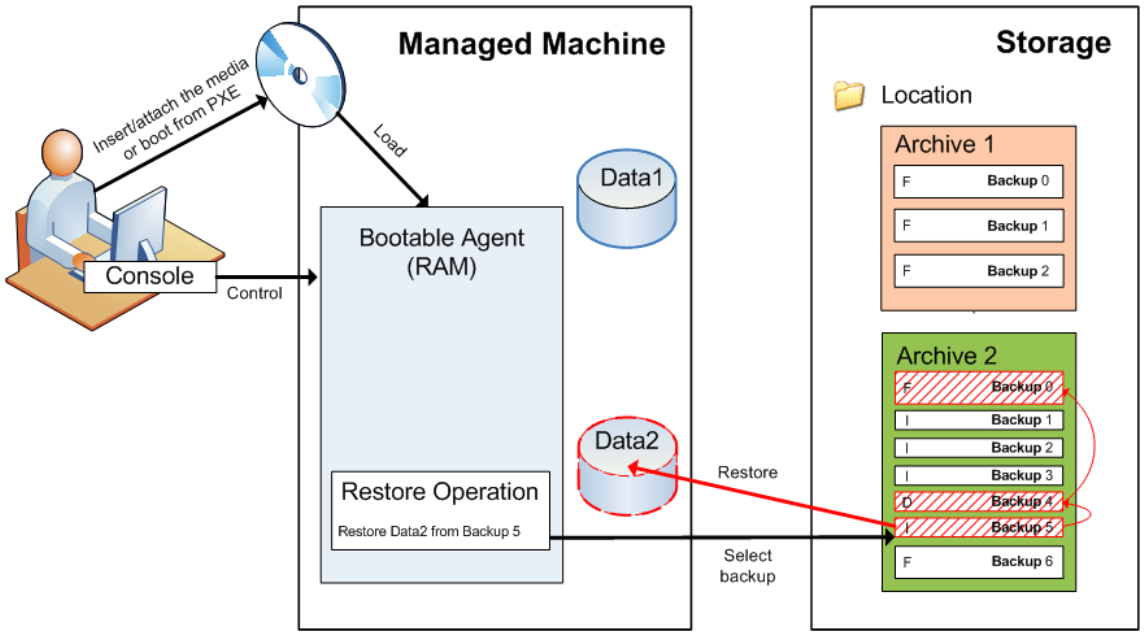

2.2. Full, incremental and differential backups

Acronis Backup & Recovery 10 provides capability to use popular backup schemes, such as Grandfather-Father-Son and Tower of Hanoi, as well as create custom backup schemes. All the backup schemes are based on full, incremental and differential backup methods. The "scheme" term in fact denotes the algorithm of applying these methods plus the algorithm of the archive cleanup.

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

Full backup

A full backup contains 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 the data from a full backup) and has the shortest recovery time as compared to incremental or differential ones. A full backup is most useful when:

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

Synthetic full backup

You have an option to minimize the resources required for creating full backups by backing up data changes with subsequent automated consolidation. For more information see Additional settings (p. 115).

Incremental backup

An incremental backup stores changes to the data against the latest backup. This kind of backup takes minimal storage space, but 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 possibility to roll back to any one of multiple saved states
- the data changes tend to be little as compared to the total data size.

The common knowledge says that incremental backups are less reliable than full ones because if one backup in the "chain" is corrupted, the next ones cannot be used anymore. However, you will have to store multiple full backups if you need multiple prior versions of your data and the reliability of an oversized archive is even more questionable. Differential backups are said to be the medium solution which might mean combination of shortcomings as well as combination of advantages.

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 this 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 little as compared to the total data size.

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

The following table summarizes advantages and shortcomings of each backup type as they appear from 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 criterion for choosing the optimal backup scheme.

Parameter	Full backup	Differential backup	Incremental backup
Storage space	Maximal	Middle	Minimal
Creation time	Maximal	Middle	Minimal
Recovery time	Minimal	Middle	Maximal

2.3. User rights 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.

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

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
- 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 operations and log files on the machine
- Manage local backup plans and tasks owned by any user registered in the operating system

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

2.4. Owners and credentials

This section explains the concept of owner and the meaning of the backup plan (or task) 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 the Administrator privilege 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 for editing a plan or task owned by another user, all passwords set in the task are cleared. This prevents the "modify settings, leave passwords" trick that enables a user access data or location that they normally cannot access. A user, or rather an intruder, cannot, say, redirect a database backup or recover a database to the location controlled by the intruder. The program displays a warning each time you are trying to edit a plan (task) last modified by another user. On seeing the notification, you have two options:

- Click "Cancel" and create your own plan or task. The original task will stay 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 a destination location. Exactly, this is the user whose account was specified when creating the backup plan in the Destination step. By default, the plan credentials are used.

Plan (task) credentials

Any task running on a machine runs on behalf of some user. When creating a plan or a task, you have an 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 **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 the 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 check box for **Advanced view**.
2. Select **General -> Plan (Task) account -> 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 the plan (task) credentials after clicking the final **OK**.

Why does the program compel me to specify credentials?

A scheduled or postponed task has to run anyway, no matter if no user is logged on (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 log on and start the task.

2.5. GFS backup scheme

Backup scheme defines when and how often to back up the data and (optionally) retention rules and when to apply the rules, that is, the cleanup schedule. This section covers implementation of the Grandfather-Father-Son (GFS) backup scheme in Acronis Backup & Recovery 10.

GFS overview

GFS is a popular backup scheme aimed to keep up the optimal balance between a backup archive size and the number of recovery points available from the archive. GFS enables recovering with daily resolution for last several days, weekly resolution for last several weeks and monthly resolution for any time in the past.

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 lifetimes for the daily, monthly and weekly backups. The daily backups are referred as “sons”; weekly backups are referred as “fathers”; the longest lived monthly backups are called “grandfathers”.

GFS as a tape rotation scheme

GFS was initially created and is often referred 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 and roll-back period
- which of the tapes you should overwrite with the next coming 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. Some useful links:

http://www.filevaultusa.com/tape-rotation-schemes.php

http://www.jack-of-all-trades.org/?m=200607.

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

The following is the explanation of the settings 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 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”.

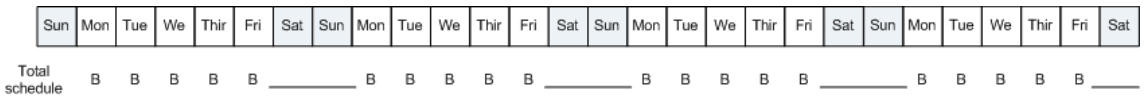

The total schedule.
Schedule: Workdays at 8:00 PM

Weekly/Monthly

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

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

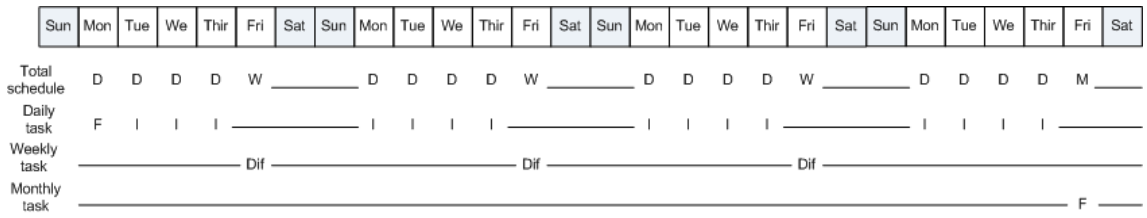

Backup tasks created according to the GFS scheme by Acronis Backup & Recovery 10.
Schedule: Workdays at 8:00 PM
Weekly/Monthly: Friday

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 the age 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 the age you specify. The weekly backups' lifetime cannot be less than the daily backups' lifetime. It is usually set several times larger.

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 the age you specify. The monthly backups' lifetime cannot be less than the weekly backups' lifetime. It is usually set several times larger. You have the option to keep the monthly backups infinitely.

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 week. For each day of 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.

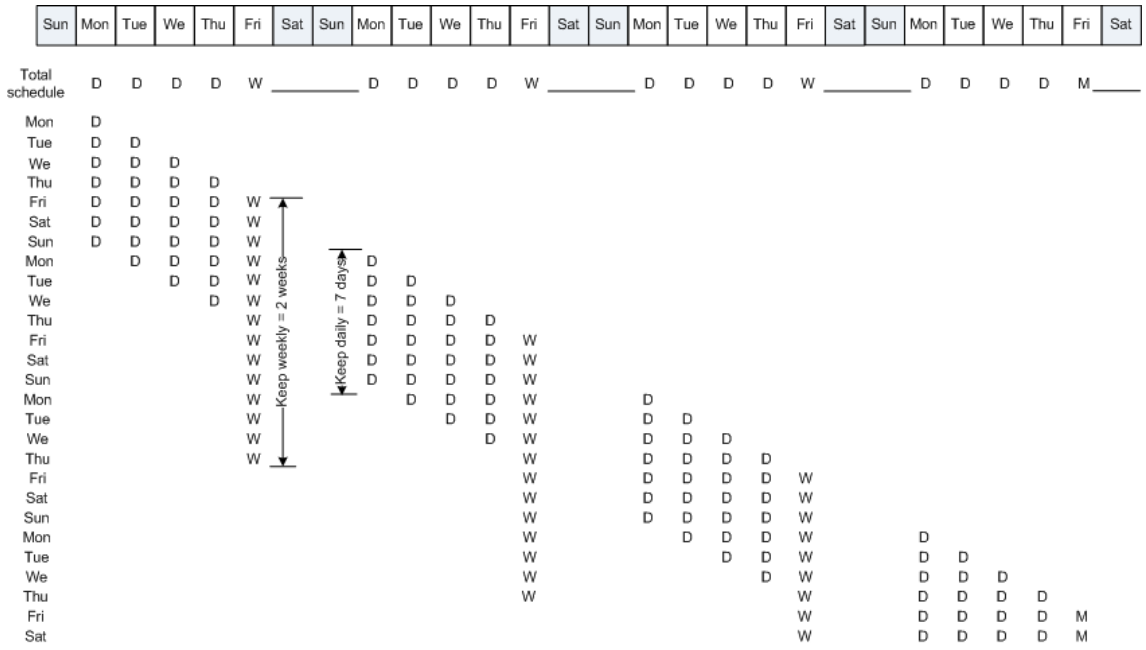

An ideal archive created according to the GFS scheme.
Schedule: Workdays at 8:00 PM
Weekly/Monthly: Friday
Keep daily backups: 7 days
Keep weekly backups: 2 weeks
Keep monthly backups: 6 months

Starting from the third week, weekly backups will be regularly deleted. After 6 months, monthly backups will start to retire. The diagram for weekly and monthly backups will look similar against 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 base on this backup. Regular consolidation is unacceptable because it takes too much of 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.

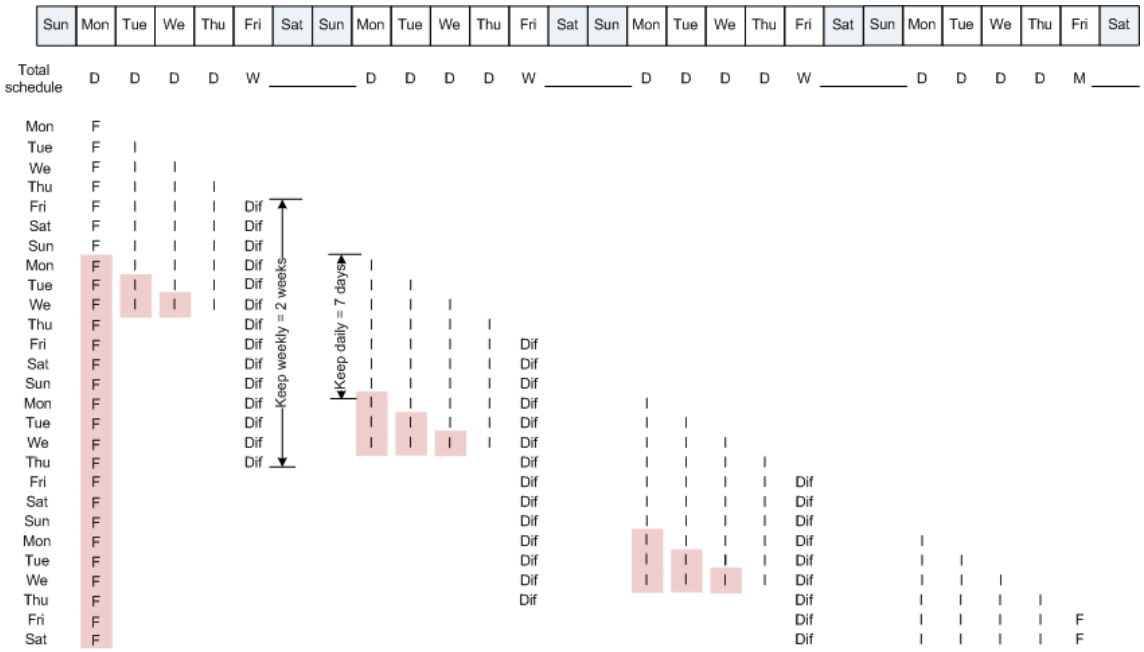

An archive created according to the GFS scheme by Acronis Backup & Recovery 10.
Schedule: Workdays at 8:00 PM
Weekly/Monthly: Friday
Keep daily backups: 7 days
Keep weekly backups: 2 weeks
Keep monthly backups: 6 months

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. But it is allowed to move only one ring at a time, and prohibited to place a larger ring above a smaller ring. The solution is to shift the first ring every other move (moves 1, 3, 5, 7, 9, 11...), the second ring at intervals of four moves (moves 2, 6, 10...), the third ring at intervals of eight moves (moves 4, 12...), and so on.

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

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

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

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

The figure below shows the 16-session pattern used in the five-level ToH backup scheme:

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

The Tower of Hanoi backup scheme implies keeping only one backup per level. All the outdated backups have to be deleted. So the scheme allows for efficient data storage: more backups accumulate towards recent time. Having four backups, we could recover data as of today, yesterday, half a week, or a week ago. For five-levels scheme we could recover data backed up two weeks ago as well. So every additional backup level doubles a period we can go back in for a data.

Tower of Hanoi by Acronis

The Towers 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 for any type of backup destination. You can set up the backup scheme while creating backup plan.

Acronis implementation for the scheme have the following features:

- up to 16 backup levels
- incremental backups on first level (A) - to get time and storage saving for the most frequent backup operations; but a data recovery from such backups is more durational because it generally requires access to three backups
- full backups on the last level (E for five-level pattern) - the most rare backups in the scheme, get more time and occupy more size in storage
- differential backups on all intermediate levels (B, C and D for five-level pattern)
- since the very first session backup is incremental one that cannot exist without a previous full backup, so a full backup is created instead this incremental one at the session
- the scheme forces that every backup level has to keep only recent backup, other backups from the level have to be deleted; however a backup deletion is postponed in cases the backup is base for another incremental or differential one
- old backup on a level is kept until new backup has been successfully created on the level - to secure storing of old backup when new backup creating is failed

The pattern circle for the five-level backup scheme in the implementation by Acronis consists of 16 sessions. In the table the circle includes sessions with numbers from 1 till 16:

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

As the result of using incremental and differential backups it is possible the situation 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 the moment of 25 session because the differential backup (D) created at session 9 is still actual and is based on this backup. In the table all cells with deleted backups are grayed:

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

Note that the previous differential backup on level 4 (D) is kept in backup archive at the moment and will be deleted only when a new backup on the level is created successfully.

So 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.

To get information on using the Tower of Hanoi backup scheme for tape libraries see the Tape rotation section.

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 of the backups.

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.

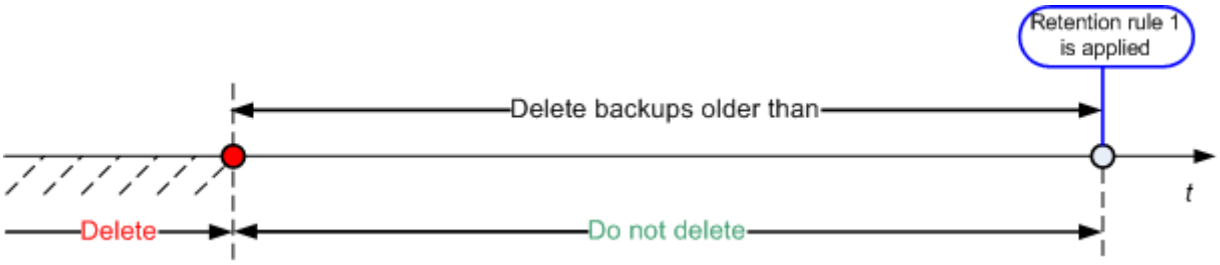

2. Keep the archive size within

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.

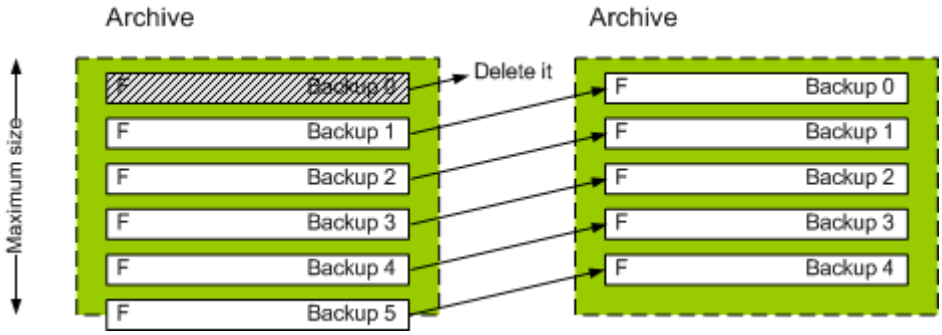

There is a certain risk that all backups but one will be deleted if the maximum archive size is set improperly (too small) or a regular backup turns out to be too large. To protect the recent backups from deletion, select the box for **Never delete backups younger than** and specify the maximum age of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

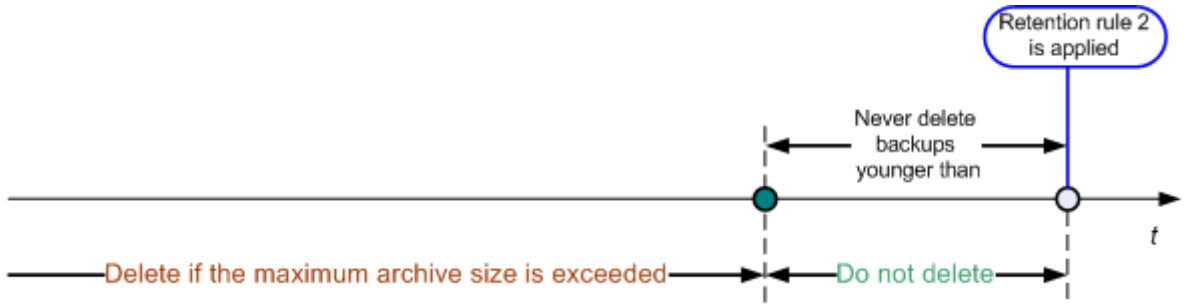

Combination of rules 1 and 2

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

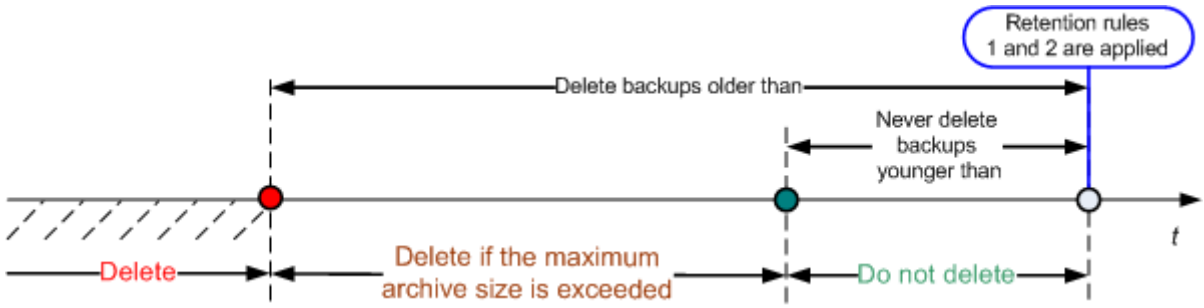

Example

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, 4 weeks) 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 location 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 the deletion. The resulting backup will not contain data that was present in the deleted backup and was absent in 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 the reasonable reserve. Then 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 Locations (p. 125) page provides you with information about free space available in each location. Check this page from time to time. If the free space (which in fact is the storage device free space) comes near to zero, you might need to toughen the restrictions for some or all archives residing in this location.

2.8. Support for logical volume management

2.8.1. Microsoft LDM (Dynamic volumes)

This section explains in outline how you would back up and recover dynamic volumes (see "Dynamic volume" on page 322) 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 (see "Dynamic disk" on page 320), or more exactly, on a disk group (on page 320). Acronis Backup & Recovery 10 supports the following dynamic volume types: simple/spanned, striped, mirrored, RAID 5 and RAID 0+1.

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 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 volumes DYN1 and DYN2 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.

Dynamic volumes recovery

A dynamic volume can be recovered

- over an existing volume of any type
- 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 the target volume (basic, simple/spanned, striped, mirrored, RAID 5, RAID 0+1) will not change. The target volume size has to be enough to accommodate the backup content.

Recovery to a 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 also matters.

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 some amount of unallocated space. The total size of unallocated space is 40GB. You will not have problems if the unallocated space is distributed evenly among the disks (20GB and 20GB.)

If one of the volumes has 10GB and the other has 30GB you cannot be sure because you need at least 15GB on each disk to enable the data to be distributed evenly. You only have chance if the data takes less than 20GB of the striped volume. One disk will hold, say, 10GB, the other will hold the remaining 10GB, but there will be no free space on the recovered volume.

The following table shows the types of volumes resulting from the recovery.

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

Moving and resizing volumes during recovery

You can resize the resulting basic volume, both MBR and GPT, during recovery, or change the volume's location on the disk. A dynamic resulting 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 5, RAID 0+1.) 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 more about Acronis Disk Director Lite, see the Disk management (p. 231) section.

2.8.2. HP LVM (Linux)

This section explains in outline how you would back up and recover volumes managed by Logical Volume Manager (LVM) using Acronis Backup & Recovery 10.

Acronis Backup & Recovery 10 Agent for Linux can access, back up and recover such volumes when running in Linux with 2.6.x kernel. On a machine booted from bootable media, LVM volumes can be backed up and recovered sector-by-sector only, because, as contrasted to the Microsoft LDM, the LVM volumes configuration (the Logical Volume Manager metadata) is contained within the operating system and cannot be accessed while the operating system is offline.

You can back up data of one or more LVM volumes and recover it to a previously created LVM volume or a basic MBR disk or volume, likewise it is also possible to recover MBR volume data to an LVM volume. In each case, the program stores and recovers volume contents only. The type or other properties of the target volume will not change.

A system, recovered from an LVM volume backup to a basic MBR disk, cannot boot because its kernel tries to mount the root file system at the LVM volume. To boot the system, change the loader configuration and /etc/fstab so that LVM is not used and reactivate your boot loader as described in the Bootability recovery section.

When recovering an LVM volume over a basic MBR volume, you can resize the resulting volume.

How to select LVM volumes to back up

LVM volumes appear at the end of the list of volumes available for backup. Basic volumes included in LVM volumes are also shown in the list with None in the Type column. If you select to back up such partitions, the program will image it sector-by-sector. Normally it is not needed. To back up all available disks, specify all dynamic volumes plus basic volumes not belonging to them.

Here is an example of a volumes list obtained with the --list command. The GUI displays a similar table.

```
Num  Partition             Flags       Start      Size      Type
_____________________________________________________________________
Disk 1:
1-1  hda1 (/boot)         Pri,Act     63         208782    Ext3
1-2  hda2                 Pri         208845     8177085   None
Disk 2:
2-1  hdb1                 Pri,Act     63         8385867   None
Disk 3:
3-1  hdd1                 Pri,Act     63         1219617   Ext3
3-2  Acronis Secure Zone  Pri         1219680    2974608   FAT32

Dynamic & GPT Volumes:
DYN1 VolGroup00-LogVol00                         15269888  Ext3
DYN2 VolGroup00-LogVol01                         1048576   Linux Swap
```

The system has three physical disks (1, 2, 3). Two LVM volumes DYN1 and DYN are arranged across basic volumes 1-2 and 2-1. Hard drive 3 includes Acronis Secure Zone which is not normally backed up.

To back up dynamic volume DYN1, select volume DYN1.

To back up all three physical drives, select volumes 1-1, 3-1, DYN1, DYN2.

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

Helpful link:

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

2.9. Backing up md and block devices (Linux)

Acronis Backup & Recovery 10 Agent for Linux can back up and recover software and hardware RAID arrays.

Software RAID arrays

Software RAID arrays under Linux combine several basic volumes and make solid block devices (/dev/md0, ... /dev/md31), information of which is stored in /etc/raidtab or in dedicated areas of that volumes.

Backup

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

Basic volumes included in software arrays are listed as if they had a corrupted file system or without a file system at all. Backing up such volumes does not make sense when a software array is mounted, as it won't be possible to recover them.

Example

Here is an example of a volumes list obtained with the **--list** command. The GUI displays a similar table.

The system has RAID-1 configured on two basic volumes: sdc1, sdd1.

```
Num  Partition             Flags       Start       Size        Type
---- -------------------- ----------- ---------- ---------- -----------
Disk 1:
1-1  sda1                 Pri,Act      63          208782     Ext3
1-2  sda2                 Pri          208845      15550920   ReiserFS
1-3  sda3                 Pri          15759765    1012095    Linux Swap
Disk 2:
    Table                             0                      Table
    Unallocated                       1           16771859   Unallocated
Disk 3:
3-1  sdc1                 Pri          63          16755732   Ext3
    Unallocated                       16755795    16065      Unallocated
Disk 4:
4-1  sdd1                 Pri          63          16755732   None
    Unallocated                       16755795    16065      Unallocated
Disk 5:
    Table                             0                      Table
    Unallocated                       1           16771859   Unallocated
Dynamic & GPT Volumes:
DYN1 md0                                          33511168   Ext3
                         Disk: 5      0           63
                         Disk: 4      0           63
```

You can back up the RAID array as follows:

```
aticmd --backup --volume:DYN1 --filename:/tmp/raid.tib --progress:on
```

In the Graphical User Interface you can select the check box for DYN1.

Recovery

Parameters of software disk arrays are not backed up, so they can only be recovered over a basic volume, to unallocated space, or to a previously configured array. Recovery can be performed in the operating system or using bootable media.

When started from bootable media, the bootable agent tries to access parameters of a software disk array and configure it. However, if the necessary information is lost, the array cannot be configured automatically. In this case, create a software array manually and restart the recovery procedure.

Hardware RAID arrays

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

Physical drives that are part of hardware disk arrays are 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.

2.10. Backing up virtual machines

Acronis Backup & Recovery 10 Virtual Edition for Hyper-V and Acronis Backup & Recovery 10 Virtual Edition for ESX enable backing up virtual machines from host.

Preparation

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

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

On VMware ESX Infrastructure 3.5 Update 2:

- Install the Agent for ESX on the ESX host.
- VMware Tools 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) or offline (stopped) or switch between the two states during backup.

A virtual machine has to be offline (stopped) during the recovery to this machine.

Virtual machine backup vs. physical machine backup

Backing up an entire virtual machine or its volumes yields a standard disk backup (p. 320). 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, recover disks and volumes from the backup to a physical machine.

The same way, 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, the physical to virtual and virtual to physical machine migration becomes available.

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 disks size) but numerous legacy servers such as ones resulting from workloads consolidation. A separate archive will be created for each machine.

Backing up volumes within a virtual machine is similar to backing up physical machine's volumes. With this source type, you select the machine and then select 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 not be backed up.

Guest operating systems

The following guest operating systems are supported.

Microsoft Windows platform:

- Microsoft Windows Server 2003 Standard (x86, x64)
- Microsoft Windows Server 2003 Enterprise (x86, x64)
- Microsoft Windows Server 2003 R2 Standard (x86, x64)
- Microsoft Windows Server 2003 R2 Enterprise (x86, x64)
- Microsoft Windows Small Business Server 2003 R2 (x86, x64)
- Microsoft Windows Server 2008 Standard (x86, x64)
- Microsoft Windows Server 2008 Enterprise (x86, x64)

Linux platform.

Guest HDD

The following virtual disk configurations are supported.

Partitioning style: MBR

Volume types: basic and dynamic volumes.

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 Application Event Log (Event ID = 8193).

Cause: This happens because there is no registry key:

```
HKEY_CLASSES_ROOT\Wow6432Node\CLSID\{F2C2787D-95AB-40D4-942D-298F5F757874}
```

Solution: Add this key to 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. Tape support

Acronis Backup & Recovery 10 supports tape libraries, autoloaders, SCSI and USB tape drives as storage devices.

A tape library (robotic library) is a high-capacity storage device containing one or more tape drives, multiple slots to hold tape cartridges, and one or more loaders (robotic mechanisms) intended for relocating the tape cartridges between the slots and the tape drives. No human concern is required to place a tape cartridge in a library in one of its drives. Small tape libraries with only one drive and one loader are known as autoloaders.

Access to tape 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. 24).

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.

Backup to a locally attached tape device

A tape drive that is locally attached to a managed machine can be used by local backup plans as storage device. The functionality of a locally attached autoloader or tape library is limited to the ordinary tape drive. Specifically the robotics are not supported. This means that the program can only work with the currently mounted tape and you have to mount tapes manually.

When creating a backup plan, you are able to select the locally attached tape device as the backup destination. The tape device appears in the list of locations under Local folders. 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 the 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.

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. Backups located on a tape cannot be validated.
5. Personal locations cannot be created on tape devices.
6. Because the presence of an operating system cannot be detected in a backup located on a tape, the Acronis Universal Restore (p. 326) is proposed at every disk or volume recovery, even when recovering a Linux or non-system Windows volume.
7. Acronis Active Restore (p. 313) is not available when recovering from a tape.

Recover 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 locations under Local folders and then select the backup. After recovery is started, you will be prompted for other tapes if the tapes are needed for recovery.

2.12. Proprietary Acronis technologies

2.12.1. Acronis Secure Zone

The Acronis Secure Zone is a hidden and secure partition that enables keeping backup archives on a managed machine disk space. Windows applications, except for dedicated disk management tools, cannot access the zone, and so the archives stored in the zone are protected from software malfunction and user errors.

Should the disk have a physical failure, the zone and the archives could be lost. That's why the 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 ordinary location is temporary unavailable or connected through a slow or busy channel.

Advantages

- The zone enables recovery of a disk to the same disk where the disk's backup resides
- The zone offers a cost-effective and handy method for protecting data from software malfunction, virus attack, operator error.

- The zone eliminates the need for a separate media or network connection to back up or recover the data. This is especially handy for mobile users.
- The zone enables using the Acronis Startup Recovery Manager (F11) to recover the operating system that cannot boot.
- The zone can serve as intermediate location when using dual destination backup.

Limitations

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

Managing the Acronis Secure Zone

In terms of backup archives management, Acronis Secure Zone is considered as personal location. Once created on a managed machine, the zone is always present in the list of Personal locations shortcuts. 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 location 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 a backup content
- export a backup from the zone
- 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 managing locations, see the Locations (p. 125) section.

Upgrade from Acronis True Image Echo

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

To upgrade the activated Acronis Startup Recovery Manager, deactivate it and activate again. No action is needed to upgrade the Acronis Startup Recovery Manager if it is not activated.

2.12.2. Acronis Startup Recovery Manager

A modification of the bootable agent (p. 316) can be put in the Acronis Secure Zone and configured to start at boot time on pressing F11. This eliminates need for a rescue media or network connection to start the rescue utility. The feature has a trade name "Acronis Startup Recovery Manager".

Acronis Startup Recovery Manager is especially handy 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 the ordinary bootable media.

Activation and deactivation of the Acronis Startup Recovery Manager

The operation that enables using Acronis Startup Recovery Manager is called “activation". You are suggested to activate the Acronis Startup Recovery Manager when creating the Acronis Secure Zone. See Creating Acronis Secure Zone (p. 217).

Once the Acronis Secure Zone is created, you can activate or deactivate the Acronis Startup Recovery Manager at any time using the Manage Acronis Secure Zone tool. The deactivation will disable the boot time prompt "Press F11 for Acronis Startup Recovery Manager…". This means you will need bootable media in the case the system fails to boot.

Limitations

- The recovery manager resides in the Acronis Secure Zone that cannot be organized on a dynamic disk or a disk using the GPT partitioning style
- The recovery manager does not allow for configuring static TCP/IP settings and so it can access a network share only if there is a DHCP server on the network
- The recovery manager activation requires manual reconfiguration of boot loaders, such as LILO and GRUB, on Linux-based machines
- The recovery manager activation requires re-activation of third-party loaders.

2.12.3. 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 disk image can be deployed easily on the hardware where it was created or to identical hardware. However, if you change a motherboard or use another processor version — a likely possibility in case of hardware failure — the deployed 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. Real-to-virtual and virtual-to-real computer migration for system recovery, test and other purposes.

The Universal Restore principles

1. Automatic HAL and mass storage drivers selection.

   Universal Restore searches the Windows default driver storage folders (in the system image) for HAL and mass storage device drivers and installs drivers that better fit the target hardware. You can specify a custom driver repository (a folder or folders on a network drive or CD) which will also be used for drivers search.

   The Windows default driver storage folder is determined in the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\Current version\DevicePath. This storage folder is usually WINDOWS/inf.

2. Manual selection of 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 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 the Universal Restore

Universal Restore comes free with Acronis Backup & Recovery 10 for Microsoft Windows Small Business Server and Acronis Backup & Recovery 10 - Virtual Edition.

Universal Restore for the rest of the 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.

2.12.4. 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 local drive, or more exactly, any device available through the machine's BIOS. This may be Acronis Secure Zone, USB hard drive, 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 image. Acronis Backup & Recovery 10 scans the selected disks or volumes in the image. If this scan finds a supported operating system, the option of Acronis Active Restore becomes available.

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

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

Once the system recovery is started, the operating system boots from the backup image. 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 temporary performance downgrade.

Usage scenarios

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

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

2. The system/storage space ratio is heavily misbalanced towards 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 the 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, 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, USB hard drive, flash drive or any internal hard drive.

   If your operating system and its loader reside on different volumes, always include both volumes in the image. 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 the 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 check box for the Acronis Active Restore option.

   Acronis Active Restore will choose for the boot-up and subsequent recovery the first operating system found during the image 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 image. 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 recovered yet.

   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.

   If you try to log off, shut down or hibernate the machine through the Start Menu commands, the end of the current session will be automatically postponed until the recovery is completed. Should you decide to switch off the machine with the Power button though, all the changes made to the system since the last boot up would be lost, the system will not be recovered, not even partially, and the only possible solution it this case will be to start the recovery process anew, 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. Before you decide to reboot or switch off the machine being recovered, make sure that the background recovery is completed.

2.13. Understanding centralized management

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

2.13.1. Basic concepts

Applying backup policies and tracking their execution

To protect data on a single machine, you install on the machine an agent (p. 314) or multiple agents for various data types you want to protect. You connect the console to the machine and create a backup plan (p. 315) 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. 324) and register (p. 325) 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. 316).

Once you apply the policy to a group of machines, the management server deploys the policy to each of the machines. The agents find on their machines the items they have to back up and create the centralized backup plans (p. 317) on each machine. You will be able to monitor the policies' statuses on the single screen and navigate, if need be, to each machine, plan or task to see their status and log entries. The management server also enables you to monitor and manage the locally originated agent's activities on the machines.

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. 318).

The centralized management does not rule out the direct management (p. 319) of each machine. You can connect the console to each machine and perform any direct management operation, if need be. However, the centralized backup plans can be managed through the management server only, since a well-thought policy functions automatically and rarely requires human intervention.

Using the management server, you can create a common shortcut to a networked storage location which is shared by the registered machines. Such location is called a centralized location (p. 317) because the shortcut is the same on all of the machines. Through this shortcut, the location can be used by any backup policy as well as any backup plan created on the registered machines using the direct management.

Organizing a managed archive storage

What the capacity of your centralized location should be? Will not transferring sizeable backups to the location cause the 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. 325) and configure it to manage a centralized location or multiple centralized locations. Such locations are called managed locations (p. 323).

The storage node helps the agent deduplicate (p. 319) backups before transferring them to managed locations and deduplicates the backups already saved in the locations. Deduplication makes for reducing the backup traffic and saving the 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 location for storing backup archives.

More than one storage nodes, each managing a number of locations, 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. 24).

2.13.2. Setting up the 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 timely deleted from the storage. These goals can be achieved by regular backup of the desired data items to a centralized location with deduplication.

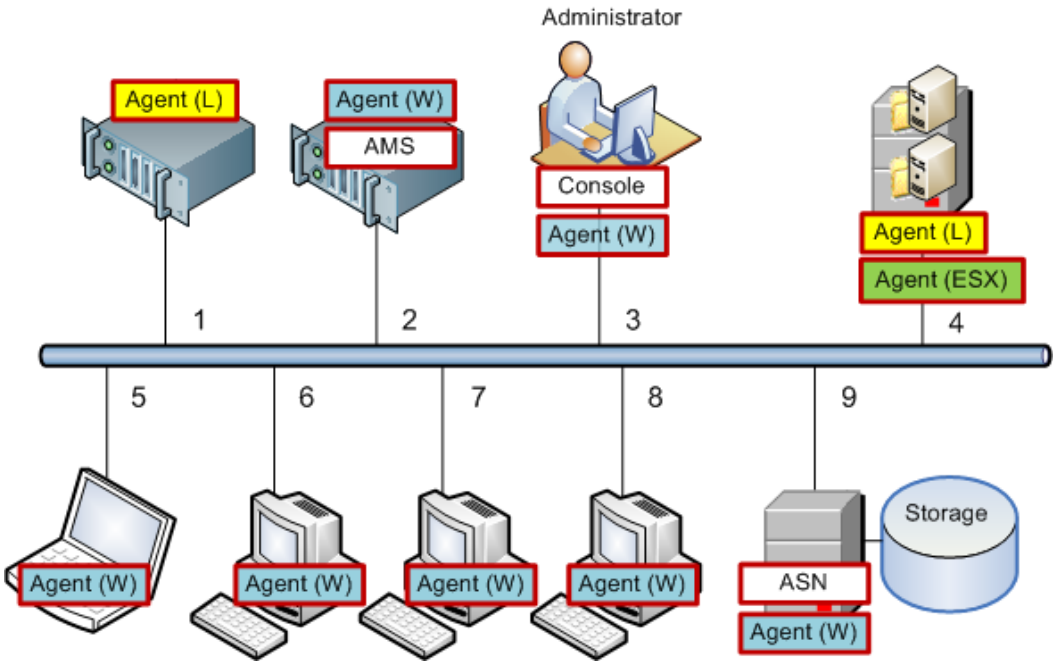

Setting up the Acronis infrastructure

1. Install Acronis Backup & Recovery 10 Management Console **[Console]** on the machine from which you prefer to operate **(3)**. The console enables you to access and manage other Acronis components through the 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.

     The Agent for Linux can be installed on the ESX server since the product is based on Linux Red Hat. If the server uses the ext2 or ext3 file system, you will be able to back up the server's disks, volumes or files. The native ESX file system can be backed up sector-by-sector only.

   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 server's administrator credentials in one of the installation wizard's windows.
4. Install Acronis Backup & Recovery 10 Agent for ESX **[Agent (ESX)]** on the ESX server **(4)** to back up the virtual machines from the host.
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 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 storage locations.
- Multiple Acronis Backup & Recovery 10 components can be installed on a machine with the 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 locations 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 the 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 the node administrators. You can add more accounts later as required.

1. Run the console, connect to the management server.
2. Create a managed location as described in Operations with centralized locations (p. 128). Enable the deduplication when creating.

Setting up groups and policies

The detailed explanation when and why you need to organize groups of machines can be found in the Grouping the registered machines (p. 63) section. Here are some scenarios supported by the described Acronis Backup & Recovery 10 implementation.

- **Protecting the servers.**

Most likely you will 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 and the like. 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 location 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. (Storage node can be added in case the managed location 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.
4. 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.

2.13.3. Grouping the registered machines

As soon as a machine is registered (p. 325) on the management server, the machine appears in the **All machines** built-in group (p. 317). 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 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.

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. 325) since their content never changes unless the administrator explicitly adds or deletes a machine.

The 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. 321) since their content can change automatically.

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

- Operating system
- Active Directory organization unit
- IP address range.

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

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.

Grouping helps organize the data protection by the company departments, the Active Directory domains or organizational units within a domain, by various populations of users, by the site locations and the like. To make the best use of the AD OU criterion, consider reproducing of the Active Directory hierarchy in the management server. Grouping by the IP range enables taking account of the network topology. The diagram below presents an example of group hierarchy.

Example

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.

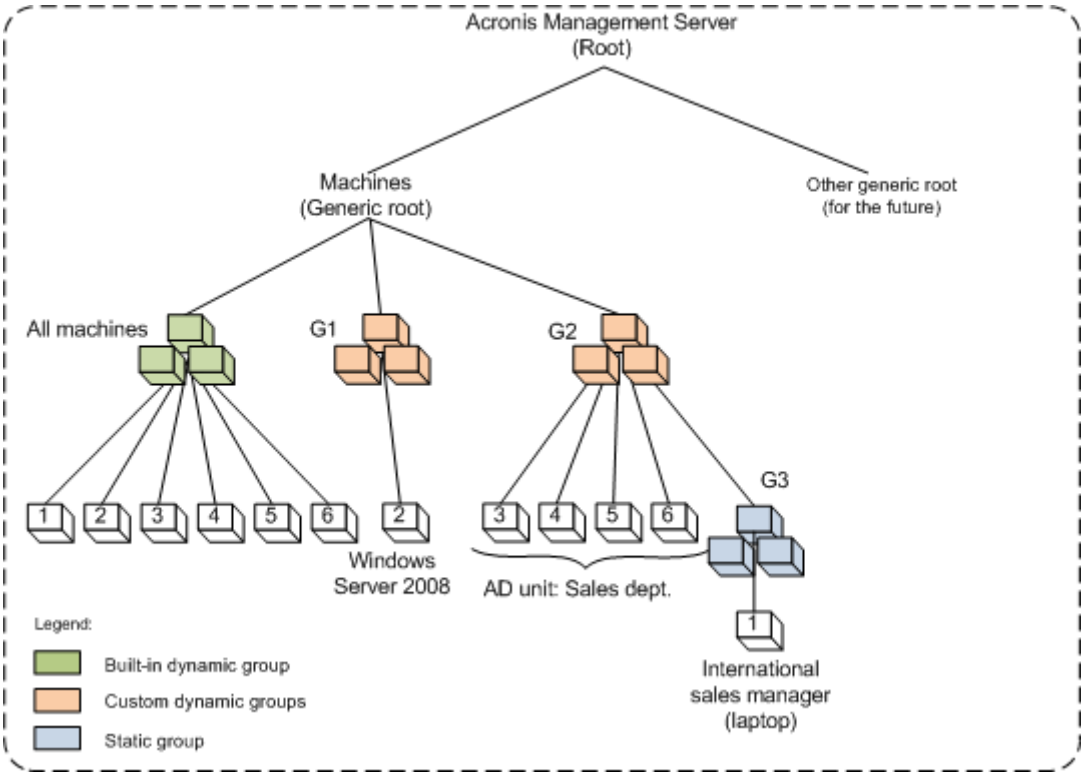

An example of group hierarchy

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 will be 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 it dynamic nature is considered intact.

In real life, the administrator most likely would prefer to protect the manager's machine by applying the policy directly to the machine, without including it to 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.

In the future, Acronis Backup & Recovery 10 Management Server will have generic roots for other entities such as virtual machines or Microsoft SQL servers. These entities will have there own grouping criteria depending on their properties.

Operations with custom groups

You create empty groups in the generic root (Machines) or within existing groups and populate them by manual adding machines (static groups) or 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
- 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 the group.

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 gets available.

For information on how to perform the operations please see Operations with groups (p. 264).

2.13.4. 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; 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: once the correspondence between the policy and one or more machines is removed, the tasks are removed from the machines.

The policy deployment has to be synchronized to the managed 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 cumulative policy's status.

A policy on a machine or a group (p. 67)

Operations with a machine (p. 68)

Inheritance of policies (p. 69)

2.13.4.1. A policy on a machine or a group

On 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

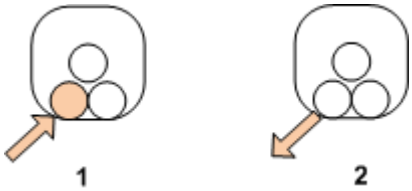

1. A policy can be applied to a machine.
2. A policy can be revoked from the machine.

Policy on a group

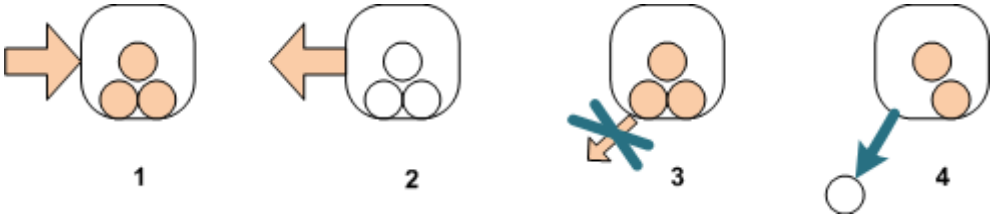

1. A policy can be applied to a group.
2. A policy can be revoked from the group.
3. A policy applied to a group cannot be revoked from a machine.
4. To revoke the policy from the machine, remove the machine from the group.

The same policy on a group and on a machine

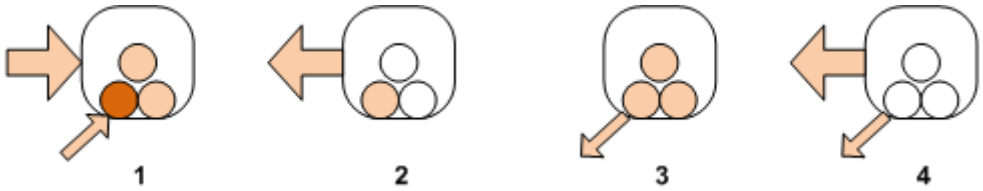

1. The same policy can be applied to a group and to a machine. Nothing changes on the machine at second application of the same policy, but the server remembers that the policy has been applied twice.
2. A policy, revoked from the group, retains on the machine.
3. A policy, revoked from the machine, retains 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.

2.13.4.2. 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.

On 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.

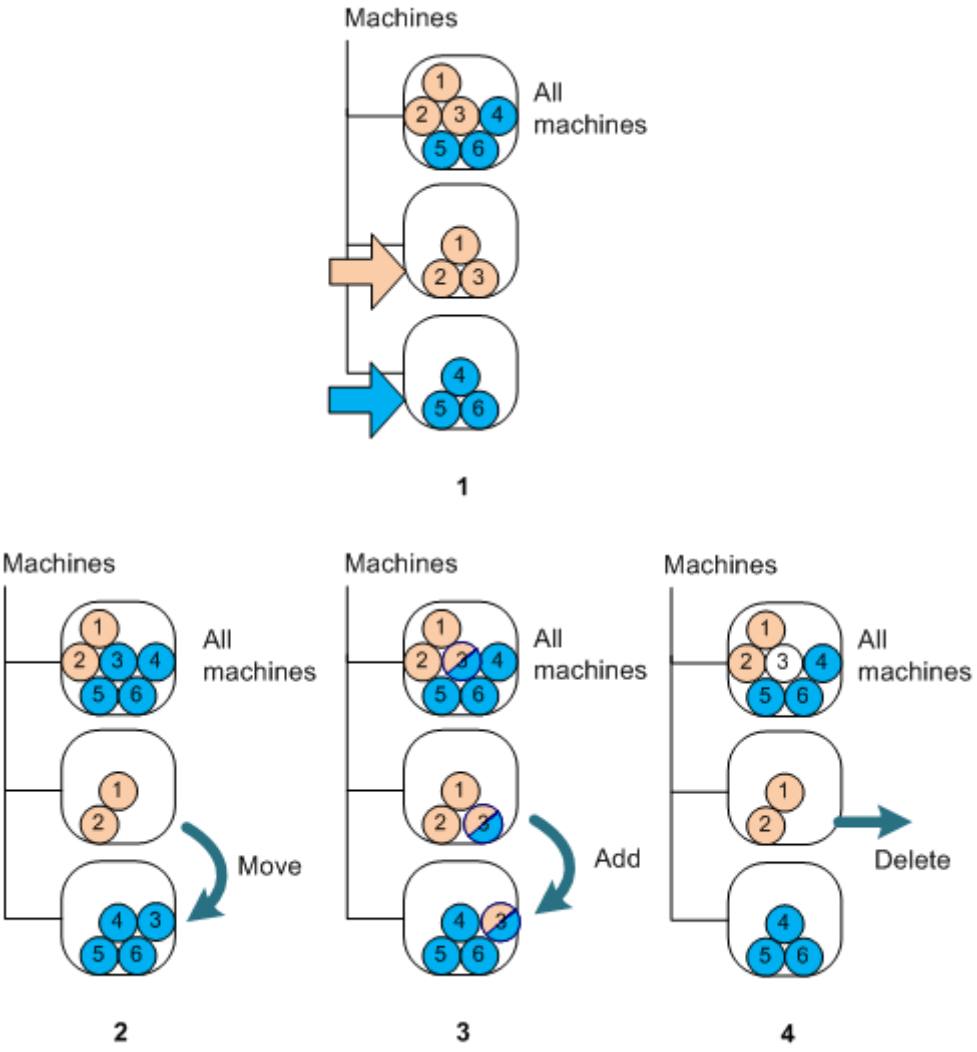

2.13.4.3. Inheritance of policies

The 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.

On 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.

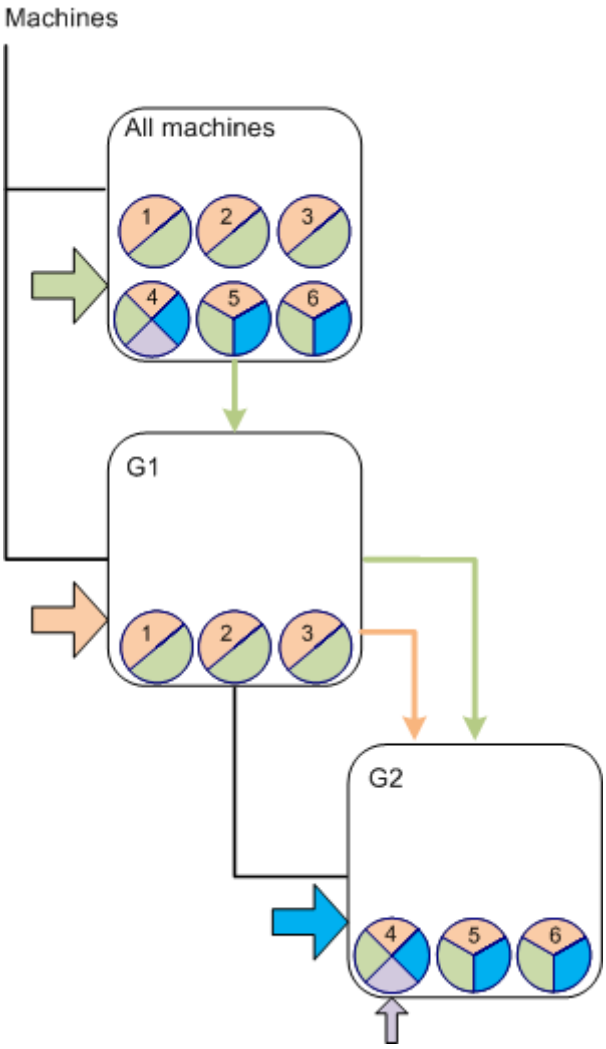

Besides the **All machines** group, we have the custom G1 group in the root and the custom G2 group that is G1's child.

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

The "red" 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 the machine #4. It will exist on #4, as well as the other three policies, in each and every group the machine will be included in; even if no policy will be applied to that group.

Let's assume we create the G3 group in the root. Formally, it inherits only the "green" policy. If no other policies are applied to the group, all its members are supposed to be "green". But if we add, say, the #1 machine to G3, it will bear both "red" and "green" policies, in spite of the fact that G3 has nothing to do with the "red" 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 **Applied to** column shows the origin of all policies existing on the machine. In our example, the policy name and the **Applied to** columns will appear as follows:

For machine	**Name of the policy**	**Applied to**
#1, #2, #3	"green"	All machines
	"red"	G1

#4	"green"	All machines
	"red"	G1
	"blue"	G2
	"violet"	This machine
#5, #6	"green"	All machines
	"red"	G1
	"blue"	G2

2.13.5. 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.

2.13.5.1. 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**.

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 the group the machine is included in. Once you confirm the changes, the management server revokes the policy from the machines the policy was deployed to. 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

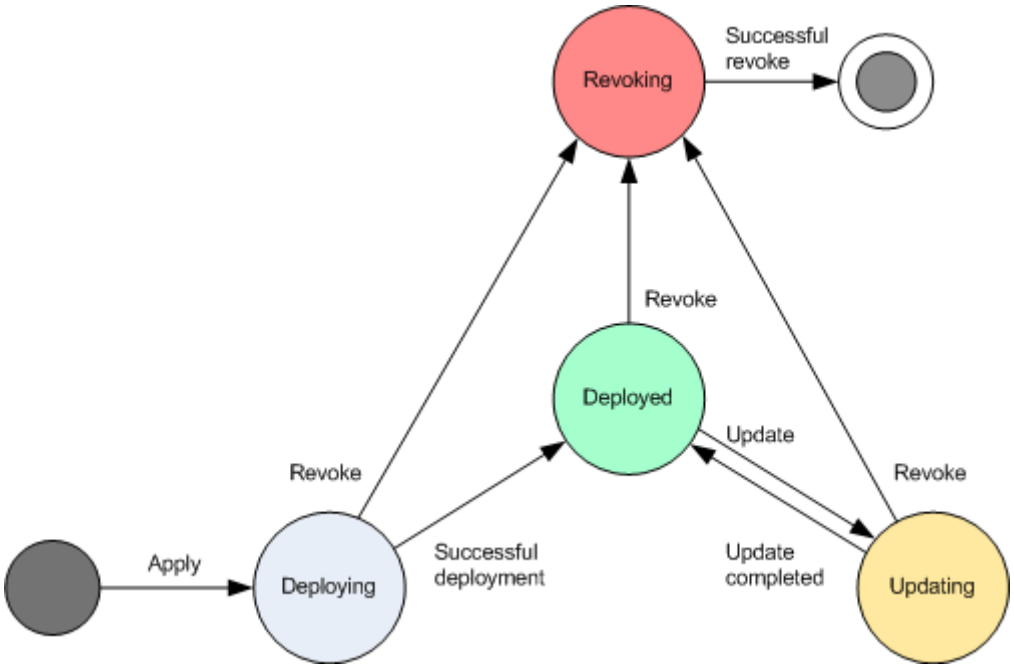

2.13.5.2. Policy status on a machine

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, in fact, 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.

In the **Deploying** state, the policy logs error if none of the items the policy has to back up is found on the machine. If at least one of a number of items is found, the policy logs warning.

Assume, the item selection rule (on page 325) says the policy has to back up volumes D: and F:. The policy is applied to both Linux and Windows machines. The policy logs error on the Linux machines and on the Windows machines that do not have such volumes. The policy logs warning on Windows machines have either D: or F: volume.

The policy that has to back up the [System] and the /dev/sda1 volumes, will log warning 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 log error on Linux machines that do not have a SCSI device or no such volume on the device.

The following table provides details.

State	Status	Description
Deploying	**Error**	The deployment log has errors: none of items to back up is found, disk space runs out… Sometimes deployment with errors may complete successfully, but in most cases the administrator's interaction is required
	Warning	The deployment log has warnings: not all items to back up found; the machine got offline during the deployment; cannot connect for N days…
	OK	The deployment log does not have errors and warnings
Deployed	**Error**	The status of the corresponding backup plan is **Error** OR Updating has completed with error The policy keeps the **Error** status until the same operation (the failed task or the policy update) completes without errors
	Warning	The status of the corresponding backup plan is **Warning** OR Updating has completed with warning The policy keeps the **Warning** state until the same operation (the task or the policy update) completes without warnings
	OK	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
	OK	The updating log does not have errors and warnings
Revoking	**Error**	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 specific machine, the backup policy has deployment state and status on a group of machines and the cumulative deployment state and status of the policy.

2.13.5.3. 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 the A and the B machines. 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 continues on the other, the state will be "Deploying, Deployed". When the deployment completes on both machines, the state will be "Deployed".

2.13.5.4. 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".

2.13.5.5. Policy cumulative state and status

In addition to the deployment state and status as related to a specific machine or group, the backup policy has the cumulative deployment state and 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 the A machine. The policy was successfully deployed. Then you modify the policy and immediately apply it to the group consisting of the B and the C machines. 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".

2.13.6. Deduplication

2.13.6.1. Overview

Deduplication is a process that aims at reducing the size of archives in a centralized location by eliminating redundant data such as duplicate files or disk blocks.

For example, if a location contains two instances of the same file—whether in the same archive or in different archives—the file's content is stored only once, and a link to that content is stored instead of the second file.

Deduplication also reduces 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 a backup. Acronis Backup & Recovery 10 Agent uses the storage node to determine what data can be deduplicated, and deduplicates the data before transferring it to the location.

Deduplication at target

Performed in the location after a backup is completed. The storage node analyses the location's content and deduplicates data in it.

When creating a backup plan, you have an option to turn off deduplication at source for that plan. As a result, the backed-up data will be deduplicated only after it is transferred to the location and the backup is completed. This may lead to lesser backup times but a greater load of the network and the storage node.

A managed centralized location where deduplication is enabled is called a *deduplicating location*. When you create a managed centralized location, you can specify whether to enable deduplication in it.

Deduplication database

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

The deduplication database is stored in the folder specified by **Database path** in the **Create centralized location** view.

The size of the deduplication database is about one percent of the total size of archives in the location. In other words, each terabyte of new (not duplicate) data adds about 10 GB to the database.

2.13.6.2. How deduplication works

Block-level deduplication

When performing a disk backup to a deduplicating location, Acronis Backup & Recovery 10 Agent reads disk blocks being backed up, and computes a fingerprint of each block. Such a fingerprint—often called a *hash value*—uniquely represents the block's content within the location.

Acronis Backup & Recovery 10 Storage Node which manages the location maintains a database of the hash values of all blocks whose content is stored in the location—called the deduplication database.

Before sending the content of a block to the location, the agent queries the storage node whether the block's hash value is the same as that of some already stored block.

If so, the agent sends only the block's hash value; otherwise, it sends the block's content.

File-level deduplication

When performing a file backup to a deduplicating location, Acronis Backup & Recovery 10 Agent computes a fingerprint of each *large* file—a file whose size is 4 KB or more—being backed up. Such a fingerprint, often called a *hash value,* uniquely represents the file's content within the location.

Acronis Backup & Recovery 10 Storage Node which manages the location maintains a database of the hash values of all large files that are stored in the location—called the deduplication database.

Before sending the content of a large file to the location, the agent queries the storage node whether the file's hash value is the same as that of some already stored file. If so, the agent sends only the file's hash value to the location; otherwise, it sends the file's content.

For files smaller than 4 KB, the agent does not compute their hash values, thus always sending the files' content to the location.

2.13.6.3. Deduplication restrictions

Block-level deduplication restrictions

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

- If the volume is a dynamic or 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 cluster size on most NTFS volumes equals 4 KB and so allows for block-level deduplication. Other examples of allocation unit sizes allowing for block-level deduplication include 8 KB, 12 KB, and 64 KB.*

- If you protected the archive with a password

Disk blocks for which deduplication was not performed are stored in the archive as they would be in a non-deduplicating location.

File-level deduplication restrictions

During a file backup to an archive in a deduplicating location, 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 for which deduplication was not performed are stored in the archive as they would be in a non-deduplicating location.

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

2.13.6.4. Deduplication ratio

Deduplication ratio shows the size of archives in a deduplicating location relatively to the size they would occupy in a non-deduplicating location. The higher the deduplication ratio, the more advantageous is the use of deduplication.

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 location will be 2 GB, but this size will be just 1 GB in a deduplicating location. This gives a deduplication ratio of 2:1.

Conversely, if the two files had different content, the backup sizes in a non-deduplicating and deduplicating locations would be the same, and the deduplication ratio would equal 1:1.

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 the ratios of 3:1, 4:1, etc.), a reasonable expectation for a typical environment is a ratio between 1.3:1 and 1.8:1.

As a more real-life 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 for each machine occupy the other 50% (another 1 GB).

In a deduplicating location, 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 location, 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 location. 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.

2.13.7. Privileges for centralized management

2.13.7.1. 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.

On stand-alone machines, the local connection is established automatically when the console starts, provided that both the console and the agent are installed on the machine.

You can establish the local connection manually—for example, to connect to the local machine when the console is disconnected or is connected to a different machine.

To establish the 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; to save the password for the specified user name, select the **Save password** check box. If you leave the **User name** box empty, the credentials under which the console is running on the local machine will be used.

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

The local connection can be established in Acronis Backup & Recovery Management Console by the user who has the root privileges on the machine (the root user), or by a user whom the root user has allowed to manage the machine.

To start the console as a root user

- As the root user, run the following command:
  `trueimageremote`

To allow non-root users to start the console

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

To start the console 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.
2. Run the following command:
   `sudo trueimageremote`

2.13.7.3. 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 the remote connection is established, the user has management rights on the remote machine as described in User rights on a managed machine (p. 34).

Note: *On a remote machine running Windows Vista with enabled User Account Control (UAC)—and which is not part of a domain—only members of the Administrators group can back up data. To overcome the restriction, include the machine into a domain or disable UAC on the machine (by default, UAC is enabled).*

For information about Acronis security groups and their default members, see Acronis security groups (p. 81).

2.13.7.4. 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 strongly 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 http://www.kernel.org/pub/linux/libs/pam/library/.

Remote connections by 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_rootok.so
account  required     pam_unix.so
```

Remote connections by non-root users

To allow non-root users to connect to the machine remotely, the root user can create an authentication policy that will allow certain users to manage the machine with the root privileges.

As a result, the specified non-root users will be able to manage the machine as if they were root users. The root user takes all the risks of using such Linux-PAM configuration.

The following are two examples of authentication policies for non-root users.

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.

1. 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 previous 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
```

2.13.7.5. 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.

Acronis Backup & Recovery 10 Agents

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

Acronis Remote Users

A user who is a member of this group can manage the machine remotely by using Acronis Backup & Recovery 10 Management Console.

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

Acronis Centralized Admins

A user who is a member of this group has the same management rights as do the members of the Administrators group on the machine.

Membership in this group does not imply permission to connect to the machine remotely by using the management console—only members of the Acronis Remote Users group have such permission.

By default, this group is empty. It should remain empty unless you want to allow certain non-administrators to manage the machine as an administrator. In this case, simply add the correspondent users to the group.

The users' management rights on a machine are described in Users' privileges on a managed machine (p. 34).

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

Members of this group are the management server's administrators. They have the same management rights on the registered machines as do the members of the Administrators groups on those machines—regardless of the contents of Acronis security groups there.

No user—even a local administrator—can be an administrator of the management server without being a member of this 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 and Acronis Centralized Admins groups.

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 from 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 had been 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.*

2.13.7.6. Users' 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, can:

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

A user who is a member of the Adminsitrators group additionally can:

- View and manage any archive in any centralized location which is managed by the storage node
- Create centralized locations 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. 269), 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 locations 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.

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.

2.13.7.7. Rights for Acronis services

In Windows, most Acronis components run as services. A service runs under an account—either a user account, such as Administrator; or a system account, such as Local System.

A security best practice is to run each service under a dedicated user account which has only a minimal set of user rights needed for that service.

When installing a component that runs as one or more services, you can specify the name of the account under which these services will run. These accounts will be assigned the rights that are necessary to run the services.

The following table shows the necessary user rights and the default names of the user accounts for each component's services.

Component name	Service names	Necessary user rights	Default user account
Acronis Backup & Recovery 10 Agent	**Acronis Remote Agent** **Acronis Management Machine Service**	Log on as a service Back up files and directories Log on locally Restore files and directories	**Acronis Agent User**
Acronis Backup & Recovery 10 Management Server	**Acronis Management Server**	Log on as a service	**AMS User**
Acronis Backup & Recovery 10 Storage Node	**Acronis Storage Server**	Log on as a service	**ACSS User**

These users are also granted access to the following registry key (called the Acronis registry key): HKEY_LOCAL_MACHINE\SOFTWARE\Acronis

Note: *The Acronis Scheduler service, which provides scheduling for components' tasks, runs under the Local System account. It cannot run under any other account.*

Important: *It is only during installation that you can specify a user account for an Acronis service. A service may fail to start if you change the account manually later, such as via the Services console in Windows.*

Management server administrator's rights

Acronis Backup & Recovery 10 Management Server administrator operates on a registered machine on behalf of the Acronis Management Machine Service that is running on that machine, and has the same privileges on the machine as the service does.

The management server administrator has an 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.

2.13.8. Communication between Acronis Backup and Recovery components

2.13.8.1. Secure communication

Acronis Backup & Recovery 10 provides a 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 credentials needed to establish a connection, by using the Secure Sockets Layer (SSL) protocol.
- **Data encryption** provides secure transfer of information—notably, 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. 86).

For instructions on how to manage SSL certificates used for secure authentication, see SSL certificates (p. 89).

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 secure authentication and data encryption settings.*

2.13.8.2. 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

2.13.8.3. Configuring communication options

You can configure communication options, 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 apply the Administrative Template, see How to apply Acronis Administrative Template (p. 88).

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

To configure communication options

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:

   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 bootable environment, see Network port configuration (p. 88).

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:
  - **Use if possible**

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

- **Don't use**

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

- **Always use**

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

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 Acronis certificates.

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

- **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:

  - **Use if possible**

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

  - **Don't use**

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

  - **Always use**

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

- **Disabled**

  The same as **Not configured**.

1. For the new communication options to take effect, restart any running Acronis components—preferably, by restarting Windows. If restart is unfeasible, make sure to 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 Agents or Acronis Backup & Recovery 10 Management Server, are running, restart their correspondent services from the **Services** console.

2.13.8.4. Network port configuration

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

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

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, which provides the temporary port configuration.

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

Configuring the port in operating system

Windows

To be able to change the ports numbers or disable the ports, apply the Administrative Template, provided by Acronis, as described in Configuring communication options (p. 86), 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 bootable environment

While creating Acronis bootable media, you have an 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.

2.13.8.5. How to apply Acronis Administrative Template

The Administrative Template, provided by Acronis, enables fine-tune of some security related features, including encrypted communication settings. Through the Microsoft Group Policy mechanism, the template can be applied to a single computer as well as to a domain.

To apply the Acronis Administrative Template:

1. Run Windows Group Policy Objects Editor (%windir%\system32\gpedit.msc.)
2. Open the Group Policy object you want to edit.
3. Expand Computer Configuration.
4. Right click **Administrative Templates**.

5. Click **Add/Remove Templates**.
6. Click **Add**.
7. Browse to the Acronis Administrative Template (\Program files\Common Files\Acronis\Agent \acronis_agent.adm or \Program files\Acronis\TrueImageConsole\acronis_agent.adm), and click **Open**.

Once the template is added, you can open it and edit the desired settings. After applying the template or editing the options, 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

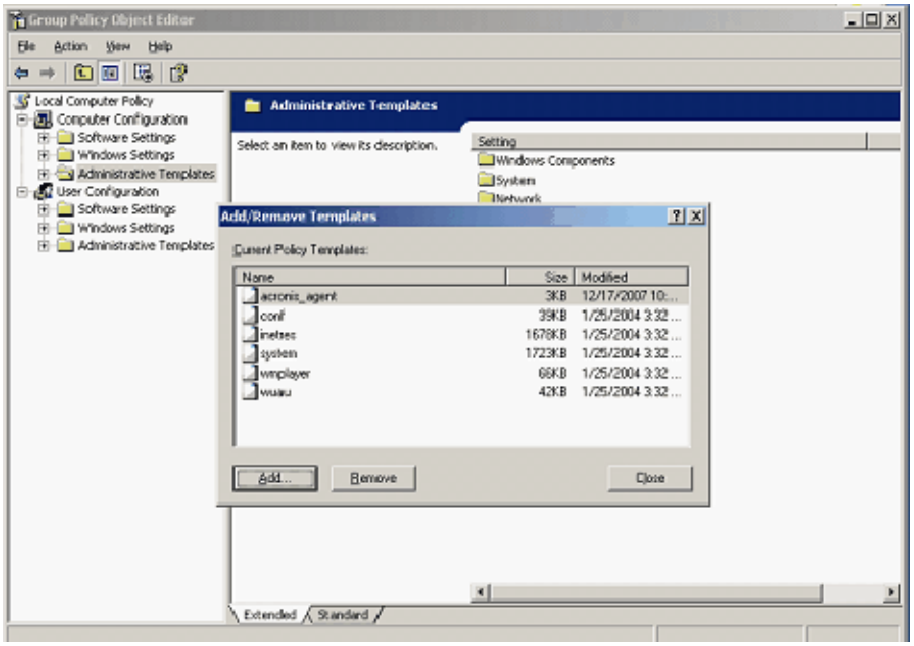

2.13.8.6. SSL certificates

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

An SSL certificate for the components can be of one of two types:

- **Third-party certificate**, a certificate issued by a Certificate Authority (CA)—for example, by a public CA such as VeriSign® or Thawte™, or by deploying a Public Key Infrastructure (PKI).
- **Acronis certificate**, a self-signed certificate created during the installation of an Acronis component.

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

To ensure reliability, the certificate is stored in Windows Certificate Store at the following location: Certificates (Local Computer)\Acronis Trusted Certificates Cache. A hash is used for future host identification.

In case the list of certificates for the local machine is not displayed in the **Certificates** console, you can use the following procedure.

To open the list of a computer's certificates

1. Click **Start**, then click **Run**, and then type **mmc**
2. In the console, on the **File** menu, click **Add/Remove Snap-in**.
3. In the **Add/Remove Snap-in** dialog box, click **Add**.
4. In the **Add Standalone Snap-in** dialog box, double-click **Certificates**.
5. Click **Computer account**, and then click **Next**.
6. Click **Local computer**, and then click **Finish**.

   Tip: *Alternatively, you can manage the list of certificates of a remote machine. To do this, click* ***Another computer*** *and then type the remote machine's name.*

7. Click **Close** to close the **Add Standalone Snap-in** dialog box, and then click **OK** to close the **Add/Remove Snap-in** dialog box.

Acronis 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 other than Acronis Backup & Recovery 10 Management Console.

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

To generate a new Acronis certificate

1. As an administrator, open the folder with the old certificate, whose location is specified by the certificate path—such as C:\Program Files\Common Files\Acronis\Agent\certificate.pem
2. By editing the permissions for the certificate's file, add the Full Control permission to the user under which you are logged on (by default, only the NETWORK SERVICE account has a permission to access the file).
3. Delete or rename the file.
4. Reinstall any of the installed components that can act as a server: Acronis Backup & Recovery 10 Agent, Acronis Backup & Recovery 10 Management Server, or Acronis Backup & Recovery 10 Storage Node.
5. Restart Windows, or restart the running Acronis services.

Third-party certificates

You have an option to use trusted third-party certificates as an alternative to self-signed Acronis certificates, by using Acronis Certificate Commandline 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:

   ```
   %ProgramFiles%\Common Files\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

3.1. Console options

The console, connected to a managed machine or to the management server, provides a suitable GUI for direct or centralized machine management.

Every time you launch the console and connect it to a machine or to server, it displays startup page, fonts, pop-up messages, alerts and table in the **Task** view in accordance with the **Console options** used by default.

To redefine the **Console options**, select **Options > Console options** from the menu.

3.1.1. Startup page

The preset is to display the **Welcome** view on connecting the console to a managed machine or to the management server.

To display the **Dashboard** view on the connection, select the related check box.

To display the **Welcome** view on the connection, remove selection from the related check box.

Note that the same option, influencing on this one, is on the **Welcome** view. If you select **At startup, show the Dashboard view instead of the current view** check box on the **Welcome** view, the respective setting of the console options will be updated.

3.1.2. Pop-up messages

The presets are the following:

- Pop up messages when interaction is required
- Enable pop-up messages upon warnings
- Enable pop-up messages upon errors
- Enable pop-up messages upon task completion

These options enable to adjust displaying for pop-up messages.

Each time a task needs human interaction (for example, to change a media, to provide additional space on a location, to ignore a read error...), an alert is displayed on the **Dashboard** view of the console. You have an option to pop up the **Tasks Need interaction** messages in addition to the alerts.

Note, there is the **Show Need Interaction windows** option on the **Dashboard** view. If you change this option, the respective setting of the console options will be updated.

Acronis Backup & Recovery 10 logs a history of actions the program does on a managed machine or an user does on a machine using the program. Each log entry is a structure with data about an event occurred on a machine. There are three types of events in the program: Error, Warning and Information. You can enable or disable the console pop-up messages upon error events, warning events, and task completion information events. These options relate only to the console connected to a managed machine.

3.1.3. Time-based alerts

Presets:

- alert if the last successful backup on a machine was completed more than **3 days** ago.
- alert if the last machine's connection to the management server was more than **3 days** ago.

When the console is connected to a managed machine or to the management server, issues occurred during data protection operations are displayed in the **Alerts** section of the console's **Dashboard**.

Last backup

The data changes occurred after the last backup are under a risk of loss. The console can alert you that no data was backed up on the machine for a relatively long period of time. You can configure the time period that is considered as critical.

Last connection

The console can alert you that no connection was established between a registered machine and the management server for a relatively long period of time and so the machine might not be centrally managed. You can configure the time period that is considered as critical.

This option is effective when console is connected to the management server or to a machine that is registered on the server.

When the console is connected to the management server, these options' settings also influence on coloring of the **Last connect** and **Last backup** columns' values of the **Machines** view and are used when generating reports.

3.1.4. Tasks number

The preset maximum is 500 tasks in the **Task** view of the console connected to the management server.

The management server provides management for number in the thousands of machines. And there are at least one backup plan as well as at least one task on each managed machine. So the **Tasks** view of the console connected to the management server might contain thousands of tasks and as the result it might be slow and awkward. To avoid the problem use the option of limitation for the tasks number in the console **Task** view.

Use the **Number of tasks** box to specify the value. Note that it don’t allow to type in value more than 500.

3.1.5. Fonts

The preset is **System Default** font both for menu and application interface items.

You can view and edit the font properties for the product interface items. It is possible to adjust separately font for menu items (drop-down and context menus) and font for other interface elements of the application including the menu.

3.2. Machine options

To access the Machine Options, select **Options > Machine Options** from the menu.

Use the options to:

- specify the machine management type (p. 94)
- adjust the machine event tracing (p. 96)

These options influence tracing of the Acronis Backup & Recovery 10 events for backup (Creating a backup plan > Backup options > Event tracing (p. 109)) and recovery (Recovering data > Recovery options > Event tracing (p. 121)) operations on the machine.

3.2.1. Machine management

Preset is **Stand-alone management**.

It is possible to manage a machine directly or centrally in Acronis Backup & Recovery 10.

Direct management is provided by means the management console connected to the machine Agent. Centralized management for the machine registered on the management server is realized through the console connected to the management server.

To set centralized management on the machine, register it on the management server:

1. Select **Centralized management** option.
2. Specify **Management Server IP/Name**.
3. Click **OK**. If the specified IP/Name is correct, you will go to the next step.
4. If you have rights enough for the registration, the machine will be registered on the management server and will be managed centrally through the console connected to the management server.
5. If you don't have rights enough for the registration, the **Credentials** window will be displayed. Enter the credentials (user name and password) and click **OK**.

To manage the machine directly select **Stand-alone management** option.

3.2.2. Event tracing

It is possible to duplicate log entries for the events issued during the Acronis Backup & Recovery 10 operations in the Windows Application Event Log (Windows event log) or send the event notifications using SNMP (SNMP notifications).

3.2.2.1. Windows event log

This option relates to Windows operating systems.

The preset is: Disabled.

Acronis Backup & Recovery 10 Agent can record events to the Windows Application Event Log (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**).

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

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

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

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

3.2.2.2. SNMP notifications

This option applies to both Windows and Linux operating systems.

This option is not available when operating under bootable media.

The preset is **Disabled**.

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

1.3.6.1.4.1.24769.100.200.1.0 - string identifying a type of occurred event (Information, Warning, Error)

1.3.6.1.4.1.24769.100.200.2.0 - string containing text description of occurred event (it looks identically to messages published by Acronis Backup & Recovery 10 in its log.)

To set up sending SNMP messages

1. Select the check box for **Send messages to SNMP server**.
2. Specify the appropriate options as follows:
   - **Types of events to send** - choose types of events to be reported: **All events**, **Warnings and Errors**, or **Errors only**.
   - **Server's name or IP address** – type the name or IP address of the host running the SNMP management application, to which messages will be sent.
   - **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".

To disable sending SNMP messages, clear the check box for **Send messages to SNMP server**.

The messages are sent over UDP.

Setting up SNMP services on a recipient

3.3. Management server options

To access the Acronis Backup & Recovery 10 Management Server options, select **Options > Management server options** from the menu.

Use the options to:

- set the management server Logging level (p. 96)
- adjust the management server Event tracing (p. 96).

3.3.1. Logging level

The preset is **Enabled** for **All events**.

The management server collects log entries from registered machines into centralized log stored in a dedicated database.

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

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

To disable the log entries collecting, remove selection from the **Collect log** check box.

3.3.2. Event tracing

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

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

3.3.2.1. Windows event log

This option relates to Windows operating systems.

The preset is: Disabled.

Acronis Backup & Recovery 10 Management Server can record events to the Windows Application Event Log (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**).

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

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

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

- **Warnings and errors**
- **Errors only**.

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

3.3.2.2. SNMP notifications

The preset is **Disabled**.

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

1.3.6.1.4.1.24769.100.200.1.0 - string identifying a type of occurred event (Information, Warning, Error)

1.3.6.1.4.1.24769.100.200.2.0 - string containing text description of occurred event (it looks identically to messages published by Acronis Backup & Recovery 10 in its log.)

To set up sending SNMP messages

1. Select the check box for **Send messages to SNMP server**.
2. Specify the appropriate options as follows:
   - **Types of events to send** - choose types of events to be reported: **All events**, **Warnings and Errors**, or **Errors only**.
   - **Server's name or IP address** – type the name or IP address of the host running the SNMP management application, to which messages will be sent.
   - **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".

To disable sending SNMP messages, clear the check box for **Send messages to SNMP server**.

The messages are sent over UDP.

3.4. Default backup and recovery options

3.4.1. Default backup options

Each of Acronis agents has its own default backup options with preset values. To access the default options, select **Options - Agent options - Default backup options** from the menu. Once you have changed a preset value, it becomes default for this agent and will be used as default in further backup operations. Nevertheless, while backing up you can override the default settings with the custom settings that will be specific for this plan only.

Availability of the backup options

The set of available backup options differs depending on:

- The environment the agent operates in (Windows, Linux, bootable media)
- The type of the data being backed up (disk, file, storage group, mailbox)

- Backup destination (networked location or local disk)
- Backup scheme (Back up now or using the scheduler)

The following table shows the summary of the backup options availability.

	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. 100) (password + encryption)	+	+	+	+	+	+
Source files exclusion (p. 100)	+	+	+	+	+	+
Pre-post backup commands (p. 101)	+	+	+	+	PE only	PE only
Pre-post data capture commands (p. 103)	+	+	+	+	-	-
Multi-volume snapshot (p. 106)	+	+	-	-	-	-
File-level backup snapshot (p. 105)	-	+	-	+	-	-
Use VSS (p. 106)	+	+	-	-	-	-
Compression level (p. 106)	+	+	+	+	+	+
Backup performance:						
Backup priority (p. 107)	+	+	+	+	-	-
HDD writing speed (p. 107)	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD
Network connection speed (p. 107)	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share
Fast incremental/differential backup (p. 110)	+	-	+	-	-	-
Backup splitting (p. 110)	+	+	+	+	+	+
File-level security (p. 111):						
Preserve files' security settings in archives	-	+	-	-	-	-
In archives, store encrypted files in decrypted state	-	+	-	-	-	-
Media components (p. 112)	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	-	-

Error handling (p. 112):						
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Ignore bad sectors	+	+	+	+	+	+
Dual destination (p. 113)	+	+	+	+	-	-
Task start conditions (p. 114)	+	+	+	+	-	-
Task failure handling (p. 115):						
Stop executing the backup plan	+	+	+	+	-	-
Restart the failed task	+	+	+	+	-	-
Additional settings (p. 115):						
Overwrite data on a tape without prompting user for confirmation	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape
Dismount media after backup is finished						
Ask for first media while creating backup archives on removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media
Validate backup after creation	-	-	-	-	+	+
Reset archive bit	-	+	-	-	-	+
Reboot after the backup	-	-	-	-	+	+
Create full backups as synthetic backups	+	+	+	+	+	+
Notifications:						
E-mail (p. 108)	+	+	+	+	-	-
Win Pop-up (p. 108)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 109)	+	+	-	-	-	-
SNMP (p. 109)	+	+	+	+	-	-

3.4.1.1. Archive protection

The preset is **Disabled**.

To protect the archive data from unauthorized access

1. Select the **Set password for the archive** check box.
2. In the **Password** field, type a password
3. In the **Confirm the password** field, re-type the password
4. Choose one of the following:
   - **None** – do not use encryption
   - **AES 128** – to use 128 bit encryption
   - **AES 192** – to use 196 bit encryption
   - **AES 256** – to use 256 bit encryption
5. Click **OK**.

The bigger the key size, the longer time it will take for program to encrypt the archive and the greater your data security will be.

AES cryptographic algorithm operates in CBC mode and uses the randomly generated key with a user-defined size of 128, 192 or 256 bits. The encryption key is in turn encrypted with AES-256 using SHA-256 hash of the password as a key. The password itself is not stored anywhere on the disk or in the backup, the password hash is used for verification purposes. With this two-level security, the backup data is protected from any unauthorized access, but the lost password recovery is not possible.

3.4.1.2. Source files exclusion

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

This option can also be defined on Create Backup Plan page.

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

Predefined exclusions

To use the predefined sets of exclusions, check any of the following:

- **Exclude all hidden files and folders** – all files and folders marked with a "hidden" attribute will be excluded. Hidden files and folders contain user preferences, program and operating system-related data.
- **Exclude all system files and folders** - all files and folders marked with a "system" attribute will be excluded. System files have .sys extension in Windows.

Custom exclusions

If required, you can create your own exclusions for the specific files and/or file types.

To do so, proceed as follows

1. Select the **Exclude files matching the following criteria** check box
2. Click **Add…**
3. In the **Add file exclusion criterion** window specify an exclusion criterion, and then click **OK**. Both explicit rules (by entering an exact file name, or path) and common Windows masking rules (with using wildcard characters) are supported. See the table below for examples of exclusions.
4. The created exclusion appears in the bottom area of the window.
5. Click **OK** to save your settings.

*Note: *.bak, *.~, *.tmp file types are excluded by default.*

Exclusion examples

Criterion	Example	Description
By name	File1.log	Excludes all files named File1.log.
By path	C:\Finance\test.log	Excludes file named test.log and located on disk C:.
Mask (*)	*.log	Excludes all files with .log extension.
Mask (?)	my???.log	Excludes all .log files with names consisting of five symbols and starting with “my”.

Editing custom exclusions

Since you create a custom exclusion, you are able to edit and/or delete it. To do so, click the respective buttons.

3.4.1.3. Pre/post commands

This option is not available when backing up from Linux-based bootable media.

Specify commands or executable files to be automatically executed before and after the backup procedure.

Use cases

- remove some .tmp 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 following scheme illustrates when pre/post commands are executed.

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 option by checking the following options:
   - Execute before backup process
   - Execute after backup process
2. Do any of the following:
   - Enter a new command or specify a batch file by clicking the **Edit** button
   - Select the existing command/batch file from the drop-down list
3. Click **OK**.

Pre-backup command

To specify a command/batch file to be executed before 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.

	Selection			
Fail the task if the command execution fails	☑	☐	☑	☐
Do not back up until the command execution is complete	☑	☑	☐	☐
Result	**Default** Start the backup after the command is successfully executed. Do not start the backup and fail the task if the command execution failed.	Start the backup after the command is executed despite of the execution failure or success.	N/A	Run the backup concurrently with the command and independently of the command execution result.

Post-backup command

To specify a command/executable file to be executed after backup

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 the command successful execution is critical for your backup strategy, select the check box for **Fail the task if the command execution fails**. In case the command execution fails, the program will remove the resulting .tib file and temporary files if possible, and fail the task.

   By default the check box is not selected. With this setting, 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 not available when backing up from bootable media.

It is recommended that you use pre/post data capture commands if there is no Microsoft Volume Shadow Copy Service (VSS) in server operating system, or your databases are incompatible with VSS. Otherwise, use Volume Shadow Copy Service options. Pre/post data capture commands ensure the transactions completion by executing commands, executable files or scripts that suspend the appropriate Windows services and automatically resume them after data capture.

It is critical to note that these commands, as opposed to Pre/post commands, will be executed before and after the data capture process, which takes seconds, while the entire backup procedure may take much longer, depending on the amount of data to be imaged. Therefore, the database idle time will be minimal.

The following scheme illustrates on which stage of backup procedure pre/post data capture commands take part.

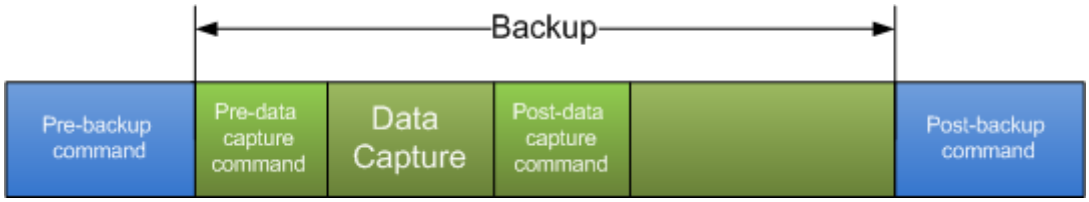

Before/after data capture commands can also be used for other purposes, especially if VSS support is enabled. You may want to suspend an application other than a database, for example. The commands execution and the VSS actions will be sequenced as follows:

“before” commands -> VSS Suspend -> data capture -> VSS Resume -> “after” commands.

To specify pre/post data capture commands

1. Enable pre/post commands execution option by checking the following options:

- **Execute before data capture**
- **Execute after data capture**

2. Do any of the following:
   - Enter a new command, or specify a batch file by clicking the **Edit** button
   - Select the existing command/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.

	Selection			
Fail the backup task if the command execution fails	☑	☐	☑	☐
Do not perform data capture until the command execution is complete	☑	☑	☐	☐
Result	**Default** Start the data capture after the command is successfully executed. Fail the task if the command execution failed.	Perform the data capture after the command is executed despite of the execution failure or success.	N/A	Perform the data capture concurrently with the command and independently 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.

	Selection			
Fail the task if the command execution fails	☑	☐	☑	☐
Do not back up until the command execution is complete	☑	☑	☐	☐
Result	Continue the backup after the command is successfully executed. Once the command execution fails, do not continue operation, remove the resulting .tib file and temporary files if any, and fail the task.	**Default** Continue the backup after the command is executed despite of the command execution failure or success.	N/A	Continue the backup concurrently with the command execution and independently of the command execution result.

3.4.1.5. File-level backup snapshot

The preset is **Create snapshot if it is possible**.

Applies to file-level backup only.

File-level backup can be performed through taking a disk snapshot or directly, by copying the files to the backup. When using a snapshot, you can back up any file including files opened for exclusive access. All the files will be backed up at the same point in time and so files in the backup will be time-consistent.

Choose one of the following:

- **Always create a snapshot**

  Select this only if the above factors are important, that is, backing up files without a snapshot does not make sense. To use a snapshot, the backup plan account must have 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. Trying to back up files 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

The preset is **Disabled**.

The option applies to Windows operating systems. This option applies to disk-level backups and file-level backups performed through taking a snapshot. (The File-level backup snapshot option determines whether a snapshot will be taken during file-level backup.)

When enabled, 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, such a Oracle database.

When disabled, the volumes snapshots will be taken one by one. Therefore, the data spanned across the volumes may be not consistent.

3.4.1.7. Volume Shadow Copy Service

The preset is **Disabled**.
The option applies to file-level and disk-level backups and Windows operating systems.

Microsoft Volume Shadow Copy Service (VSS) provides the infrastructure for backup on running systems by keeping up coordination between user applications that update data on disk and backup applications. VSS is available in Microsoft Windows XP and Microsoft Windows Server 2003 operating systems. The examples of VSS-aware databases are Exchange, Oracle, SQL Server.

If your database is compatible with Microsoft Volume Shadow Copy Service (VSS), then choose Enable Microsoft VSS support. This will ensure completion of all transactions before the backup process starts. Then the database will be ready to access immediately after recovery.

If your database is incompatible with VSS, use Pre/post data capture commands (p. 103) option.

3.4.1.8. Compression level

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 compress more than other file types.

Choose the necessary compression level:

- **None** – the data will be copied as is, without any compression. The resulting backup will have practically the same as for the data being backed up. Takes minimum time to create.
- **Normal** – recommended in most cases. The resulting backup will be approximately a half less than the initial data size.
- **High** – the data will be compressed with high level. The resulting backup will be approximately s less than the initial data size. Takes quite long time to create.
- **Maximum** – the data will be compressed as much as possible. Takes longest time to create. You might want to select maximum compression for removable media to reduce the number of blank disks required.

3.4.1.9. Backup performance

Backup performance options might have a more or less noticeable effect on the backup process speed. This depends on overall system configuration and physical characteristics of devices.

Backup priority

This option is not available when backing up under bootable media.

The preset is: Low

The priority of a process running in a system determines the amount of CPU usage and system resources allocated to that process. Decreasing the backup priority will free more resources for other CPU tasks. Increasing the backup priority might speed up the backup due to taking resources from the other currently running processes. The effect will depend on total CPU usage and other factors.

To specify backup priority

Choose one of the following:

- **Low** – to minimize the backup process speed, unleashing resources for 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 the other processes.

HDD writing speed

This option is available if the machine's hard disk drive is selected as a backup destination.

Backing up in the background to an internal hard disk (for example, to the Acronis Secure Zone) may slow other programs' performance because of the large amounts of data transferred to the disk. You can limit the hard disk usage by the program to the desired level.

To set the desired HDD writing speed for data being backed up

Do any of the following:

- drag the slider
- enter the writing speed in kilobytes per second

Network connection speed

This option is available if a network share is selected as a backup destination.

If you frequently backing up data to network drives and need to reserve a part of bandwidth for other network activities, think of limiting network connection speed. By default the speed is set to maximum, i.e. backup transferring allocates maximum of network bandwidth. This setting is also applied to an FTP connection, if an FTP server is selected as backup destination device.

To set the desired network connection speed for data being backed up

Do any of the following:

- drag the slider
- 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 recipients when the backup is finished through:

- E-mail (p. 108)
- Windows Messenger (WinPopup) (p. 108)

E-mail

Preset is **Disabled**.
This option is not available when backing up from bootable media.

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 options as follows:
   - **Upon operation's successful completion** – to send notification when the operation is completed successfully
   - **Upon operation failure** – to send notification when the operation failed
   - **Add full log to the notification** - to add the full operation log to the message
4. 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 intact, messages will be constructed as if they are from the destination address.
   - **Use encryption** – to activate encryption for the notification message. SSL and TLS encryption methods are available for selection.
   - Check **Use specified outgoing e-mail server** to enable SMTP server and to set up its settings
     - **Outgoing mail server (SMTP)** - provide the outgoing SMTP server name.
     - **Port** – set port of the outgoing mail server. By default the port is set to 25.
     - **User name** - enter the user name.
     - **Password** – enter the password
5. Click **Send Test E-mail Message** to check if the settings are correct.

Windows Messenger (WinPopup)

Preset is **Disabled**.
This option is not available when backing up from bootable media.

Before configuring Windows Messenger notification, make sure the Messenger service is enabled on both the machine executing the task and the machine that will receive messages.

The Messenger service is disabled by default in Windows Server 2003 family. Change the service Startup mode to Automatic and start the service.

To configure Windows Messenger:

1. Select the **Send WinPopup notifications** check box to activate notifications.
2. In the **Computer name** field, provide the name of the machine to which notifications will be sent.
3. Under **Send notifications**, select the appropriate options as follows:
   - **When backup/recovery completes successfully**– to send notification when the operation is completed successfully
   - **When backup/recovery fails** – to send notification when the operation failed
   - **When user interaction is required** – to send notification during the operation when user interaction is required (always selected)
4. Click **Send Test Message** to check if the settings are correct.

3.4.1.11. Event tracing

It is possible to duplicate log entries for the events issued during the Acronis Backup & Recovery 10 backup operations in the Windows Application Event Log (Windows event log) or send the event notifications using SNMP (SNMP notifications).

Windows event log

This option relates to Windows operating systems only.

This option is not available when backing up under bootable media.

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

Acronis Backup & Recovery 10 Agents can record events concerned with the backup operations in the Windows Application Event Log (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**).

To select whether to log 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 option setting specified for the machine (Machine options (p. 94)).
- **Log the following event types** - to log events concerned with the backup operations in the Windows Application Event Log. Specify types of the 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 concerned with the backup operations in the Windows Application Event Log.

SNMP notifications

This option applies to both Windows and Linux operating systems.

This option is not available when operating under bootable media.

The preset is **Disabled**.

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

1.3.6.1.4.1.24769.100.200.1.0 - string identifying a type of occurred event (Information, Warning, Error)

1.3.6.1.4.1.24769.100.200.2.0 - string containing text description of occurred event (it looks identically to messages published by Acronis Backup & Recovery 10 in its log.)

To select whether to send SNMP messages related to the backup operations

Choose one of the following:

- **Use the setting set in the Machine options** - to use options specified for the Machine (Machine options (p. 94)).
- **Send SNMP notifications individually for backup operation events** - to ignore the Machine options (for operations performed within the backup plan only) and select the appropriate options as follows:
  - **Types of events to send** - choose types of events to be reported: **All events**, **Warnings and Errors**, or **Errors only**.
  - **Server name or its IP** – type the name or IP address of the host running the SNMP management application, to which messages will be sent.
  - **Community** – type the name of SNMP community to which both the host running SNMP management application and the machine executing the backup plan belong. The typical community is "public".
- **Do not send SNMP notifications** - to disable SNMP messages about operations performed within the backup plan.

3.4.1.12. Fast incremental/differential backup

The preset is **Enabled**.
The option applies only to the disk-level backup.
This option is available when backing up file systems used by both Windows and Linux.

Incremental/differential backup captures only changes in data occurred since the last backup. To speed up the backup process, the program determines whether the file has changed by file size and the date/time when the file was last saved. Disabling this feature will make the program compare the entire file contents to that stored in the archive.

3.4.1.13. Backup splitting

The preset is **Automatic**.

Sizeable backups can be split into several files that together make the original backup. A backup file can be split for burning to removable media or saving on an FTP server (data recovery directly from an FTP server requires the archive to be split into files no more than 2GB in size).

To configure archive splitting, choose

- **Automatic** - with this setting, Acronis Backup & Recovery 10 will act as follows:

  When backing up to the hard disk: The program will create a single archive file if the selected disk has enough space and its file system allows the estimated file size.

  The program will automatically split the backup into several files if the storage disk has enough space, but its file system does not allow the estimated file size.

  FAT16 and FAT32 file systems have a 4GB file size limit. However, the existing hard drive's capacity can reach as much as 2TB. Therefore, an archive file might easily exceed this limit if you are going to back up the entire disk.

  If you do not have enough space to store the backup on your hard disk, the program will warn you and wait for your decision as to how you plan to fix the problem. You can try to free some additional space and continue or click Back and select another disk.

  When backing up to a diskette, CD-R/RW or DVD+R/RW: Acronis Backup & Recovery 10 will ask you to insert a new disk 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. That comes in handy when backing up to a hard disk with a view to burning the archive to CD-R/RW or DVD+R/RW later on.

  Creating a backup directly on CD-R/RW or DVD+R/RW generally will take considerably more time than it would on a hard disk.

3.4.1.14. File-level security

Preset is Preserve file's security settings in archives.

The option applies only to file-level backups.

To configure file-level security settings, select the appropriate options as follows:

- **Preserve file's security settings in archives**
  Preset is: enabled

  By default, files and folders are saved in the archive with their original Windows security settings (i.e. permissions for read, write, execute and so on for each user or user group, set in file **Properties -> Security**). 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.

  You can disable preserving the files' security settings in archives to completely eliminate this kind of problem. The recovered files/folders will always inherit the permissions from the parent folder to which they are recovered (parent folder or disk, if recovered to the root).

  Alternatively, you can disable files' security settings during the recovery (p. 120), even if they are available in the archive. The result will be the same - the files will inherit the permissions from the parent folder.

- **In archives, store encrypted files in decrypted state**
  Preset is: disabled

  Simply ignore this option if you do not use the encryption feature available in Windows starting from Windows XP and Windows 2003 Server. (Files/folders encryption is set in **Properties -> General -> Advanced Attributes -> Encrypt contents to secure data**).

Check the option if there are encrypted files 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 another machine.

3.4.1.15. Media components

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

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

- **Acronis Backup and Recovery 10 Bootable Agent** is a bootable rescue utility (based on Linux kernel) that includes most of functionality of the Acronis Backup & Recovery 10 Agent. Adding this component to the media will enable you to boot a machine from the media or Acronis PXE Server and recover data from the backup stored on the same media or in another accessible location.
- **One-click Recovery** is a minimal addition to the disk backup stored on removable media, allowing one-click recovery from this backup. At boot from the media and clicking **Recover** the disk will be silently recovered from the media.

  Because the one-click approach does not presume user selections, such as selecting volumes to recover, Acronis One-Click Recovery always recovers the entire disk. Therefore, if your disk consists of several volumes and you are planning to use Acronis One-Click Recovery, all the volumes must be included in the image. Any volumes missing from the image will be lost.

- **Acronis Disk Director Lite** is a disk management tool that enables such operations as cloning disks, creating, deleting and converting basic and dynamic volumes, and some additional operations like converting a disk partitioning style from MBR to GPT and vice versa or changing disk label. Adding this component to the media will enable you to prepare a machine disk configuration before recovering the data.

3.4.1.16. Error handling

Specify how to handle errors that might occur during backup.

- **Do not show messages and dialogs while processing (silent mode)**

  The preset is **Disabled**.

  Corporate administrators need an option to continue backup despite any errors that might occur without the system popping up an error box. Details of the operation, including errors, if any, could be found in the operation log.

  With the silent mode enabled, the program will not display interactive windows. Instead, it will automatically handle situations requiring user intervention such as running out disk space (except for handling bad sectors, which is defined as a separate option). No prompts will be displayed, including those for removable media or overwriting data on a tape. If an operation cannot continue without user action, it will fail.

  Therefore, enable this feature if you do not want unattended backup operations hang on pop-ups and errors.

- **If an error occurs, re-attempt**

  The preset is **Enabled**.
  This option is not available when backing up from bootable media.

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 networked location becomes unavailable or not reachable, the program will attempt to reach the location every 5 seconds, but no more than 30 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 enabled, the program will display a pop-up window each time it comes across a bad sector and ask for user decision 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 image and extract valid files to another disk.

3.4.1.17. Dual destination

This option is available when the primary destination is a **local folder or Acronis Secure Zone** and the secondary destination is **other than a managed location**.

This option is not available when backing up under bootable media.

The preset is: Disabled

When 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 laptops to the Acronis Secure Zone while on the road. When the laptop 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 the 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 backup option determines the program behavior in case a triggering event for a backup task occurs but the condition (or any of multiple conditions) is not met.

Wait until the conditions are met (preset)

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 getting risky, you can set the time interval after which the task will run irrespective of the condition. Select the check box for Run the task anyway after and specify the time interval. The task will start as soon as the conditions are met OR the maximum delay time lapses, depending on which comes first.

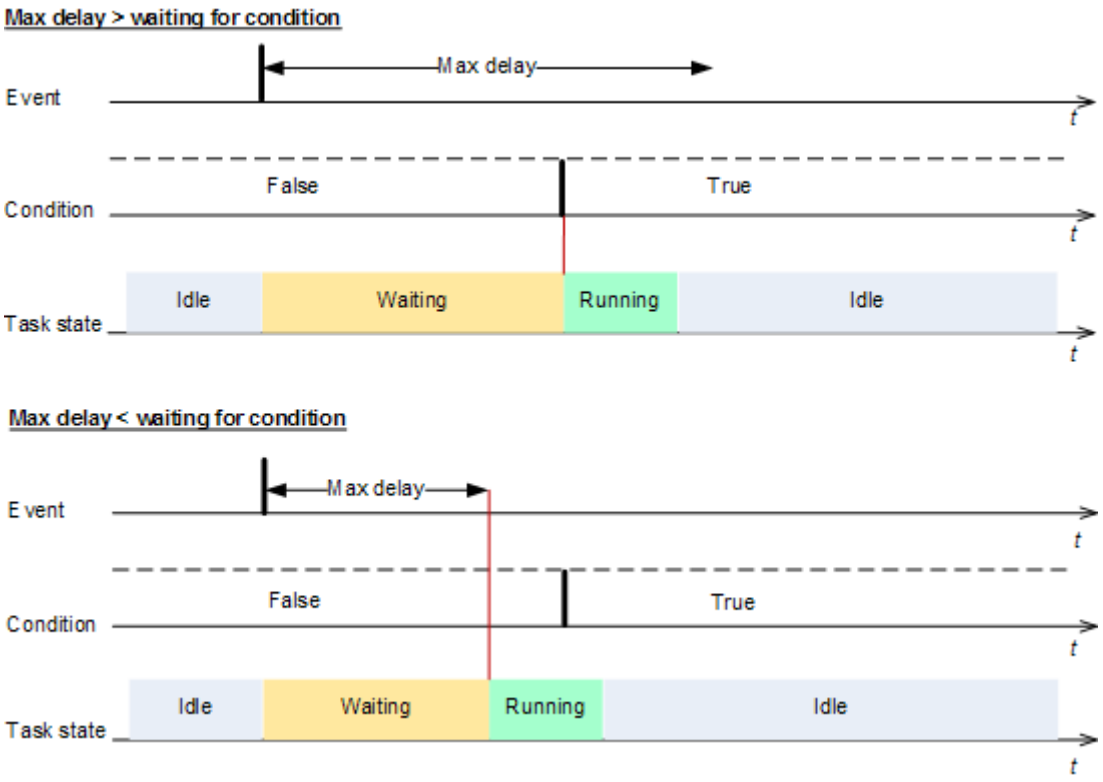

Skip the task execution

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 are relatively often.

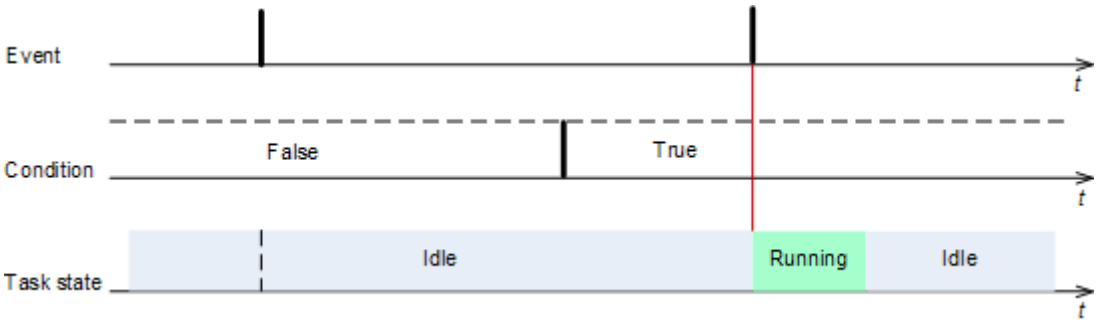

3.4.1.19. Task failure handling

This backup option determines the program behavior when any of backup plan's tasks fails.

Stop executing the backup plan

The plan's schedule will be temporary disabled so that you have time to find out and eliminate the failure cause. None of the plan's tasks will start on schedule until you resume the plan execution with the Restart button on the Backup plans page.

Continue executing the backup plan (preset)

The plan's tasks will be executed on schedule as though the failure had not occur. In addition, the program can try to execute the failed task again. Specify the time interval between attempts and number of attempts. The number of attempts should be reasonable so that the task could start on the next coming event. The program stops trying as soon as an attempt completes successfully OR the specified number of attempts is performed, depending on which comes first.

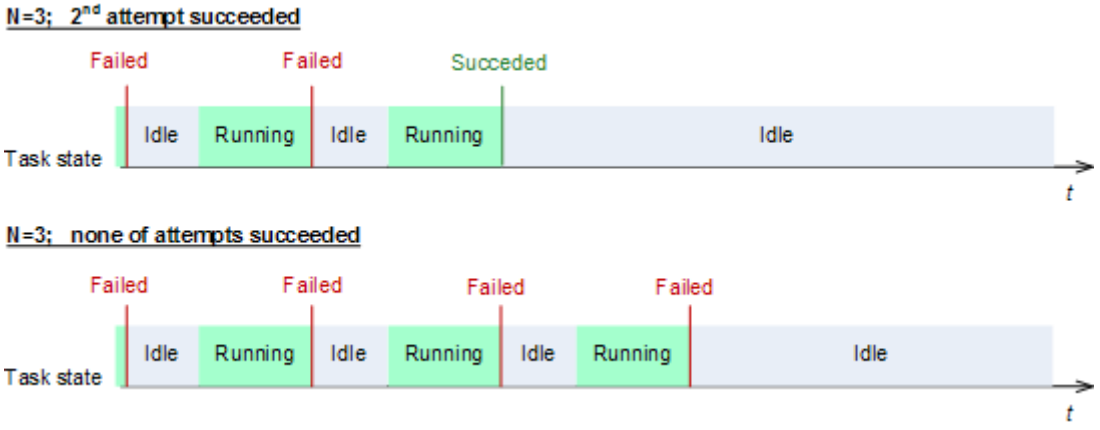

3.4.1.20. Additional settings

Specify the additional settings for the backup operation.

- **Overwrite data on a tape without prompting user for confirmation**

The preset is **Disabled**.
This option is available only when backing up to a tape device.

When pulling a tape from the Imported media pool, Acronis Backup & Recovery 10 will warn that you are about to lose data on the tape. To disable this warning, select this check box.

- **Dismount media after backup is finished**

  The preset is **Disabled**.
  This option is available only when backing up to a removable media (tape, CD/DVD, floppy disc.)

  The CD/DVD can be ejected or the tape can be dismounted after backup completes.

- **Ask for first media while creating backup archives on removable media**

  The preset is Enabled.
  This option is available only when backing up to removable media.

  You can choose whether to display the Insert First Media prompt when backing up to removable media. With the default setting, 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. Therefore, you should disable the prompt when scheduling a backup to removable media. Then, if the removable media is available (for example, CD-R/RW inserted), the task can run unattended.

- **Reset archive bit**

  The preset is **Disabled**.
  The option relates only to file-level backup and Windows operating systems.

  In Windows operating systems, each file has an attribute File is ready for archiving, available at selecting file -> Properties -> General -> Advanced -> Archive and Index attributes. This attribute, also known as 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 copy. Archive bit value is used by various applications such as databases.

  With Reset archive bit enabled, Acronis Backup & Recovery 10 will reset 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.

- **Create full backups as synthetic backups**

  The preset is **Disabled**.

  When this option is enabled, each full backup created by the backup plan is obtained by consolidating the newly backed up data changes (incremental or differential backup) with the latest backups up to the latest full backup. Thus the full backup is created without transferring all backed up data through the network. The maximum resource saving is achieved when backing up to a managed location, since the consolidation will be performed by the storage node without taking resources from the managed machine.

  This option enables you to have an archive consisting of full backups which makes for faster recovery and faster cleanup because full backups do not depend on each other and can be readily recovered and easily deleted.

3.4.2. Default recovery options

Each of Acronis agents has its own default recovery options with preset values. To access the default options, select **Options - Agent options - Default recovery options** from the menu. Once you have changed a preset value, it becomes default for this agent and will be used as default in further recovery operations. Nevertheless, while creating a recovery task, you can override the default settings with the custom settings that will be specific for this task only.

Availability of the recovery options

The set of available recovery options differs depending on:

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

The following table shows the summary of the recovery options availability.

	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. 118)	+	+	+	+	PE only	PE only
Recovery priority (p. 119)	+	+	+	+	-	-
File-level security (p. 120):						
Recover files with their security settings	-	+	-	+	-	+
Error handling (p. 122):						
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Additional settings (p. 123):						
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. 120)	+	+	+	+	-	-
Win Pop-up (p. 121)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 121)	+	+	-	-	-	-
SNMP (p. 122)	+	+	+	+	-	-

3.4.2.1. Pre-post commands

Specify commands or executable files to be automatically executed before and after the recovery procedure. For example, you may want to 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'.)

To specify pre/post commands

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

Pre-recovery command

To specify a command/batch file to be executed before 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.

	Selection			
Fail the task if the command execution fails	☑	☐	☑	☐
Do not recover until the command execution is complet	☑	☑	☐	☐
Result	**Default** Start the recovery after the command is successfully executed. Do not start and fail the task if the command execution failed.	Start the recovery after the command is executed despite of the execution failure or success.	N/A	Run the recovery concurrently with the command and independently 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 the command successful execution is critical for you, select the check box for **Fail the task if the command execution fails**. In case the command execution fails, the task run result will be set to Failed.

   By default the check box is not selected. With this setting, 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.

Keep in mind that the path to the batch file and the working directory you specify must exist and be accessible from the managed machine after the recovery is completed.

3.4.2.2. Recovery priority

This option is not available when recovering under bootable media.

The preset is: Normal

The priority of a process running in a system determines the amount of CPU usage and system resources allocated to that process. Decreasing the recovery priority will free more resources for other CPU tasks. Increasing the recovery priority might speed up the recovery due to taking resources from the other currently running processes. The effect will depend on total CPU usage and other factors.

To specify the recovery priority

Choose one of the following:

- **Low** – to minimize the recovery process speed, unleashing resources for 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

The preset is **Recover files with their security settings**.
The option relates to recovering files from file-level backups.

If the files' security settings were preserved during backup (p. 111), you can choose whether to recover files' security settings or let the files inherit the security settings of the folder where they will be recovered.

3.4.2.4. Notifications

Acronis Backup & Recovery 10 provides the ability of notifying recipients when the recovery is finished through:

- E-mail (p. 120)
- Windows Messenger (WinPopup) (p. 121)

E-mail

The preset is **Disabled**.
This option is not available when recovering from bootable media.

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 options as follows:
   - **Upon operation's successful completion** – to send notification when the operation is completed successfully
   - **Upon operation failure** – to send notification when the operation failed
   - **Add full log to the notification** - to add the full operation log to the message
4. 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 intact, messages will be constructed as if they are from the destination address.
   - **Use encryption** – to activate encryption for the notification message. SSL and TLS encryption methods are available for selection.
   - Check **Use specified outgoing e-mail server** to enable SMTP server and to set up its settings
     - **Outgoing mail server (SMTP)** - provide the outgoing SMTP server name.
     - **Port** – set port of the outgoing mail 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.

Windows Messenger (WinPopup)

Preset is **Disabled**.
This option is not available when recovering from bootable media.

Before configuring Windows Messenger notification, make sure the Messenger service is enabled on both the machine executing the task and the machine that will receive messages.

The Messenger service is disabled by default in Windows Server 2003 family. Change the service Startup mode to Automatic and start the service.

To configure Windows Messenger:

1. Select the **Send WinPopup notifications** check box to activate notifications.
2. In the **Computer name** field, provide the name of the machine to which notifications will be sent.
3. Under **Send notifications**, select the appropriate options as follows:
   - **When backup/recovery completes successfully**– to send notification when the operation is completed successfully
   - **When backup/recovery fails** – to send notification when the operation failed
   - **When user interaction is required** – to send notification during the operation when user interaction is required (always selected)
4. Click **Send Test Message** to check if the settings are correct.

3.4.2.5. Event tracing

It is possible to duplicate log entries for the events issued during the Acronis Backup & Recovery 10 restore operations in the Windows Application Event Log (Windows event log (p. 121)) or send the event notifications using SNMP (SNMP (p. 122)).

Windows event log

This option relates to Windows operating systems only.

This option is not available when recovering data under bootable media.

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

Acronis Backup & Recovery 10 Agents can record events concerned with the recovery operations in the Windows Application Event Log (to see this log, run **eventvwr.exe** or select **Control Panel > Administrative tools > Event Viewer**).

To select whether to log recovery 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 option setting specified for the machine (Machine options (p. 94)).

- **Log the following event types** - to log events concerned with the recovery operations in the Windows Application Event Log. Specify types of the 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 concerned with the recovery operations in the Windows Application Event Log.

SNMP notifications

This option applies to both Windows and Linux operating systems.

This option is not available when operating under bootable media.

The preset is **Disabled**.

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

1.3.6.1.4.1.24769.100.200.1.0 - string identifying a type of occurred event (Information, Warning, Error)

1.3.6.1.4.1.24769.100.200.2.0 - string containing text description of occurred event (it looks identically to messages published by Acronis Backup & Recovery 10 in its log.)

To select whether to send SNMP messages related to the recovery operations

Choose one of the following:

- **Use the setting set in the Machine options** - to use options specified for the Machine (Machine options (p. 94)).
- **Send SNMP notifications individually for recovery operation events** - to ignore the Machine options (for operations performed within the recovery task only) and select the appropriate options as follows:
  - **Types of events to send** - choose types of events to be reported: **All events**, **Warnings and Errors**, or **Errors only**.
  - **Server name or its IP** – type the name or IP address of the host running the SNMP management application, to which messages will be sent.
  - **Community** – type the name of SNMP community to which both the host running SNMP management application and the machine executing the recovery task belong. The typical community is "public".
- **Do not send SNMP notifications** - to disable SNMP messages about operations performed within the recovery task.

3.4.2.6. Error handling

Set up how the program should act if errors occur.

- **Do not show messages and dialogs while processing ("silent" mode)**

  The preset is **Disabled**.

With the silent mode enabled, the program will not display interactive windows. Instead, it will automatically handle situations requiring user intervention. No prompts will be displayed, including those for inserting removable media or the next tape. If an operation cannot continue without user action, the task will fail. Details of the operation, including errors, if any, could be found in the operation log.

- **If an error occurs, re-attempt**

  The preset is **Enabled**.

  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 networked location becomes unavailable or not reachable, the program will attempt to reach the location every 5 seconds, but no more than 30 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

Set up additional settings:

- **Set current date and time for recovered file**

  The preset is **Enabled**.

  Choose whether to recover files' date and time from the archive or assign the files the current date and time.

- **Validate backup archive before recovery**

  The preset is **Disabled**.

  Before data is recovered from the archive, Acronis Backup & Recovery 10 can check its integrity. If you suspect that the archive might have been corrupted, select this check box.

- **Check file system after recovery**

  The preset is **Disabled**.

  Having recovered a disk/partition from an image, Acronis Backup & Recovery 10 can check the integrity of the file system. To do so, select this check box.

  ***Windows**:*
  Verification of the file system is available only when recovering disk/partitions under Windows and for FAT16/32 and NTFS file systems.

  ***Linux**:*
  Verification of the file system is available only when recovering disk/partitions under Linux (i.e. not booted from the rescue CD) and only for Ext2, Ext3, Reiser4, ReiserFS, Linux Swap, XFS and JFS file systems.

- **Reboot the computer automatically if it is required for recovery**

  The preset is **Disabled**.

  The system disk's recovery requires rebooting the machine. Acronis Backup & Recovery 10 may can reboot the target machine automatically without user interaction. To do so, select this check box.

- **Change SID after the recovery is finished**

  The preset is **Disabled**.

Acronis Backup & Recovery 10 can generate a unique security identifier (SID) for the recovered system. You do not need a new SID when recovering a system on the same computer where the image was taken from or when creating a full duplicate 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.

The SID can remain unchanged if there is no computer with the same SID in the same workgroup or domain. Also, it is recommended not to change the SID if the system image is recovered on the same computer where the image was taken from.

3.4.2.8. VM power management

These options apply to virtual machines residing on the virtualization servers.

These options are available only if an Acronis agent for the virtual environment 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 setting 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 virtual machine recovery, there is a chance of appearing of two machines with the same security identifiers in the domain. To be on the safe side, power on the machine manually after you take the necessary precautions.

Select the check box for this setting if automatic powering on the machine is required.

4. Archive locations

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

A location 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 location size or number of backups in a location. You can limit each archive size using the cleanup, but the total size of archives stored in a location is limited by the storage size only.

Why create locations?

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

Quick access to the location

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 locations will be available for quick access without drilling down through the folders tree.

Easy archives management

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

- get a list of backups included in each archive
- recover data from a backup
- examine a backup content
- validate all archives in the location or individual archives or backups
- export any backup, including incremental and differential ones, to a new archive containing a single self-sufficient full backup
- mount a volume backup to copy files from the backup to a physical disk
- safely delete archives and backups from the archives.

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

The operation of creating a location results in adding the location name to the Locations section of the Navigation pane.

Centralized and personal locations

A shortcut to a location can be created directly on a managed machine or deployed from the management server.

Personal location

A location is called personal if the shortcut to this location was created directly on a managed machine, either locally or remotely, using the direct console-agent connection. Personal locations are specific for each managed machine.

See the Personal locations (p. 133) section for details.

Centralized location

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

See the Centralized locations (p. 126) section for details.

Key elements of the Locations view

 **Locations** (on the navigation pane) - top element of the locations tree. Clicking this item displays a list of both centralized and personal location groups:

 **Centralized** – for centralized locations. Clicking this item displays a list of centralized locations added by the management server administrator.

Double-click any personal location to open the detailed view of the centralized location (p. 127) and perform operations with the location (p. 128), archives (p. 136) and backups (p. 137) stored in there.

 **Personal** – for personal locations. Clicking this item displays a list of local or networked archive locations created using direct console connection to a managed machine.

Double-click any personal location to open the detailed view of the personal location and perform operations with the location (p. 134), archives (p. 136) and backups (p. 137) stored in there.

4.1. Centralized locations

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

As soon as the management server administrator commits creating a centralized location, the location path and name are distributed to all machines registered on the server. The shortcut to the location appears on the machines in the Centralized locations list. Any backup plan existing on the machines, including local plans, can use the centralized location.

On a machine that is not registered on the management server, a user having the privilege to back up to the centralized location can do so by specifying the full path to the location. In case the location is managed, user's archives will be managed by the storage node as well as other archives stored in the location.

4.1.1. Key elements of the "Centralized location" view

The key elements of the **Centralized location** view

Location's toolbar

Contains operational buttons and lets you perform operations with the location.

Related topic: Operations with centralized locations (p. 128)

The flow chart with legend

<picture>

The **flow chart** lets you estimate the location's daily load. It demonstrates the location's free space and used space with daily resolution over a period of time.

"Gray" - Free space: space on the storage device (selected for location) which is available for creating backup archives.

"Green" - Used space: total size of backup archives in this location.

A sum of "used" and "free" is not total size of a location, because if another user creates an archive in this location, the free size will be decreased, but the used space will not be changed.

The **legend** displays the location's information actual on the date selected in the flow chart:

- [for managed location only] the name of the storage node that manages the location
- [both] the full the path to the location
- [for managed locations only] deduplication (if enabled) and the ratio of the deduplicated data
- [for managed locations only] encryption (if enabled)
- [both] total number of the archives and the backups stored in the there

Location content

Contains archives table and toolbar. The archives table displays archives and backups that are stored in the location. Archives toolbar lets you perform operations with selected archives and backups. The list of backups is expanded by clicking a "plus" 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/volume backups (images).
- The **File archives** tab lists all the archives that contain file backups.

Related topics:

Operations with archives stored in a location (p. 136)

Operations with backups (p. 137)

Filtering and sorting archives (p. 138)

Bars of the Actions and Tools pane

- **[Location Name] Actions** bar available when clicking the location in the locations tree. Duplicates operations of the location's toolbar.
- **[Archive Name] Actions** bar available when you select an archive in the archives table. Duplicates operations of the archives list.
- **[Backup Name] Actions** bar available when you expand the archive and click on any backup within the archive. Duplicates operations of the archives table.

4.1.2. Operations with centralized locations

All the operations described here are performed by clicking the corresponding buttons on the location's toolbar. These operations can be also accessed from the **[Location Name] Actions** bar (on the **Actions and Tools** pane) and from the **Location** item of the main menu.

The following is a brief guidance for you to perform operations with centralized locations.

To	Do
Create a managed or an unmanaged location	1 Click **Create**. 2 In the **Type** field, select the location's type: **Managed** or **Unmanaged** The procedure of creating centralized locations is described in-depth in the following sections: • Create a managed centralized location (p. 129) • Create an unmanaged centralized location (p. 131)
Edit a managed or an unmanaged location	1 Select the location. 2 Click **Edit**. Depending on the location you select (managed or unmanaged), the respective edit page will be opened: • **Edit managed location** page lets you edit the location comments, and change the encryption password (if the location is encrypted) • **Edit unmanaged location** page lets you edit the location comments only.
Validate a location	1 Select the location. 2 Click **Validate**. You will be taken to the Validation (p. 208) page with already pre-selected location as a source. The location validation checks all the archives within the location.
Delete a location	1 Select the location. 2 Click **Delete**. You'll be asked whether to keep the archives stored in there, or delete the location along with the archives. The plans and tasks that use this location will fail. If you choose to keep archives for a managed location, the location will be detached from the storage node. Later on, you'll be able to attach this location to the same or to another storage node.
Explore unmanaged location.	1 Select the location. 2 Click **Explore**.

	The location will be available for examination in the standard Explorer window. Exploring of managed locations is not available.
Attach the managed location that was deleted without removing its content.	Click **Attach**. The procedure of attaching a managed location to a storage node is described in-depth in the Attaching a managed location (p. 132) section.
Change user's credentials for location	Click **Change user**. Changing user's credentials is available for locations that reside on shared storages only.
Refresh a location's information	Click **Refresh**. While you are reviewing the location content, archives can be added to the location, deleted or modified. Click **Refresh** to update the location information with the most recent changes.
Define tape labels and perform inventorying of a tape library on a managed location	Click **Manage tapes**. In the **Define tape labels** window, define labels for that tapes and refresh inventory. For more details see the Managing tape library section.
Rescan tapes in a managed location on a tape library	Click **Rescan tapes**. Rescan performs reading information about content of user-selected tapes and updating the database. This operation is described in-depth in the Rescan section.

4.1.2.1. Creating a managed centralized location

To create a managed centralized location, perform the following steps

Location

Name

Specify a unique name for the location. Creating of two centralized locations with the same name is prohibited.

Comments

Enter the distinctive description of the location.

Type

Select the **Managed** type.

Storage node

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

Path (p. 130)

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

Database path (p. 131)

The storage node will create in its local folder a location-specific database. This database will store the metadata required for cataloguing the archives and performing deduplication. Specify path to the folder where to create the location's database. By default, the database is created in %ALLUSERSPROFILE%\Application Data\Acronis\Storage Node.

Deduplication

Select whether to enable archives deduplication in the location. Deduplication minimizes storage space taken by the archives, backup traffic and network usage during backup. It reduces the size of archives in the location 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. 74) section.

Encryption (p. 131)

Select whether to protect the location with encryption. Every write to the location will be encrypted and every read from it will be decrypted transparently by the storage node, using a location-specific encryption key stored on the storage node.

After you performed all the required steps, click **OK** to commit creating the location.

The newly created centralized location will appear in the **Centralized locations** list on both the management server and the registered machines. Any backup policy or a backup plan created on either the management server or on the machines, including local plans, will be able to use the location.

On a machine that is not registered on the management server, a user having the privilege to back up to the centralized location can do so by specifying the full path to the location. In case the location is managed, user's archives will be managed by the storage node as well as other archives stored in the location.

Location path

To specify path where the managed location 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 locations can be organized on:
   - **Local folders** of the storage node
   - **Tape device**, locally attached to the storage node.
   - **Network places**, such network shares, SAN or NAS.

   To create a new folder for the location, click **Create folder**.

   Location can be created only in empty folder.

2. Click **OK**.

Location database path

To specify the path where the location'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 **Create folder**.

2. Click **OK**.

Location encryption

Every write to the managed location will be encrypted and every read from it will be decrypted transparently by the storage node, using a location-specific encryption key stored on the storage node. In case the storage medium is stolen or accessed by unauthorized person, the malefactor will not be able to decrypt the location 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 location with encryption

1. Select the **Encrypt** check box.
2. In the **Password** field, type a password
3. In the **Confirm the password** field, re-type the password.
4. Choose one of the following:
   - **None** – do not use encryption
   - **AES 128** – to use 128 bit encryption
   - **AES 192** – to use 196 bit encryption
   - **AES 256** – to use 256 bit encryption
5. Click **OK**.

The bigger the key size, the longer time it will take for program to encrypt the data and the greater your location's content security will be.

AES cryptographic algorithm operates in CBC mode and uses the randomly generated key with a user-defined size of 128, 192 or 256 bits. The encryption key is in turn encrypted with AES-256 using SHA-256 hash of the password as a key. The password itself is not stored anywhere on the disk or in the location, the password hash is used for verification purposes. With this two-level security, the location's data is protected from any unauthorized access, but the lost password recovery is not possible.

4.1.2.2. Creating an unmanaged centralized location

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

Location

Name

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

Comments

Enter the distinctive description of the location.

Type

Select the **Unmanaged** type.

Path (p. 132)

Specify where the location will be created.

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

As soon as the Management Server administrator commits creating a centralized location, the location path and name is distributed to all machines registered on the server. The shortcut to the location appears on the machines in the Centralized locations list. Any backup plan existing on the machines, including local plans, can use the centralized location. On a machine that is not registered on the Management Server, a user can add a centralized location to the Personal locations list, if the user has the privilege to back up to the location.

Location path

To specify path where the managed location 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 locations can be organized on:
   - **Network folders**, such network shares, SAN or NAS.
   - **FTP** and **SFTP** servers

   An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

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

   To create a new folder for the location, click  **Create folder**.

   Location can be created only in empty folder.

2. Click **OK**.

4.1.2.3. Attaching a managed location

When deleting a managed location from a storage node, you have the option to keep the archives contained in the location. A location that was deleted without removing the archives, can be attached to the same or another storage node. You might need to do so when retiring the storage node server, when the storage node is lost, when balancing loads between storage nodes.

Personal or centralized unmanaged locations cannot be attached.

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

Location

Storage node

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

Path

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

If the location is encrypted, provide the encryption password in the pop-up dialog box.

Database path

The storage node will create in its local folder a location-specific database. This database will store the metadata required for cataloguing the archives and performing deduplication. Specify path to the folder where to create the location's database. By default, the database is created in %ALLUSERSPROFILE%\Application Data\Acronis\Storage Node.

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

The attached location will appear in the **Centralized locations** list on both the management server and the registered machines. Any backup policy or a backup plan created on either the management server or on the machines; including local plans; will be able to use the location.

On a machine that is not registered on the management server, a user having the privilege to back up to the centralized location can do so by specifying the full path to the location. User's archives will be managed by the storage node as well as other archives stored in the location.

4.2. Personal locations

A location is called personal if the shortcut to this location was created directly on a managed machine, either locally or remotely, using the direct console-agent connection. Personal locations are specific for each managed machine.

Personal locations can be organized on detachable or removable media. Acronis Secure Zone, removable media, locally attached tape devices are considered as a personal locations available to all users that can log on the system.

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

Sharing a personal location

Multiple machines can refer to the same physical location, say, to the same shared folder, but each of the machines has its own shortcut in the **Locations** tree. Users that back up to a shared folder can see and manage each other's archives according to their access permissions in that folder. To ease the archives identification, the Locations page has the Owner column that displays the owner of each archive. To find out more about the owner concept see Owners.

4.2.1. Key elements of the "Personal location" view

Location's toolbar

Contains operational buttons and lets you perform operations with selected personal location.

Related topic: Operations with personal locations (p. 134)

Pie chart with legend

The **pie chart** lets you estimate the location's load: it demonstrates the ratio between location's free space and used space.

-  - Free space: space on the storage device (selected for location) which is available for creating backup archives.
-  - Used space: total size of backup archives in this location.

A sum of "used" and "free" is not total size of a location, because if another user creates an archive in this location, the free size will be decreased, but the used space will not be changed.

The **legend** displays the full path to the location and total number of the archives and the backups stored in the there.

Location content

Contains archives table and toolbar. The archives table displays archives and backups that are stored in the location. Archives toolbar lets you perform operations with selected archives and backups. The list of backups is expanded by clicking a "plus" 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/volume backups (images).
- The **File archives** tab lists all the archives that contain file backups.

Related topics:

Operations with archives stored in a location (p. 136)

Operations with backups (p. 137)

Filtering and sorting archives (p. 138)

Bars of the "Actions and tools" pane

- **[Location Name] Actions** bar available when clicking the location in the locations tree. Duplicates operations of the location's toolbar.
- **[Archive Name] Actions** bar available when you select an archive in the archives table. Duplicates operations of the archives list.
- **[Backup Name] Actions** bar available when you expand the archive and click on any backup within the archive. Duplicates operations of the archives table.

4.2.2. Operations with personal locations

To perform any operation (except for creation) with a location, you must select it first.

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

The following is a brief guidance for you to perform operations with personal locations.

To	Do
Create a personal location	Click **Create**. The procedure of creating personal locations is described in-depth in the Creating a personal location (p. 135) section.
Change user account for accessing a location	Click **Change user**. In the appearing dialog box, provide the credentials required for accessing the location.
Create Acronis Secure Zone	Click **Create ASZ**. The procedure of creating personal locations is described in-depth in the Creating Acronis Secure Zone (p. 217) section.
Explore a location's content	Click **Explore**. In the appearing Explorer window, examine the selected location's content.
Validate a location	Click **Validate**. You will be taken to the Validation (p. 208) page with already pre-selected location as a source. The location validation checks all the archives within the location.
Delete a location	Click **Delete**. Deleting a location actually removes only a shortcut to the folder from the **Locations** view. The folder itself remains untouched. You have the option to keep or delete backups contained in the folder.
Refresh a location's information	Click **Refresh**. While you are reviewing the location content, archives can be added to the location, deleted or modified. Click **Refresh** to update the location information with the most recent changes.

4.2.2.1. Creating a personal location

Adding a personal location will create a shortcut to the place where the backup archives are stored or will be stored.

To add a personal location

1. Do one of the following:
   a. On the **Actions** bar, click **Add Personal Location**;
   b. Right-click the **Locations** tree or **Personal** group, and then select **Add Personal Location** from the context menu.
2. On the **Add Personal Location** page

1. In the **Path** field, click **Change...** In the opened **Selecting Path** window, specify a path to the location being added, and then click **OK**.
2. In the **Name** field, click **Change...** and type a new name for the location, then click **OK**.
3. Click **OK**.

3. As a result, the created location appears in the **Personal** group of the locations tree.

4.2.2.2. Merge or move locations

What if I need to move the existing location from a one place to another?

Proceed as follows

1. Make sure the none of the backup plans uses the existing location over the time of moving files, or temporary disable schedules of the given plans (see Temporary disabling a backup plan).
2. Move the location folder with all its archives to a new place manually by means of a third-party file manager.
3. Add the new location.
4. Edit backup plans and tasks that use the existing location so they use the new one.
5. Delete the old location.

How can I merge two locations?

Suppose, you have two locations A and B in use. Both locations are used by backup plans. You decide to leave only the location B, moving in there all the archives from the location A.

To do this, proceed as follows

1. Make sure the none of the backup plans uses the location A over the time of merging, or temporary disable schedules of the given plans (see Temporary disabling a backup plan).
2. Move the archives to the location B manually by means of the third-party file manager.
3. Edit the backup plans that use the location A: redirect the backup destination to the location B.
4. In the locations tree, select the location B to check whether the archives are displayed. If not, click **Refresh**.
5. Delete the location A.

4.3. Common operations

4.3.1. Operations with archives stored in a location

To perform any operation with an archive, you must select it first. If the archive is protected with a password, you will be asked to provide it in the pop-up dialog box.

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 **Location** item of the main menu.

The following is a guidance for you to perform operations with archives stored in a location.

To	Do
Validate an archive	Click **Validate**. The **Validation (p. 208)** page will be opened with the pre-selected archive as a source. Validation of an archive will check all the archive's backups.
Delete a single archive or multiple archives	1 Select the archive or one of the archives you want to delete. 2 Click **Delete**. The program duplicates your selection in the **Backups deletion** (p. 138) 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 location	Please be aware that if filters have been applied to the locations list, you see only a part of the location content. Be sure that the location does not contain archives you need to retain before starting the operation. Click **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 must select it first. To select a backup, expand the archive, then click on the backup. If the archive is protected with a password, you will be asked to provide it in the pop-up dialog box.

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 **Location** item of the main menu.

The following is a guidance for you to perform operations with archives backups.

To	Do
View backup's content in a separate window	Click **View content**. In the **Backup Content** window, examine the backup's content.
Recover	Click **Recover**. The **Recovering Data** (p. 188) page will be opened with pre-selected archive/backup as a source.
Validate a backup	Click **Validate**. The **Validation (p. 208)** 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 to a new archive	Click **Export**. The **Export backup** page will be opened with pre-selected backup as a source. All you need to do is to set up the export destination.

	Use the export operation when you need to get a new archive containing a single self-sufficient full backup from any existing backups, including incremental and differential ones.
Delete a single or multiple backups	1 Select the backup one of the backups you want to delete. 2 Click **Delete**. The program duplicates your selection in the **Backups deletion** (p. 138) 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 location	Please be aware that if filters have been applied to the locations list, you see only a part of the location content. Be sure that the location does not contain archives you need to retain before starting the operation. Click **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.

4.3.3. Deleting archives and backups

The **Backups deletion** window displays the same tab as for the location view, but with check boxes for each archive and each backup. The archive or backup you have chosen to delete has the check mark.

The window inherits filters from the archive list. Thus, if filters have been applied to the archives list, here you see see only the archives and backups corresponding to the filters. To see all the tab's content, clean all the filter fields.

Review the archive or backup that you have selected for deletion. If you need to delete other archives and backups select the respective check boxes, then click **Delete selected** to confirm the deletion.

What happens if I delete a backup that is a base of an incremental or differential backup?

To preserve the 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 the deletion. The resulting backup will not contain data that was present in the backup you have deleted and was absent in the incremental or differential backup you have retained.

There should be enough space in the location 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 guidance for you to filter and sort archives in the archives table.

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

Configuring the archives table

The table has seven columns that are displayed by default: **Archive name**, **Owner**, **Machine**, **Created**, **Size**, **Backups**, **Comments**. The **Encrypted** column is hidden. If required, you can hide the shown 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 column headers presenting 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.

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 time interval)	+	+
User logon (any user, current user, specify the user's account)	+	-
User logoff (any user, current user, specify the user's account)	+	-
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)	+	-

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 the 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 screensaver is running or the machine is locked)	+	-
Archive location is available	+	+
Specific network is available (select the network connection)	+	+
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 had 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 condition backup option.

"What if"s

- What if an event occurs (and a condition, if any, is met) while a task initiated by the previous event is running?

  The task will be waiting and start as soon as the running task completes. The state of condition at this moment does not matter. Any event that may occur while the task is waiting 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 very long time?

  When a task is waiting for a condition or for human interaction, the Dashboard shows a message like "Task_N is waiting for...". If delaying a backup is getting risky, you can force the condition (tell the users to log off) or run the task manually on seeing the warning. 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

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 **Frequency** area, choose 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 for 12 times: from 10 AM to 10 PM during a day.

In the **Duration** area, set the following settings:

Start date: <...>	Set up a date when this schedule will be enabled (an effective date). If this check box is deselected the task will be started on the nearest day and time you have specified above.
End date <...>	Set up a date when this schedule will be disabled. If this check box is deselected, the task will be run for infinite days.

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

Examples

"Simple" daily schedule

Run the task every day at 6PM.

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

1. Every **1** day(s).
2. Once at **06:00:00 PM**.
3. Start date: **empty**. The task will be started on the current day, if it has been created before 6PM. If you have created the task after 6PM, the task will be started for the fist time on the next day at 6PM.
4. End date: **empty**. The task will be performed for infinite 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, 09.15.2009), 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 **00:00:00 AM** until **09:00:00 PM** - thus, the task will be performed 8 times a day with 3 hour time interval. After the last daily recurrence at 9PM, the next day comes and the task starts over again from 0AM.
3. **Start date**: **09/15/2009**. If September 15, 2009 is the current date of the task's creation and, say, the **1:15PM** is the task's creation time, the task will be started when the nearest time interval comes: at **3PM** in our example.
4. **End date**: **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 a one task

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

For example, the task has to be run every 3 days, starting from 09/20/2009. During a day the task should be repeated 5 times:

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

The obvious way is to add five simple schedules. If you spend a 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 fifths is equal to 2 hours. In this case, the optimal way is to add two schedules to the task.

First daily schedule

1. Every **3** day(s).
2. Every **4** hours
   From **08:00:00 AM** until **12:00:00 PM**.
3. Start date: **09/20/2009**.
4. End date: **empty**.

Second daily schedule

1. Every **3** day(s).
2. Every **2** hours
   From **03:00:00 PM** until **07:00:00 PM**.
3. Start date: **09/20/2009**.
4. End date: **empty**.

5.2. Weekly schedule

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 days of 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 **Frequency** area, choose one of the following:

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

In the **Duration** area, set the following settings:

Start date: <...>	Set up a date when this schedule will be enabled (an effective date). If this check box is deselected the task will be started on the nearest day and time you have specified above.
End date <...>	Set up a date when this schedule will be disabled. If this check box is deselected, the task will be run for infinite months.

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

Examples

"One day in a week" schedule

Run the task every Friday at 10PM, starting from a certain date (say 05/15/2009) 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. Start date: **05/15/2009**. The task will be started on the nearest Friday at 10PM.
4. End date: **11/15/2009**. On this date the task will be performed for the last time, but the task itself is still available in the Tasks view.

This schedule is widely used when creating a custom backup scheme. The "One day in a week"-like schedule is added to the full backups, while the incremental backups are scheduled to perform on workdays. For more details, see the Full and incremental backups plus cleanup example in the Custom backup scheme (p. 183) section.

"Workdays" schedule

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

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

1. Every **1** week(s) on: **<workdays>** - selecting the <workdays> check box automatically marks out all the workdays of a week (Mon, Tue, Wed, Thu, Fri), leaving weekend (Sat, Sun) unmarked.
2. Once at: **09:00:00 PM**.
3. Start date: **empty**. If you have created the task, say on Monday at 11:30AM, the task will be started on the same day at 9PM. If the task was created, say on Friday after 9PM, then it will be started for the fist time on the nearest marked out workday (Monday in our example) at 9PM.
4. End date: **empty**. The task will be restarted for infinite 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 perform on day in a week. For more details, see the Full and incremental backups plus cleanup example in the Custom backup scheme (p. 183) section.

Several weekly schedules for a one task

In case when the task needs to be run on different days of 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 12PM and 9PM
- Tuesday: every 3 hours from 9AM till 9PM
- Wednesday: every 3 hours from 9AM till 9PM
- Thursday: every 3 hours from 9AM till 9PM
- Friday: twice at 12PM and 9PM
- Saturday: once at 9PM
- Sunday: once at 9PM

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. Start date: **empty**.
4. End date: **empty**.

Second schedule

1. Every **1** week(s) on: **Tue, Wed, Thu**.
2. Every **3** hours
   From **09:00:00 PM** until **09:00:00 PM**.
3. Start date: **empty**.
4. End date: **empty**.

Third schedule

1. Every **1** week(s) on: **Sat, Sun**.
2. Once at: **09:00:00 PM**.
3. Start date: **empty**.
4. End date: **empty**.

5.3. Monthly schedule

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 tasks on. You may also select the last day of the month, despite of its actual date.
On: <...> <...>	The task will be run on the selected days and weeks of the month you specified above.

In the **Frequency** area, choose one of the following:

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

In the **Duration** area, set the following settings:

Start date: <...>	Set up a date when this schedule will be enabled (an effective date). If this check box is deselected the task will be started on the nearest day and time you have specified above.

End date <...>	Set up a date when this schedule will be disabled. If this check box is deselected, the task will be run for infinite months.

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

Examples

"Last day of every month" schedule

Run the task once at 10PM 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 of its actual date.
3. Once at: **10:00:00 PM**.
4. Start date: **empty**.
5. End date: **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 perform 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. 183) section.

"Season" schedule

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

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. Start date: **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.
5. End date: **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.

Several monthly schedules for a one task

In 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.

The task goes into effect on the 11/01/2009.

- During winter the task runs once at 10PM on every workday.
- During spring and autumn the task runs every 12 hours on all workdays.

- During summer the task runs every first, fifteenth and last day of every month at 10PM.

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

First schedule

1. Months: **December, January, February**.
2. On: **<All> <workdays>**
3. Once at **10:00:00 PM**.
4. Start date: **11/01/2009**.
5. End date: **empty**.

Second schedule

1. Months: **March, April, May, September, October, November**.
2. On: **<All> <workdays>**.
3. Every **12 hours**
   From **12:00:00 AM** until **12:00:00 PM**.
4. Start date: **11/01/2009**.
5. End date: **empty**.

Third schedule

1. Months: **June, July, August**.
2. Days: **1, 15, Last**.
3. Once at: **10:00:00 PM**.
4. Start date: **11/01/2009**.
5. End date: **empty**.

5.4. At Windows Event Log event

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, such as **Microsoft Office Sessions**

Event source

Specifies the event source, which typically indicates the program or the system component that caused the event.

Event type

Specifies the event type, such as **Information**, **Warning**, or **Error**.

Event ID

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

Examples

"Bad block" emergency backup

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

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

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

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

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

Pre-update backup in Vista

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

Having downloaded one or more updates and scheduled their installation, the Microsoft Windows Vista operating system records an event with the event source **Microsoft-Windows-WindowsUpdateClient** and event number **18** into the **System** log; the type of this event is **Information**.

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

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

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

How to view events in Event viewer

To open a log in Event Viewer

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

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

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

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

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

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

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

5.5. Conditions

Conditions add more flexibility to the scheduler, enabling to execute backup tasks with respect to the certain conditions. Once a specified event occurs (see the Scheduling section for the list of the 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 condition backup option. There, you can specify how important are the conditions 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 is more prior - put the task on hold for the specified time interval. If the time interval lapses and the conditions are still not met, run the task anyway. With this setting, the program will automatically handle the situation when the conditions are not met for too long and further delaying the backup is undesirable.
- backup task start time matters - skip the backup task if the conditions are not met at the time when the task should be started. Skipping the task run makes sense when you need to back up data strictly at the specified time, especially if the events are relatively often.

Adding multiple conditions

Multiple conditions must be met simultaneously to enable the task execution.

Example:

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

- Event: When free space changed, Run task if free space has changed by at least: 1 GB.
- Condition: User logged off
- Condition: Time since last backup, Time since the last backup: 12 hour(s).
- Task start conditions: Wait until the conditions are met.

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

5.5.1. User is idle

Applies to: Windows

"User is idle" means that a screensaver 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.

5.5.2. Location is available

Applies to: Windows, Linux

"Location is available" means that the destination specified for storing archives on a networked drive is accessible for creating a backup.

Example:

Backing up data to the networked location is performed on workdays at 9PM. If the location is not available at that moment (for instance, due to maintenance works), 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 is available**
- Task start conditions: **Skip the task execution**.

As a result,

(1) if 9PM comes and the archive location is available, the backup task starts right on time.

(2) if 9PM comes but the archive location is unavailable, the backup task will start on the next workday if the location is available.

(3) If the location will never be available on workdays at 9PM, the task never starts.

5.5.3. Specific network

Applies to: Widows, Linux

Enables to put a backup task on hold, until the specific network connection becomes enabled.

Example:

Backing up data on the laptop to the corporate networked drive after 12 hours passed since the last backup has started, but only when the laptop is connected to the local corporate network (say, Corporate Local Area Connection). If the network connection is disabled (if the laptop is out of office), wait until it becomes available, and then start the backup task.

- Event: **Upon time since last backup**, Time since a previous backup: **12** hour(s).
- Condition: **Specific network**, Run the task only if the following network is available: **Corporate Local Area Connection**.
- Task start conditions: **Wait until the conditions are met**.

As a result, the backup task, scheduled to run after 12 hours since the last backup, may actually start later, as soon as the Corporate Local Area Connection is established. When another network connection is used (say, the user connects to the neighborhood network at home), the backup will not start.

5.5.4. Fits time interval

Applies to: Windows, Linux

Enables to put a backup task run within the specified time interval.

Example:

A company uses different locations on the same network-attached storage for backing up users data and servers. Workday starts at 8AM and end at 5PM. Users data should backed up as soon as users log off, but not earlier than 4:30 PM and not later than 10PM. 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 10PM, it is supposed that the backing up of users data does not take more than one hour. If a user is still logged in within the specified time interval, or logs off any other time – do not back up the user's data, i.e. skip the task execution.

- Event: **When logging off**, The following user: **Any user**.
- Condition: **Fits 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 04:30:00PM and 10:00:00PM, the backup task will start immediately.

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

What if the task is scheduled for execution at a certain time and this time is outside the specified time interval?

For example:

Event: **Weekly**, Every **1** day(s); Once at **03:00:00 PM**.

Condition: **Fits time interval**, from **06:00:00 PM** until **11:59:59 PM**.

Task start conditions: **Wait until the conditions are met**.

As a result, the task scheduled to run at 3PM will always start at 6PM.

5.5.5. User logged off

Applies to: Windows

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

Example:

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

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

5.5.6. 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 backup successful completion passes.

Example:

Run the backup task after free space on the managed machine is changed by at least 1GB, but only if more than 12 hours have passed since the last successful backup.

- Event: **When free space changed**, Run task if free space has changed by at least: **1 GB**.
- 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 free space changes for more than 1GB before 12 hours pass since the successful completion of the last backup, the scheduler will wait until 12 hours to pass and then start the task.

(2) if the free space changes for more than 1GB after 12 hours pass since the last backup successful completion, the backup task will start immediately.

(3) if the free space never changes for more than 1GB, the task will never start. You can start the backup manually, if need be, in the Backup plans and tasks view.

6. Direct management

6.1. Administering a managed machine

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 issues.

System state

Draws your attention to issues occurred on the machine and offers you ways of fixing or examining them. The most critical issues are displayed on the top. The examples below illustrate the types of messages you may observe.

	Description	Offer	Comment
	X tasks failed	Show tasks	**Show tasks** will open the **Tasks** view with failed tasks, where you can examine the reason of failure.
	X task(s) need(s) user interaction	Show tasks	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.) The similar information is displayed when a scheduled backup is waiting for the condition to start. If delaying the backup is getting risky, you can force the condition, say, tell the users to log off, or run the task manually.
	License check failed: "X days remaining until the software stops working. Please make sure you have a valid license on Acronis License Server"		
	Trial period is over	Enter the license key	
	Low free space on X locations	Show list	**Show location** will take you to the **Locations** view where you can examine the location size, free space, content and take the necessary steps to increase the free space.
	Bootable rescue media was not created	Create now	To be able to recover an operating system when the machine fails to boot, you must: 1 Back up the system volume (and the boot volume, if it is different) 2 Create at least one bootable media (p. 317). **Create** will launch the Bootable Media Builder (p. 324).
	No backups performed within last 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. The **Back up now** will take you to the **Create a Backup**

			Plan page where you can instantly configure and run the backup operation. To configure the time interval that is considered as critical, select **Options - Console options - Time-based alerts**.
	Not connected to the management server for 3 days	Show log	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. Click **Show log** to open the **Log** view, where you can examine the log entries and find out the reason of the issue.

If there are no issues, the system state is OK.

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 the date.

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

	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.

6.1.2. Backup plans and tasks

The Backup plans view keeps you informed of data protection on a given machine.

To find out what a backup plan is currently doing on the machine, check the backup plan state. Status of a backup plan helps you to estimate whether the data is successfully protected. You can start, stop, edit, delete local backup plans. Use the filtering and sorting capabilities to adjust the presentation to your liking.

A task is a set of sequential actions to be performed on a machine when a certain event occurs. The Backup plans and tasks view lets you monitor and manage tasks. You can view tasks' details, their states and execution results, as well as run, stop and delete tasks.

The key elements of the Backup plans and tasks view

- **Toolbar** with operational buttons - lets you perform operations with the backup plan or task you select.
- Backup plans and tasks table - displays a list of backup plans and tasks according to filters applied.
- **Information** pane (below the table) - displays the detailed information on the backup plan/task which is currently selected. This information is also duplicated in the Plan details (Task details) window. The pane is collapsed by default. To expand the panel, click the chevron.
- **Backup plan/task actions** bar (on the Actions and Tools pane) - contains operations you can perform with selected backup plan/task. The list of actions is also accessible from the Actions menu.

Way of working

- Use filters to display the desired backup plans/tasks in the backup plans table (see Filter and sort backup plans and tasks (p. 163)). By default, the table displays all the plans of the managed machine sorted by name.
- Select the backup plan/task to perform operation on it (see Perform actions with backup plans and tasks (p. 160)).
- [OPTIONAL] Configure backup plans and tasks table appearance (see Configuring backup plans and tasks table)

6.1.2.1. Concepts

Backup plan execution 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 location; 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 setting the maximum delay or forcing the condition (tell the user to log off, enable the needed 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 occasional 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 tasks overlapping may result from 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.

	State	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 the future failure in case a local plan has failed • Edit the backup policy on the management server in case a centralized plan has failed When creating a backup plan or policy the administrator can turn on the option to stop executing the backup plan as soon as the backup plan gets the Error status. The backup plan's execution can be resumed using the Restart button.
2	**Warning**	At least one task has succeeded with warnings. Otherwise, see 3.	View the log to read the warnings -> [optionally] Perform actions to prevent the future warnings or failure.
3	**OK**	All the tasks are completed successfully.	No action is required. Note that a backup plan can be OK in case none of the tasks has been started yet or some of the tasks are stopped or being stopped. These situations are considered as normal.

Task states and statuses

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 goes to the **Running** state when the event specified by the schedule occurs AND the condition set in the backup plan is met AND no other task that locks the necessary resources is running. In this case, nothing prevents the task from running.

Waiting

A task goes to the **Waiting** state when the task is about to start, but another task using the same resources is already running. Particularly this is done because two tasks cannot lock the same volume or make a snapshot at the same time. Once the other task unlocks the resource, the waiting task enters the **Running** state.

A task may also go 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. 114) 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 is Reboot, that can put the task to the **Running** state.)

Stopping

The user can stop a running task or the task that needs interaction. The task goes to the **Stopping** state and then to the **Idle** state. A waiting task can also be stopped. Since the task is not running, 'stop' means removing from the queue in this case.

Task state diagram

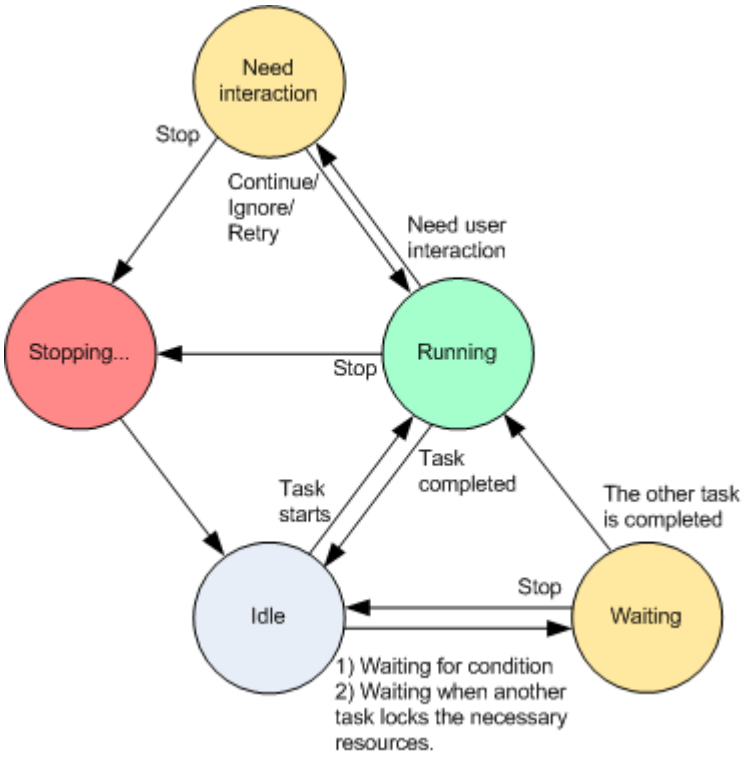

Task statuses

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 the future failure • Edit the local plan to prevent the 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	**OK**	Last result is "Succeeded", "No result yet", or "Stopped"	No action is required.

6.1.2.2. How to...

Perform actions on backup plans and tasks

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

To	Do
View details of a plan/task	**Backup plan** Click **View details**. The **Plan's Details** window will appear. **Task** Click **View details**. The **Task's Details** window will appear.
View plan's/task's log	**Backup plan** Click **View log**. You will be taken to the **Log** view containing the list of the plan-related log entries. **Task** Click **View log**. You will be taken to the **Log** view containing the list of the task-related log entries.
Run a plan/task	**Backup plan** Click **Run**. In the **Start Backup Plan** 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. *Why cannot 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.

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 – 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 plan/task	**Backup plan** Click **Edit**. *Why cannot 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 the centralized origin. The direct editing of centralized backup plans, deployed by backup policies is impossible. You need to edit the original backup policy. **Task** Click **Edit**. *Why cannot 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 **Delete**. *What will happen if I delete the backup plan?* The plan's deletion deletes all its tasks. *Why cannot I delete the backup plan?* • The backup plan in the "Running" state A backup plan cannot be deleted, if at least one its task 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 the centralized origin. A centralized plan can be deleted by the management server administrator by revoking the backup policy that produced the plan. **Task** Click **Delete**. *Why cannot 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 **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 be retrieved from the managed machine not immediately due to some latency. The manual refresh guarantees

	that the most recent data is displayed.

Filter and sort backup plans and tasks

To	Do
Set a number of tasks to display	Options -> ... … and set the desired number. Further filtering and sorting will performed with the given number of tasks only.
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/owner.	Type a plan's/task's name or an owner's name in the field below the corresponding header's 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 a field below the corresponding header, select the required value from the list.

Configuring backup plans and tasks table

The table has six columns that are displayed by default. Other columns are hidden. If required, you can hide the shown 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 column headers presenting in the table.
2. Click the items you want to be displayed/hidden.

Headers of the table are also clickable. Clicking a column header sorts the list by the corresponding attribute.

Temporary disable a backup plan

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

Applies to backup plans that use custom backup schemes only.

To disable a backup plan

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

6.1.3. Log

Log is used for viewing information on operations performed by the Acronis Backup & Recovery 10. With log, you can examine operations, results of tasks' execution including reasons for failure, if any.

The key elements of the Log view

- **Toolbar** with operational and filtering buttons, lets you to perform actions with the selected log entry.
- **Log table** - displays the list of log entries according to applied filters.
- **Information** pane (at the foot of the page) - displays the detailed information on the currently selected log entry. The panel is collapsed by default. To expand the panel, click the  chevron. This information can be also accessed by clicking on **View Details** item in the **Log Actions** bar.
- **Log Actions** bar (on the **Actions and Tools** pane) - contains actions you can perform with selected log entries. All actions are also duplicated in the menu and toolbar. Thus, you can choose the way optimal for your work.

Way of working with log entries

- Use filters to display the desired log entries only (see Filtering and sorting log entries (p. 165))
- Select a log entry (or log entries) to perform operation on it (see Operations with log entries (p. 164))
- Configure the log table appearance (see Configuring log table (p. 166))

Ways to open the Log view with the pre-filtered log entries

Having selected items in other administration views (Dashboard, Backup plans and 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 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. Operations with log entries

All the operations described below are performed by clicking the corresponding items on the Log Actions bar. In addition, the operations can be also performed with the context menu (by right-clicking the log), or with the toolbar (by clicking the corresponding buttons).

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

To	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 other 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	1 Select a single log entry or multiple log entries. 2 Click **Save Selected to File**. 3 In the opened window, specify a path and a name for the file.
Save all the log entries to a file	1 Make sure, that the filters are not set. 2 Click **Save All to File**. 3 In the opened window, specify a path and a name for the file.
Save all the filtered log entries to a file	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

6.1.3.2. Filtering and sorting log entries

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

To	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's buttons: to filter error messages to filter warning messages to filter information messages
Filter log entries by the original backup plan or task	Under **Backup plan** column header, select the backup plan from the list.
Filter log entries by owner	Type an owner's name in the field below the **Owner** header. As a result, in the drop-down list you will see the owner names 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.

6.1.3.3. Configuring log table

By default the table has five columns: Type, Date and time, Backup plan, Task, Message. If required, you can display other columns (Code, Module, and Owner) and/or hide the shown.

To show new or hide shown columns

1. Right-click any table header to open the context menu. Items corresponding to headers used in the table by default are ticked.
2. Click the items you want to be displayed/hidden.

6.1.3.4. Log entry details

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

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

Log entry data fields

A local log entry contains the following data fields:

- Type - Type of the event (Error; Warning; Information)
- Date - Date and time when the event took place
- Backup plan - The backup plan the event relates to (if any)
- Task - The task the event relates to (if any)
- Code - 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 backup plan owner (only under operating system)
- Message - The event text description.

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

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

6.2. Creating a backup plan

Before creating your first backup plan, please familiarize yourself with the basic concepts (p. 29) used in Acronis Backup & Recovery 10.

A backup plan specifies how the given data on a given machine will be protected. A backup plan specifies:

- What data to back up
- Where to store the backup archive
- Backup scheme
- [Optional] Validation rules
- Backup options

Physically, a backup plan is a bundle of tasks configured for execution on a managed machine. A backup plan is created by the local Acronis Backup & Recovery 10 administrator directly on the machine (local plan) or appears on the machine as a result of a backup policy deployment (centralized plan.)

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. 169)

[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 check box for **Advanced view**.

Plan comments

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

What to backup

Source type (p. 169)

Select the type of data to back up.

Items to backup (p. 170)

Specify the data items to back up: disks, volumes or files/folders.

Access credentials (p. 172)

[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 check box for **Advanced view**.

Exclusions (p. 172)

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

Where to back up

Archive (p. 173)

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.

Access credentials (p. 175)

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

Archive comments

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

How to back up

Backup scheme (p. 175)

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. 185)

[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 (p. 185)

[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 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 defaults**.

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

After that, you might be prompted for the password.

A scheduled or postponed task has to run regardless of users being logged on. Because 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.

6.2.1. 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 on 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 a name of an Active Directory user account, be sure to specify also the domain name (DOMAIN\Username.)
     - **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. 34) section.

6.2.2. Source type

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

Acronis Backup & Recovery 10 Agent for Windows (and for Linux)

When the Acronis Backup & Recovery 10 Agent for Windows (for Linux) is installed, you can back up:

- **Files** – select this option to back up specific files and folders.

If you are not concerned about recovery of operating system along with all 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** – select this option to back up disks and/or volumes. To be able to back up disks or volumes, you must have the 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 the backup process when it comes to backing up large volumes of data.

Acronis Backup & Recovery 10 Agent for Hyper-V or Acronis Backup & Recovery 10 Agent for ESX

When any of the Acronis Backup & Recovery 10 Agents for virtual machines is installed, you can back up:

- **Entire virtual machines** – 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 ones resulting from workloads consolidation. A separate archive will be created for each machine.
- **Volumes of a virtual machine** – 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 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. Backing up volumes within a virtual machine is similar to backing up physical machine's volumes. The virtual machine configuration will not be backed up.

Backing up an entire virtual machine or its volumes yields a standard disk backup (p. 320). Having Acronis Backup & Recovery 10 Agent for Windows or for Linux, you can mount its volumes, recover individual files from this backup, 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.

6.2.3. Items to backup

Depending on the type of the source selected previously one the following items to backup will be available:

Selecting disks and volumes (p. 171)

Selecting files and folders (p. 171)

6.2.3.1. 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.

2. [OPTIONAL] To create an exact copy of a disk or volume on a physical level, select the check box for **Back up sector-by-sector**. 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 any, 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 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.

6.2.3.2. Selecting files and folders

To select files and/or folders for backing up

1. Expand the local folders' tree items in order to get to the nested folders and files.
2. Select an item by checking the corresponding check box in the tree. Selecting a folder automatically selects all the nested folders and files. You can select a random set of files, folders, volumes, disks and even computers.

A file-based backup is not sufficient for the operating system recovery. In order to recover your operating system, you must image the system disk or volume.

Use the table in the right part of the window to browse and select 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**.

6.2.4. Access credentials for source

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

6.2.5. Exclusions

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

Predefined exclusions

To use the predefined sets of exclusions, check any of the following:

- **Exclude all hidden files and folders** – all files and folders marked with a "hidden" attribute will be excluded. Hidden files and folders contain user preferences, program and operating system-related data.
- **Exclude all system files and folders** - all files and folders marked with a "system" attribute will be excluded. System files have .sys extension in Windows.

Custom exclusions

If required, you can create your own exclusions for the specific files and/or file types.

To do so, proceed as follows

1. Select the **Exclude files matching the following criteria** check box

2. Click **Add…**
3. In the **Add file exclusion criterion** window specify an exclusion criterion, and then click **OK**. Both explicit rules (by entering an exact file name, or path) and common Windows masking rules (with using wildcard characters) are supported. See the table below for examples of exclusions.
4. The created exclusion appears in the bottom area of the window.
5. Click **OK** to save your settings.

*Note: *.bak, *.~, *.tmp file types are excluded by default.*

Exclusion examples

Criterion	Example	Description
By name	File1.log	Excludes all files named File1.log.
By path	C:\Finance\test.log	Excludes file named test.log and located on disk C:.
Mask (*)	*.log	Excludes all files with .log extension.
Mask (?)	my???.log	Excludes all .log files with names consisting of five symbols and starting with “my”.

Editing custom exclusions

Since you create a custom exclusion, you are able to edit and/or delete it. To do so, click the respective buttons.

6.2.6. Archive

Specify the location where the archive will be stored and the archive name.

Selecting the location

Enter the full path to the location in the **Path** field or select the desired location in the folders tree:

- To back up data to a personal location, expand the **Personal locations** group and click the location.
- To back up data to a local folder, locally attached tape device or removable media , expand the **Local folders** group and click the required folder/device.
- To back up data to 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 it.
- To back up data to an FTP or SFTP server, expand the corresponding group and reach the appropriate server.

An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

As appears from the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that 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 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 the archives.

You can switch between displaying archives by names and displaying the physical representation of the archives as TIB files using the **Show archives** and **Show TIB files** buttons.

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 [Type_Name], this means that the name contain variables. Such might be the case when you have selected virtual machines to back up. "Type" stands for the virtualization server type (ESX, Hyper-V or other.) "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 password, the program will ask for it in the pop-up window.

By selecting the existing archive, you are meddling in the area of another backup plan that uses the archive. This is not an issue if the other plan is discontinued, but in general you should follow the rule: "one backup plan - one archive". Doing the opposite will not prevent the program from functioning but is not practical or efficient except for some specific cases.

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

1. Backing up different sources to the same archive makes using the archive difficult from the usability standpoint. When it comes to recovery, each second counts, but 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 in some way 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. Especially you cannot expect the classic behaviour of the GFS and Tower of Hanoi backup schemes.

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

6.2.7. Access credentials for archive location

Specify credentials required for access to the location where the backup archive will be stored. The user name of these credentials that 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 location.

     Specify:

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

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

6.2.8. Backup schemes

Choose one of the available backup schemes:

- **Back up now** – to create a backup task for manual start and run the task right after its creation.
- **Back up later** – to create a backup task for manual start OR schedule the 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 per 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 lifetimes for the daily (referred as "sons"), weekly (referred as "fathers") and monthly (referred as "grandfathers") backups. The backups whose lifetimes have expired 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 number of backup levels (up to 16). In this scheme, the data can be backed up more than once per 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 in the archive 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.

6.2.8.1. Back up now scheme

With the **Back up now** scheme, the backup tasks will be executed immediately (right after you click the **OK** button).

In the **Backup type** field, select whether you want to create a full, incremental or differential backup. If you have not backed up the selected data yet, or the full archive seems too old to append incremental changes to it, choose full backup. Otherwise it is recommended that you create an incremental or differential backup.

6.2.8.2. Back up later scheme

With the Back up later scheme, the backup plan will be executed 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 you have not backed up the selected data yet, or the full archive seems too old to append incremental changes to it, choose full backup. Otherwise it is recommended that you create an incremental or differential backup.
Date and time	Specify when to start the backup plan.
The task will be started manually	Select this check box, if you do not need to put the backup plan on a schedule and wish to start it manually afterwards (by clicking Run on the Tasks view).

6.2.8.3. Simple scheme

With the simple backup scheme you just schedule when and how often to back up data and set the retention rule. At first time the full backup will be created. Next backups will be incremental.

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

Backup	Set up the backup schedule - when and how often to backup the data. To learn more about setting up schedule, see the Scheduling (p. 140) section.

Retention rule	With the simple scheme, only one retention rule is available. Set the backups' lifetime.

6.2.8.4. Grandfather-Father-Son scheme

At a glance

- Daily incremental, weekly differential, and monthly full backups
- Custom day for weekly and monthly backups
- Custom retention periods for backups of each type

Description

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

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

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

- Monthly ("Grandfather") backups are full;
- Weekly ("Father") backups are differential;
- Daily ("Son") backups are incremental.

Parameters

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

Start backup at:	Specifies when to start a backup. The default value is 12:00 PM.

Back up on:	Specifies the days on which to perform a backup. The default value is Workdays.
Weekly/Monthly:	Specifies which of the days selected in the Back up on field you want to reserve for weekly and monthly backups. The default value is Friday. A monthly backup will be performed every fourth such day.
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: 1 week (recommended minimum) Weekly: 1 month Monthly: indefinite The retention period for weekly backups must exceed that for daily backups; monthly backups' retention period must be greater than 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

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

- Back up files every day, including weekends;
- Be able to recover files as of any date over the past seven days;
- Have access to weekly backups of the past month;
- Keep monthly backups indefinitely.

Backup scheme parameters can then be set up as follows.

- Start backup at: **11:00 PM**
- Back up on: **All days**
- Weekly/monthly: **Saturday** (for example)
- Keep backups:

- Daily: **1 week**
- Weekly: **1 month**
- Monthly: **indefinite**

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.

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

Let the plan allow for

- Performing backups at the end of each working day;
- A possibility to recover an accidentally deleted or inadvertently modified file if this has been relatively quickly discovered;
- Having access to a weekly backup for 10 days after it was created;
- Keeping monthly backups for half a year.

Backup scheme parameters can then be set up as follows.

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

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

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, 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, over the six months, and keep the monthly backups for five years.

The following GFS scheme would suit such purposes:

- Start backup at: **11:30 PM**
- Back up on: **Tuesday, Thursday, Friday**

- Weekly/monthly: **Friday**
- Keep backups:
  - Daily: **16 weeks**
  - Weekly: **6 months**
  - Monthly: **5 years**

Here, daily incremental backups will be created on Tuesdays and Thursdays, with weekly and monthly backups performed on Fridays. Note that, in order to choose Friday in the Weekly/monthly field, you should first enlist 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, have a five-year history of all documents, etc.

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: **indefinite**

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.

6.2.8.5. 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, weekly, or monthly schedule. Setting up schedule parameters allows create simple schedules (example of a simple daily schedule: a backup task will be run every 1 day on 10 AM) as well as more complex schedules (example of a complex daily schedule: a task will be run on every 3 days, starting form 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 towards recent 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 7 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 differential backup of level 3 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. This enables us to go as far back as 11 days, which is the best-case scenario.

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

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

This gives us only a four days' 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 falling 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 equals four days.

6.2.8.6. 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.
Retention rules	Specifies what retention rules will be applied to the archive. For example, the cleanup procedure can be set up to delete all backups older than six months.
Apply the rules (only if the retention rules are set)	Specifies when to apply the retention rules (p. 42). 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 (only if **On schedule** is selected)	Specifies a schedule for archive cleanup. 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

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).

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. 42).

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**: **Machine 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 the machine is 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. 42).

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.

6.2.9. 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 system resources.
3. **Validation schedule** (appears only if you have selected the on schedule in the step 1) - set the schedule of validation. For more information see the Scheduling section.

6.2.10. Backup options

Customize your backup plan by configuring the backup options, such as pre/post backup commands, maximum network bandwidth allocated for the backup stream or the backup archive compression level.

Default options and presets

Each of Acronis agents has its own default backup options with preset values. To access the default options, select **Options - Agent options - Default backup options** from the menu. Once you have changed a preset value, it becomes default for this agent and will be used as default in further backup operations. Nevertheless, while backing up you can override the default settings with the custom settings that will be specific for this plan only.

The plan-specific settings and the option values set by default

Before any setting is changed against default value the **Backup options** section on the **Create backup plan** page contains the only line with the settings status Default.

Clicking the **Change…** link opens **Backup Options** window, where you can change settings of the current backup operation.

After one or more settings are changed against default values, a new line that contains the text description of modified backup options appears in the section. The settings status is changed from the **Default** to **Custom**.

Clicking the **Reset to defaults** resets all backup settings to the default values.

Availability of the backup options

The set of available backup options differs depending on:

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

The following table shows the summary of the backup options availability.

	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. 100) (password + encryption)	+	+	+	+	+	+
Source files exclusion (p. 100)	+	+	+	+	+	+
Pre-post backup commands (p. 101)	+	+	+	+	PE only	PE only
Pre-post data capture commands (p. 103)	+	+	+	+	-	-
Multi-volume snapshot (p. 106)	+	+	-	-	-	-
File-level backup snapshot (p. 105)	-	+	-	+	-	-

Use VSS (p. 106)	+	+	-	-	-	-
Compression level (p. 106)	+	+	+	+	+	+
Backup performance:						
Backup priority (p. 107)	+	+	+	+	-	-
HDD writing speed (p. 107)	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD
Network connection speed (p. 107)	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share
Fast incremental/differential backup (p. 110)	+	-	+	-	-	-
Backup splitting (p. 110)	+	+	+	+	+	+
File-level security (p. 111):						
Preserve files' security settings in archives	-	+	-	-	-	-
In archives, store encrypted files in decrypted state	-	+	-	-	-	-
Media components (p. 112)	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	-	-
Error handling (p. 112):						
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Ignore bad sectors	+	+	+	+	+	+
Dual destination (p. 113)	+	+	+	+	-	-
Task start conditions (p. 114)	+	+	+	+	-	-
Task failure handling (p. 115):						
Stop executing the backup plan	+	+	+	+	-	-
Restart the failed task	+	+	+	+	-	-
Additional settings (p. 115):						
Overwrite data on a tape without prompting user for confirmation	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape
Dismount media after backup is finished						

Ask for first media while creating backup archives on removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media
Validate backup after creation	-	-	-	-	+	+
Reset archive bit	-	+	-	-	-	+
Reboot after the backup	-	-	-	-	+	+
Create full backups as synthetic backups	+	+	+	+	+	+
Notifications:						
E-mail (p. 108)	+	+	+	+	-	-
Win Pop-up (p. 108)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 109)	+	+	-	-	-	-
SNMP (p. 109)	+	+	+	+	-	-

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. 317) or using Acronis Startup Recovery Manager (p. 55). Then, create a recovery task.

The Acronis Universal Restore (p. 56) 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. 58) technology, Acronis Backup & Recovery 10 will boot the machine into the operating system found in the backup as if the system was 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 basic disk. To learn more about recovering dynamic volumes, please turn to Microsoft LDM (Dynamic volumes) (p. 45) section.

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 more about Acronis Disk Director LV, see the Disk management (p. 231) 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. 191)

[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 check box for **Advanced view**.

What to recover

Archive (p. 191)

Select the archive to recover data from.

Backup (p. 192)

Select the backup and content to be recovered.

Content type (p. 193)

Applies to: disk recovery

Choose the type of data you need to recover from the selected disk backup.

Content

If the backup content has been not selected yet or if you need to change your selection, you can do it right here.

Access credentials (p. 193)

[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 check box for **Advanced view**.

Where to recover

This section appears after the required backup is selected and the type of data to recover is defined. The parameters you specify here depend on the type of data being recovered.

Disks (p. 194)

Volumes (p. 196)

Files (p. 199)

You may have to specify credentials for the destination. Skip this step when operating on a machine booted with bootable media.

Access credentials (p. 201)

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

When to recover

Recover (p. 201)

Select when to start the recovery. The task can start immediately after its creation, be scheduled for the specified date and time in the future or simply saved for manual execution.

[OPTIONAL] Acronis Universal Restore

Applies to: Windows OS and system disk or volume recovery

Universal Restore (p. 202)

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 check box for **Advanced view**.

Recovery options

Settings (p. 203)

[OPTIONAL] Customize settings of the recovery task. If you do nothing in this section, the default values 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 defaults** resets all 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 a name of an Active Directory user account, be sure to specify also the domain name (DOMAIN\Username.)
  - **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 Users' privileges on a managed machine (p. 34) section.

6.3.2. Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field In the folders tree, or select the desired location/folder in the folders tree:
   - If the archive is stored in a personal location, expand the **Personal locations** group and click the location.
   - If the archive is stored in a local folder, locally attached tape device or removable media, expand the **Local folders** group and click the required folder/drive. Make sure the media is inserted.
   - If the archive is stored in the 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 it.
   - If the archive is stored on the **FTP** or **SFTP** server: expand the corresponding group and reach the appropriate folder on the server.

     An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

     As appears from the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that user name and password can be intercepted by an eavesdropper using a packet sniffer.
2. In the table to the right of the tree, select the archive. The table displays names of the archives contained in each location/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 the archives.

   You can switch between displaying archives by names and displaying the physical representation of the archives as TIB files using the **Show archives** and **Show TIB files** buttons.

3. Click **OK**.

When using bootable media

Keep in mind that Acronis rescue utilities identify volumes in other way than the operating systems do.

The media with bootable agent for Windows might show disk letters different from ones you see in Windows. For example, the D: volume under the rescue utility might correspond to the E: volume in Windows.

The media with bootable agent for Linux will show local volumes as unmounted (scsi1p1, scsi1p2..)

6.3.3. Backup selection

The Backup selection window's representation depends on the type of data stored in the archive you have selected.

6.3.3.1. Selecting disks/volumes

To select a backup and disks/volumes to recover:

1. Select one of successive incremental backups by its creation date and time. Thus, you can return the disk data to a certain moment.

   Specify items to recover. By default, all items of the selected backup will be selected. If you do not want to recover the certain items, just uncheck them.
   To obtain information on a disk/volume, right-click it and then click **Information**.

2. [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. 58) for details.

3. Click **OK**.

Selecting 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 MBR of a one disk to another Acronis Backup & Recovery 10 recovers Track 0 not affecting 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 MBR and Track 0 for Windows systems, unless the MBR is damaged.

6.3.3.2. Selecting files

To select a backup and files to recover:

1. Select one of successive incremental backups by its creation date/time. Thus, you can return the files/folders to a specific moment.
2. Specify 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 beside the **Name** column's header automatically selects all items in the table. Clearing this check box automatically deselects all items.
3. Click **OK**.

6.3.4. Data type

Choose what to recover from the selected disk backup:

- **Disks** - to recover disks
- **Volumes** - to recover volumes
- **Files** - to recover the specific files and folders

6.3.5. Access credentials for location

Specify 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 location.

     Specify:

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

As appears from the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that 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

The available disk destinations depend on the agents residing on the machine.

Acronis Backup & Recovery 10 Agent for Windows or for Linux

When the Acronis Backup & Recovery 10 Agent is installed, you can set the following parameters:

Disk #:

Disk # (MODEL) (p. 197)

Select the destination disk for each of the source disks.

Disk signature (p. 195)

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.

Acronis Backup & Recovery 10 Agent for Hyper-V

When the Acronis Backup & Recovery 10 Agent for Hyper-V is installed, you can choose between the three destinations.

Recover to:

Physical machine

This will recover the selected disks to the physical disks of the host server you are connected to. On selecting this, you proceed to the regular disk mapping procedure described above.

Existing virtual machine

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular disk mapping procedure described above.

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.

New virtual machine

On selecting this, you specify the virtualization server and the name of the virtual machine to be created on the server. The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **VM settings** (p. 198) section. Check the settings and make changes if necessary. Then you proceed to the regular disk mapping procedure described above.

Acronis Backup & Recovery 10 Agent for ESX

Acronis Backup & Recovery 10 Agent for ESX is installed, you don't have the ability to recover disks to the physical machine unless the Acronis Backup & Recovery 10 Agent for Linux is installed on the host. Choose between the two destinations.

Recover to:

Existing virtual machine

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular disk mapping procedure described above.

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.

New virtual machine

On selecting this, you specify the virtualization server and the name of the virtual machine to be created on the server. The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **VM settings** (p. 198) section. Check the settings and make changes if necessary. Then you proceed to the regular disk mapping procedure described above.

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 for non-backed-up data that you might need.

Disk signature

When MBR is selected along with disk backup, choose what the program should do with disk signature:

- **Create new**

  The program will generate a new disk signature for the recovered drive.

  This may be needed when you use an image backup not for disaster recovery but for cloning your Windows Vista hard drive to another one. Trying to boot Windows after cloning with both drives connected will result in a problem. During Windows booting its loader checks the disk signatures of all of the connected drives, and if it finds two identical disk signatures, the loader changes the signature of the second disk, which would be the clone disk. Once this happens, the clone disk would not be able to boot up independently of the original disk, because the *MountedDevices* fields in the clone's registry reference the disk signature of the original disk, which will not be available if the original disk is disconnected.

- **Recover from backup**

  The program will replace the existing disk's signature with one from the backup.

  Recovering 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
- Allows 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**

  Leaves the existing destination hard disk drive's signature untouched.

6.3.6.2. Volumes

The available volume destinations depend on the agents residing on the machine.

Acronis Backup & Recovery 10 Agent (for Windows and for Linux)

When the Acronis Backup & Recovery 10 Agent is installed, you can set the following parameters:

[Volume name] [Letter]:

Disk # /Volume (p. 197)

Sequentially map each of the source volumes to a volume or an unallocated space of the destination disk.

Size (p. 197):

[Optional] Change the recovered volume size, location and other properties.

Acronis Backup & Recovery 10 Agent for Hyper-V

When the Acronis Backup & Recovery 10 Agent for Hyper-V is installed, you can choose between the three destinations.

Recover to:

Physical machine

This will recover the selected volumes to the physical disks of the host server you are connected to. On selecting this, you proceed to the regular volume mapping procedure described above.

Existing virtual machine

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular volume mapping procedure described above.

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.

New virtual machine

On selecting this, you specify the virtualization server and the name of the virtual machine to be created on the server. The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **VM settings** (p. 198) section. Check the settings and make changes if necessary. Then you proceed to the regular volume mapping procedure described above.

Acronis Backup & Recovery 10 Agent for ESX

When the Acronis Backup & Recovery 10 Agent for ESX is installed, you don't have the ability to recover volumes to the physical machine unless the Acronis Backup & Recovery 10 Agent for Linux is installed on the host. Choose between the two destinations.

Recover to:

Existing virtual machine

On selecting this, you specify the virtualization server and the target virtual machine. Then you proceed to the regular volume mapping procedure described above.

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.

New virtual machine

On selecting this, you specify the virtualization server and the name of the virtual machine to be created on the server. The new virtual machine will be configured automatically, the source machine configuration being copied where possible. The configuration is displayed in the **VM settings** (p. 198) section. Check the settings and make changes if necessary. Then you proceed to the regular volume mapping procedure described above.

Volume destination

To specify a destination volume:

1. Select a volume, or unallocated space where you want the selected volume to recover 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 for non-backed-up data that you might need.

Changing the properties of the recovered volume

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 volumes being recovered. In this case, you will have to recover the volume to be reduced first.

Properties

Type

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 from it. 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. Logical volume cannot be set as active. If you recover a system volume to another hard disk with its own volumes and operating system, most likely you will 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 to recover a partition 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 partitions up to 4GB, so you will not be able to recover a 4GB FAT16 partition to a partition that exceeds that limit without changing the file system. It would make sense here to change the file system from FAT16 to FAT32.

The 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 partition and change its file system. These can be normally recovered on a FAT16 partition 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 partition, hiding it from OS.You should not assign letters to partitions inaccessible to Windows, such as to those other than FAT and NTFS.

6.3.6.3. Virtual machine configuration

The following virtual machine settings can be configured.

Storage

Initial setting: the default storage of the virtualization server.

This is the place on the server where the new virtual machine will be created. Whether you can change the storage or not, depends on the virtualization product brand and settings. VMware ESX may have multiple storages. On a Microsoft Hyper-V server, only the default storage is available.

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 commonly equals 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.

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 more than one processors 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.

6.3.6.4. 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 recreated 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 full path, unless you clear the **Recover without full path** check box.
2. Click **OK**.

Overwriting (p. 201)

Exclusions (p. 199)

Exclusions

Set up exclusions for the specific types of files you do not wish to be overwritten during the recovery.

Predefined exclusions

To use the predefined sets of exclusions, check any of the following:

- **Exclude all hidden files and folders** – all files and folders marked with a "hidden" attribute will be excluded. Hidden files and folders contain user preferences, program and operating system-related data.
- **Exclude all system files and folders** - all files and folders marked with a "system" attribute will be excluded. System files have .sys extension in Windows.

Custom exclusions

If required, you can create your own exclusions for the specific files and/or file types.

To do so, proceed as follows

1. Select the **Exclude files matching the following criteria** check box
2. Click **Add…**
3. In the **Add file exclusion criterion** window specify an exclusion criterion, and then click **OK**. Both explicit rules (by entering an exact file name, or path) and common Windows masking rules (with using wildcard characters) are supported. See the table below for examples of exclusions.
4. The created exclusion appears in the bottom area of the window.
5. Click **OK** to save your settings.

*Note: *.bak, *.~, *.tmp file types are excluded by default.*

Exclusion examples

Criterion	Example	Description
By name	File1.log	Excludes all files named File1.log.
By path	C:\Finance\test.log	Excludes file named test.log and located on disk C:.
Mask (*)	*.log	Excludes all files with .log extension.
Mask (?)	my???.log	Excludes all .log files with names consisting of five symbols and starting with “my”.

Editing custom exclusions

Since you create a custom exclusion, you are able to edit and/or delete it. To do so, click the respective buttons.

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 the 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 overwriting files, you still have an option to prevent overwriting (p. 199):

- hidden files and folders
- system files and folders
- any files you specify by names or using wildcards
- any folder you specify by path.

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 a name of an Active Directory user account, be sure to specify also the domain name (DOMAIN\Username.)
     - **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 put the task on a schedule and wish to start it manually afterwards (by clicking **Run** in the **Backup plans and tasks** view), select the **Task will be started manually (do no schedule the task)** check box.

6.3.9. Acronis Universal Restore

Use the Acronis Universal Restore when you need to recover and boot up Windows on dissimilar hardware or a virtual machine. Acronis Universal Restore handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

With Acronis Backup & Recovery 10 Virtual Edition, you might need the Acronis Universal Restore when recovering a virtualization server itself. When recovering a system to a virtual machine, the Universal Restore technology is applied in the background, because the program knows what drivers are needed for the supported virtual machines.

Acronis Universal Restore is not available when:

- a machine is booted with Acronis Startup Recovery Manager (using F11)
- the image being recovered is located in Acronis Secure Zone
- using Acronis Active Restore

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 for the operating system start. 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 extension. 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 types of the devices 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 the Acronis Universal Restore is used.

Acronis Universal Restore settings

- **Automatic driver search**

  Specify where the program should search for the Hardware Abstraction Layer (HAL), HDD controller driver and network adapter driver(s):

  - If the drivers are on the vendor's 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, Acronis 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 the available, and install them into the recovered system. Acronis Universal Restore also searches for the network adapter driver and transmits the path to the found driver to the operating system. If the hardware has multiple network interface cards, Acronis Universal Restore will try to configure all the cards drivers. In case Acronis 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 check box for **Advanced view**.

  If the target hardware has a specific mass storage controller such as RAID (specially NVIDIA RAID) or a fibre channel adapter, specify the appropriate drivers in the **Drivers** field.

  The drivers defined here will have the priority. They will be installed, with appropriate warnings, even if the program finds the better driver.

  Use this option only if the automatic drivers search does not help to boot the system.

When recovering the system to a virtual machine that uses SCSI hard drive controller, be sure to specify SCSI drivers for virtual environment. Use drivers bundled with your virtual machine software or download the latest drivers versions from the software manufacturer website.

6.3.10. Recovery options

Customize the recovery operation by configuring the recovery options, such as pre/post recovery commands, recovery priority, error handling or notification options.

Default options and presets

Each of Acronis agents has its own default recovery options with preset values. To access the default options, select **Options - Agent options - Default recovery options** from the menu. Once you have changed a preset value, it becomes default for this agent and will be used as default in further recovery operations. Nevertheless, while creating a recovery task, you can override the default settings with the custom settings that will be specific for this task only.

The task-specific settings and the default options

Before any setting is changed against default value the **Recovery options** section contains the only line with the settings status **Default**.

Clicking the **Change…** link opens the **Recovery options** window, where you can change settings of the current recovery operation.

After one or more settings are changed against default values, a new line that contains the text description of modified backup options appears in the section. The settings status is changed from **Default** to **Custom**.

Clicking this **Reset to defaults** resets all settings to the default values.

Availability of the recovery options

The set of available recovery options differs depending on:

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

The following table shows the summary of the recovery options availability.

	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. 118)	+	+	+	+	PE only	PE only
Recovery priority (p. 119)	+	+	+	+	-	-
File-level security (p. 120):						
Recover files with their security settings	-	+	-	+	-	+
Error handling (p. 122):						
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Additional settings (p. 123):						
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. 120)	+	+	+	+	-	-
Win Pop-up (p. 121)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 121)	+	+	-	-	-	-
SNMP (p. 122)	+	+	+	+	-	-

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 the 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 set bootable**

  Solution: Recover Windows to a basic, simple or mirrored volume.

- **A system volume was recovered to a disk that does not have MBR**

  When you configure recovery of a system volume to a disk that does not have 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**

  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. The 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 the 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 with the same absolute numbers as on the source disk.

  Solution: Reactivate the boot loader. You might need also 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 recovered not to their original location.

  Solution:

  Fixing this for Windows loaders comes to modification of the boot.ini or the boot\bcd files. 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 also recommended that you change /etc/fstab so that the SWAP volume could 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 LVM is not used and reactivate the boot loader.

6.3.11.1. How to reactivate GRUB or LILO and change its configuration

Generally, you should refer to the boot loader manual pages for the appropriate procedure. In case the system disk (volume) is recovered to identical hardware, the following steps would usually help.

GRUB

1. Boot the machine from the Linux rescue CD or Linux installation CD.

   The installed system and the system on the CD must have the same kernel version. The Linux distribution does not have to be the same. Opt for boot without RAM disk (without the installer):

   ```
   linux noinitrd root=/dev/ROOTDEV
   ```

   Where

   "linux" is the kernel name (the installation CD may contain multiple kernels with different names, LILO and GRUB enable you to select the kernel)

"ROOTDEV" is the device corresponding to the root partition. This is usually hda1 on IDE device and sda1 on SCSI device.

The kernel loads and the operating system starts from the specified root partition. Log in as root; or log in as a user and switch user: $ sudo su.

2. Edit the GRUB configuration file if the number of partition where the operating system resides has changed. Otherwise skip this step.

   a. Open the GRUB configuration file (usually /boot/grub/grub.conf or /etc/grub.conf):

```
# vim /boot/grub/grub.conf
```

   b. Find the record "root(hdX,Y)" in the section corresponding to the current kernel.

      In this record,

      X – number of the disk

      Y – number of partition on this disk

      Change X and Y according to the new location of the root.

   c. Save the configuration file. In VIM, this is done by pressing subsequently <esc> : w q <enter>

3. Execute the command for activating GRUB:

```
# grub
```

4. Eject the CD and reboot.

LILO

1. Perform step 1 described above.
2. Edit the LILO configuration file if the number of partition where the operating system resides has changed. Otherwise skip this step.

   a. Open the LILO configuration file (/etc/lilo.conf):

```
# vim /etc/lilo.conf
```

   b. Edit the record "root=/dev/XdaY" in the section corresponding to the current kernel. In this record,

      X – number of the disk

      Y – number of partition on this disk

      Change X and Y according to the new location of the root.

   c. Save the configuration file.

3. Execute the command for activating LILO:

```
# lilo
```

4. Eject the CD and reboot.

6.3.11.2. About Windows loaders

Windows NT/2000/XP/2003

A part of 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 loader resides in the partition boot sector, the rest is in the files bootmgr, boot\bcd. At starting Windows, boot\bcd is mounted to the registry key HKLM \BCD00000000.

6.4. Validating locations, 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 volume backup calculates a checksum for every data block saved in the backup. Both procedures are resource-intensive.

While the 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 the recovery success. At least ensure that the backup can be successfully validated using the bootable media.

Validation of an archive will validate all the archive's backups. A location validation will validate all archives in the location.

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 the permission to access.

Validation of an archive or of the latest backup in the archive can be scheduled as a part of backup plan. For more information see Backup plan.

You can access the Validation page from the Locations view. Right-click the object to validate (archive, backup or location) 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 the name for the task.

When exporting a backup from an archive, you have an option to validate the backup immediately after the export. For more information see Exporting a backup.

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

[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 check box for Advanced view.

What to validate

Validate

Choose an object to validate:

- Archive - in that case, you need to specify the archive.
- Backup - specify the archive first, and then select the desired backup in this archive.
- Location - select a location, which archives to validate.

Access Credentials

[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

Specify when and how often to perform validation.

After you configure all the required settings, click **OK** to create the validation task.

After the task is created, it will be available for examination and editing on the **Backup plans and tasks** view.

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 a name of an Active Directory user account, be sure to specify also the domain name (DOMAIN\Username.)
- **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 Users' privileges on a managed machine (p. 34) section.

6.4.2. Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field In the folders tree, or select the desired location/folder in the folders tree:
   - If the archive is stored in a personal location, expand the **Personal locations** group and click the location.
   - If the archive is stored in a local folder, locally attached tape device or removable media, expand the **Local folders** group and click the required folder/drive. Make sure the media is inserted.
   - If the archive is stored in the 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 it.
   - If the archive is stored on the **FTP** or **SFTP** server: expand the corresponding group and reach the appropriate folder on the server.

     An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

     As appears from the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that user name and password can be intercepted by an eavesdropper using a packet sniffer.
2. In the table to the right of the tree, select the archive. The table displays names of the archives contained in each location/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 the archives.

   You can switch between displaying archives by names and displaying the physical representation of the archives as TIB files using the **Show archives** and **Show TIB files** buttons.

Click **OK**.

6.4.3. Backup selection

To specify a backup to validate

1. In the upper pane, select a backup by its creation date/time.

The lower table displays the selected backup's 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 personal location, expand the **Personal locations** group and click the location.
- To select a local folder, locally attached tape device or removable media, expand the **Local folders** group and click the required folder/device.
- To select 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 it.
- To select **FTP** or **SFTP** server, expand the corresponding group and reach the appropriate server's folder.

An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

As appears from the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that 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 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 the archives.

You can switch between displaying archives by names and displaying the physical representation of the archives as TIB files using the **Show archives** and **Show TIB files** buttons.

6.4.5. Access credentials for location

Specify 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 location.

     Specify:

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

2. Click **OK**.

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

6.4.6. When to validate

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 you prefer to be immediately informed whether the data is not corrupted and can be successfully recovered, think of starting the 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
  - **Do not specify the date and time, just create task to run manually** - select this check box, if you wish to start the task manually later.
- **On schedule** - to put the task on schedule. To learn more about how to configure the scheduling parameters, please see the Scheduling (p. 140) section.

6.5. Mounting images and managing mounted images

Mounting an image (to be more precise, mounting a volume contained in a disk or volume backup) lets you access the backup as though it was a physical disk. The mount operation is available when the console is connected to a managed machine running either Windows or Linux.

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 several files and folders from a volume backup, you do not have to perform the recovery procedure.

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.

Multiple volumes contained in the same backup can be mounted within a single mount operation.

Limitation

- Mounting of volume backups stored on Acronis Backup & Recovery 10 Storage Node is not possible.

Alternative way to access mounting operation

You can also access the Mount page from the **Locations** view. Right-click the backup and select Mount from the context menu. The Mount Image page will be opened with the pre-selected backup as a source. All you need to do is to set up the mount settings.

6.5.1. Mounting images

To mount a volume, perform the following steps.

Source

Archive (p. 213)

Specify path to the archive location and select the archive.

Backup (p. 214)

Select the disk backup.

Access credentials (p. 214)

[OPTIONAL] Provide credentials for the archive location. To access this option, select the check box for **Advanced view**.

Mount settings

Volumes (p. 215)

Select volumes to mount and configure the mount settings for every volume: assign 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 volumes.

6.5.1.1. Archive selection

Selecting the archive

1. Enter the full path to the location in the **Path** field In the folders tree, or select the desired location/folder in the folders tree:
   - If the archive is stored in a personal location, expand the **Personal locations** group and click the location.
   - If the archive is stored in a local folder, locally attached tape device or removable media, expand the **Local folders** group and click the required folder/drive. Make sure the media is inserted.
   - If the archive is stored in the 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 it.
   - If the archive is stored on the **FTP** or **SFTP** server: expand the corresponding group and reach the appropriate folder on the server.

An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

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

2. In the table to the right of the tree, select the archive. The table displays names of the archives contained in each location/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 the archives.

   You can switch between displaying archives by names and displaying the physical representation of the archives as TIB files using the **Show archives** and **Show TIB files** buttons.

Click **OK**.

6.5.1.2. Backup selection

To select a backup:

1. Select one of backups by its creation date/time.
2. To assist you with choosing the right backup, the bottom table displays volumes contained in the selected backup.

   To obtain information on a volume, right-click it and then click **Information**.
3. Click **OK**.

6.5.1.3. Access credentials

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 location.

     Specify:

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

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

6.5.1.4. Mount settings - Volume Selection

Select the volumes to mount and configure the mounting parameters for each of the selected volume as follows

1. Select the check box for each volume you need to mount.
2. Click on the selected volume to set up its mounting parameters:
   - **Access mode** - choose the mode you want the image to be mounted in:
     - **Read only** - enables exploring and opening files within the volumes in without committing any changes, adding, modifying or deleting files.
     - **Read/write** - with this mode, the program assumes that the backup content will be modified, and creates an incremental backup file to capture the changes.
   - **Assign letter** - Acronis Backup & Recovery 10 will assign an unused letter to a recovered volume. If required, select another letter to assign from the drop-down list.
3. If several volumes are checked for mounting, click on every volume to set up its mounting parameters, described in step 2.
4. Click **OK**.

6.5.2. Managing mounted images

Once a volume is mounted, it appears on the Manage Mounted Images page (Actions and tools -> Manage mounted images.) From this point on, you can operate with files or folders of the mounted volume as if they were located on a physical drive.

6.5.2.1. Exploring mounted images

Exploring mounted volumes lets you view and modify (if mounted in the read/write mode) the volume's content.

To explore a volume

1. In the **Manage Mounted Images** view, select the mounted volume from the list.
2. On the toolbar, click the **Explore** button. The Windows Explorer window opens, allowing user to examine the mounted volume contents.

6.5.2.2. Unmounting images

Keeping up 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 stay mounted until the operating system restarts.

To unmount a volume

1. In the **Manage Mounted Images** view, select the mounted volume from the list.
2. On the toolbar, click the **Unmount** button.

   To unmount all the mounted volumes, click **Unmount All**.

6.6. Managing Acronis Secure Zone

The Acronis Secure Zone is a hidden and secure partition that enables keeping backup archives on a managed machine disk space. Windows applications, except for dedicated disk management tools, cannot access the zone, and so the archives stored in the zone are protected from software malfunction and user errors.

To learn more about advantages and limitations of the Acronis Secure Zone, see the Proprietary Acronis technologies -> Acronis Secure Zone section.

Managing the Acronis Secure Zone

In terms of backup archives management, Acronis Secure Zone is considered as personal location. Once created on a managed machine, the zone is always present in the list of Personal locations shortcuts. 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 location 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 a backup content
- export a backup from the zone
- 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 managing locations, see the Locations section.

6.6.1. Allocating space for Acronis Secure Zone

The Acronis Secure Zone can be located on any internal 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 volumes' free space. When creating the zone, you select the desired zone size and (optionally) the volumes to take free space from. Unallocated space is always selected.

Moving or resizing of locked volumes, such as the volume containing the currently active operating system, requires a reboot.

How the settings you do 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 volumes layout, 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.
- There should be free space on a volume however so that the operating system and applications can operate, for example, create 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, the program will continue decreasing the volumes proportionally.

As appears 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 unstable and even fail to start.

6.6.2. Creating Acronis Secure Zone

You can create Acronis Secure Zone while the operating system is running or using bootable media. Moving or resizing of locked volumes, such as the volume containing the currently active operating system, requires a reboot.

To create Acronis Secure Zone, perform the following steps.

Settings

Disk

Choose a hard disk (if several) on which to create the zone. Unallocated space is selected by default. The program displays the total space available for the Acronis Secure Zone. Compare this value with the estimated zone size. Do not select any volume and go to the next step if the space is enough.

If the space is not enough, 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 next step.

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 40MB, 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 mind the following:

- 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 unstable and even fail to start. Do not set the maximum zone size if the boot or the system volume is selected.

Password

[OPTIONAL] You have an option to protect the Acronis Secure Zone with a password. The program will ask for the password at any operation relating to the zone, such as data backup and recovery, mounting images or validating archives on the zone, using the Acronis Startup Recovery Manager with the F11 key, resizing and deleting the zone.

Acronis Startup Recovery Manager

Acronis Startup Recovery Manager is a bootable rescue utility (in fact, a modification of the Acronis Backup & Recovery 10 Bootable Agent) located in the Acronis Secure Zone and configured to start at boot time on pressing F11. It eliminates the need for a separate media or network connection to start the rescue utility.

To enable Acronis Startup Recovery Manager, select **Activate**.

Acronis Startup Recovery Manager activation overwrites the Master Boot Record (MBR) with its own boot code. You will have to reactivate third-party boot loaders, if installed. Under Linux, consider installing the LILO or GRUB loader to a Linux root (or boot) partition boot record instead of MBR before the activation. Otherwise, reconfigure these boot loaders manually after the activation.

To disable Acronis Startup Recovery Manager, select **Do not activate**.

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 Service (RIS).

See Bootable rescue media for details.

You can activate Acronis Startup Recovery Manager later from the Manage Acronis Secure Zone page.

After you configure the required settings, click **OK** to start creating the zone.

6.6.3. 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 your operating system from booting.

4. Click **OK**.

6.6.4. 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 40MB, depending on the geometry of the hard disk;
   - typing an exact value in the Acronis Secure Zone Size field.

   When reducing the zone, any unallocated space, if the hard disk has it, will be allocated to the selected partitions along with the space freed from the zone. Thus, no unallocated space will remain on the disk.
4. Click **OK**.

6.6.5. Password for Acronis Secure Zone

Setting up a password protects the Acronis Secure Zone from an unauthorized access. The program will ask for the password at any operation relating to zone, such as data backup and recovery, exploring archives, mounting images or validating archives on the zone, using the Acronis Startup Recovery Manager with the F11 key, 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**.

6.6.6. Acronis Startup Recovery Manager

Acronis Startup Recovery Manager is a bootable rescue utility (in fact, a modification of the Bootable agent) located in the Acronis Secure Zone and configured to start at boot time on pressing F11. It eliminates the need for a separate media or network connection to start the rescue utility.

Activate

Enables the boot time prompt "Press F11 for Acronis Startup Recovery Manager...". If the system fails to boot, you will be able to start bootable rescue utility located in the Acronis Secure Zone by pressing F11.

Acronis Startup Recovery Manager activation overwrites the Master Boot Record (MBR) with its own boot code. You will have to reactivate third-party boot loaders, if installed. Under Linux, consider installing the LILO or GRUB loader to a Linux root (or boot) partition boot record instead of MBR before the activation. Otherwise, reconfigure these boot loaders manually after the activation.

Do not activate

Disables boot time prompt "Press F11 for Acronis Startup Recovery Manager…". 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 Service (RIS.).

See Bootable rescue media for details.

6.6.7. Deleting Acronis Secure Zone

Acronis Secure Zone deletion will automatically disable Acronis Startup Recovery Manager if it is activated and destroy all backups stored in the zone. Thus, to keep your backups safe and sound, think of exporting them before you delete the zone.

To delete the zone without uninstalling the program, proceed as follows:

1. In the Acronis Secure Zone Actions bar, 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 volumes, the freed space becomes unallocated.

After you click **OK**, Acronis Backup & Recovery 10 will start deleting the zone.

There is an option to keep Acronis Secure Zone along with its contents (which will enable data recovery on booting from bootable media) or remove Acronis Secure Zone if you remove Acronis Backup & Recovery 10 Agent from the system.

6.7. 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 the Linux-based environment or Windows Preinstallation Environment (WinPE), without help of an operating system. Bootable media is most often used to:

- recover an operating system that cannot start
- deploy an operating system on bare metal
- create basic or dynamic volumes on bare metal
- back up logical volumes while the system volume is corrupted
- back up sector-by-sector a disk with 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 or Microsoft Remote Installation Service (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 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 the necessary devices drivers.

6.7.1. How to create bootable media

Linux-based bootable media

To create Linux-based bootable media, prepare a blank disk or install Acronis PXE Server or set up Microsoft WDS. Then start the Bootable Media Builder by selecting **Tools > Create Bootable Rescue Media**. The wizard will guide you through necessary operations. Please refer to Bootable Media Builder (p. 222) 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)

If you already have media with PE1.x distribution, unpack the media ISO to a local folder and start the Acronis WinPE ISO Builder by selecting it from the start menu -> Acronis. The wizard will guide you through necessary operations. Please refer to Adding the Acronis Plug-in to WinPE 1.x (p. 225) for details.

To be able to create or modify PE 2.x images, install Acronis WinPE ISO 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 (p. 226) 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

1. [optional] Burn the WAIK to DVD or copy to a flash drive.
2. Install the Microsoft .NET Framework v.2.0 from this kit (NETFXx86 or NETFXx64, depending on your hardware.)
3. Install Microsoft Core XML (MSXML) 5.0 or 6.0 Parser from this kit.
4. Install Windows AIK from this kit.
5. Install Acronis WinPE ISO Builder on the same machine.

It is recommended that you familiarize yourself with the help documentation supplied with Windows AIK. To access the documentation, select **Microsoft Windows AIK -> Documentation** from the start menu.

Using Bart PE

You can create a Bart PE image with Acronis Plug-in using the Bart PE Builder. Please refer to Building Bart PE with Acronis Plug-in from Windows distribution (p. 227) for details.

6.7.1.1. Bootable Media Builder

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 RIS configuration, the machine must have a network connection. Bootable Media Builder can also create an ISO image of bootable disk to burn it later on a blank disk.

When using the media builder, you have to specify:

1. The Acronis bootable components to be placed on the media.

2. [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 in 10 seconds after the menu is displayed. This enables unattended onsite operation when booting from a PXE server or RIS.
3. [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.
4. [optional] Network settings (p. 223):
   - TCP/IP settings to be assigned to the machine network adapters.
5. [optional] Network port (p. 224):
   - the TCP port that the bootable agent listens for incoming connection.
6. The type of the 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 bootable disc to burn it later on a blank disc
   - upload the selected components to Acronis PXE Server
   - upload the selected components to a RIS.
7. [optional] Windows system drivers to be used by Acronis Universal Restore (p. 224). This window appears only if the Acronis Universal Restore add-on is installed and the media other than PXE or RIS is selected.
8. Path to the media ISO file or the name or IP and credentials for PXE or RIS.

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 the 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, not only for the machine where the media is created. To do so, configure the NICs according to their slots order on that machine: NIC1 occupies the slot closest to the processor, NIC2 is in the next slot and so on. When the bootable agent starts on that machine, it will find no NICs with known MAC addresses and will configure the NICs in the same order as you did.

Example

The bootable agent could use one of the network adapters for communication with the management console through the production network. Automatic configuration could be done for this connection. The 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
- 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. The temporary port configuration is available. While connecting the console to the agent, specify the port for the given session in the URL notation <Agent-IP>:<port>.

Drivers for Acronis Universal Restore

While creating bootable media, you have an option to add Windows drivers to the media. The drivers will be used by Acronis Universal Restore when recovering Windows on a machine with a dissimilar processor, different motherboard or other mass storage device than in the backed up system.

You will be able to configure the Acronis Universal Restore:

- to search the media for the drivers that best fit the target hardware

- to get from the media the mass-storage drivers that you explicitly specify. 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 Acronis Universal Restore (p. 202).

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 conditions that:

1. The Acronis 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 flash drive. Drivers cannot be uploaded on a PXE server or 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**.

6.7.1.2. 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. Install the Acronis Plug-in for WinPE from the Acronis Backup & Recovery 10 setup file.
2. Unpack all files of your WinPE 1.x ISO to a separate folder on the hard disk.
3. Select Acronis Win PE ISO Builder from the start menu.
4. Specify path to the folder with the WinPE files.
5. Specify path to the folder with the Acronis Plug-in files (check the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\BartPE\Settings\BartPE for the plug-in location.)

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.

6.7.1.3. Adding the Acronis Plug-in to WinPE 2.x

Acronis WinPE ISO Builder provides three methods of integrating Acronis Backup & Recovery 10 with WinPE 2.x:

- Adding the Acronis Plug-in to the existing PE 2 ISO. This comes in handy when you have to add the plug-in to the previously configured PE 2 ISO that is already in use.
- Creating the PE 2 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 Acronis WinPE ISO Builder on a machine where Windows Automated Installation Kit (AIK) is installed. If you do not have such machine, prepare as described in How to create bootable media (p. 221).

Adding Acronis Plug-in to WinPE 2.x ISO

To add Acronis Plug-in to WinPE 2.x ISO:

1. Do one of the following:

   When adding the plug-in to the existing Win PE 2 ISO, unpack all files of your Win PE 2 ISO to a separate folder on the hard disk.

   When creating a new PE 2 ISO:

   - select from the start menu **Microsoft Windows AIK -> Windows PE Tools Command Prompt**
   - run the **copype.cmd** script to create a folder with Windows PE files. For example, from a command prompt, type:

```
cd Program Files\Windows AIK\Tools\PETools\

copype <arch> <destination>
```

   Where <arch> is the hardware architecture (can be x86, amd64, or ia64) and <destination> is a path to the local folder. For example,

```
copype x86 c:\winpe_x86
```

1. Select Acronis WinPE ISO Builder from the start menu.
2. Specify path to the folder with the WinPE files.
3. Specify path to the folder with the Acronis Plug-in files. (check the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\BartPE\Settings\BartPE for the plug-in location.)
4. Choose whether you want to create ISO or WIM image.
5. Specify the full path to the resulting image file including the file name.
6. Check your settings in the summary screen and click Proceed.

7. 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
```

Adding Acronis Plug-in to WinPE 2.x WIM

1. Select Acronis WinPE ISO Builder from the start menu.
2. Specify path to the source WINPE.WIM file. The standard path to this file for x86 hardware is \Program Files\Windows AIK\Tools\PETools\x86\winpe.wim.
3. Specify path to the folder with the Acronis plug-in files. (check the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\BartPE\Settings\BartPE for the plug-in location.)
4. Specify the full path to the resulting WIM file including the file name.
5. Check your settings in the summary screen and click Proceed.

For how to create a PE image (ISO file) from the resulting WIM file please see the previous section.

For more information on customizing Windows PE, see the Windows Preinstallation Environment User's Guide (Winpe.chm).

6.7.1.4. Building PE-based bootable media

1. Install the Acronis Plug-in for WinPE from the Acronis Backup & Recovery 10 setup file.
2. Get the Bart PE builder.
3. Copy the contents of \Program Files\Acronis\BartPE\BartPE (check the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\BartPE\Settings\BartPE for the plug-in location.) to the %BartPE folder%\plugins\acronis.
4. Insert your Windows distribution CD if you do not have a copy of Windows installation files on the HDD.
5. Start the Bart PE builder.
6. Specify the path to the Windows installation files or Windows distribution CD.
7. Click **Plugins** and check whether the Acronis plug-in is enabled. Enable if disabled.

8. Specify the output folder and the full path to the resulting ISO file including the file name or the media to create.
9. Build the Bart PE.
10. 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.

6.7.2. 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 in the console menu **Connect -> Manage a remote machine** 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 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.

6.7.3. Working under bootable media

Operations on a machine booted with bootable media are very much alike the 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 (scsi1p1, scsi1p2, scsi1p3...)
3. There is no **Navigation** tree in the media GUI. Use the **Navigation** menu item to navigate between views.
4. Tasks cannot be scheduled; in fact, tasks are not created at all. If you need to repeat the operation, configure it from scratch.
5. The log lifetime is limited to the current session. You can save all the log or the filtered log entries to a file.

6.7.4. Acronis PXE Server

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

The 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 time required for booting multiple machines as compared to using physical bootable media.

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. Without the DHCP, you will have to pre-configure and upload on the PXE server the bootable agent for each machine separately.

6.7.4.1. Acronis PXE Server Installation

To install Acronis PXE Server:

1. Run the Acronis Backup & Recovery 10 setup program.
2. Select Acronis PXE Server from the list of 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.

6.7.4.2. 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 the Hard Drive device. The example below shows one of reasonable BIOS configurations. If you don't insert bootable media, the computer will boot from the network.

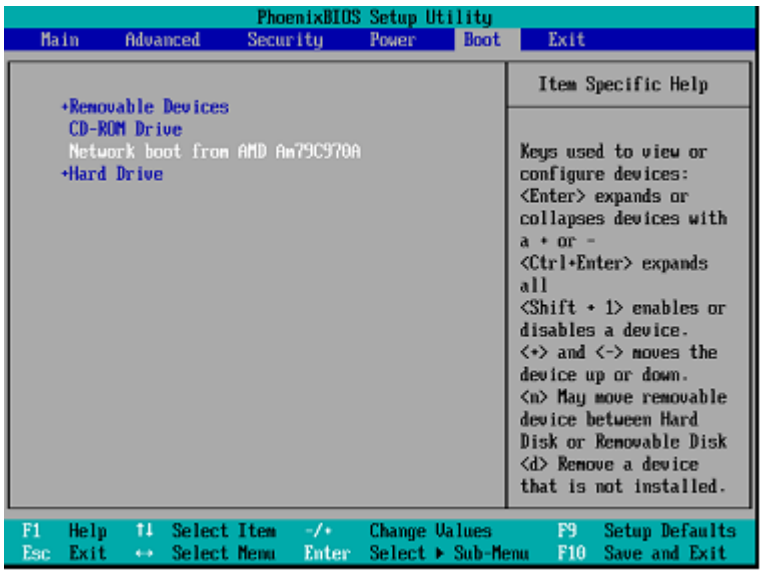

In some BIOS versions, you have to save changes to BIOS after enabling the network interface card so that the card appears in the list of boot devices.

If the hardware has multiple network interface cards, make sure that the card supported by the BIOS has the network cable plugged in.

6.7.4.3. PXE and DHCP on the same server

If Acronis PXE Server and the DHCP server are on the same computer, 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
```

6.7.4.4. Work across subnets

To enable the Acronis PXE Server to work in other 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
http://support.microsoft.com/default.aspx/kb/257579.

6.7.4.5. Configuring Acronis PXE Server

You can access Acronis PXE server configuration in the following ways:

- when creating bootable media, you can upload bootable components on the server. See details in Rescue Media Builder.
- using direct server configuration.

To perform the direct PXE server configuration:

1. Start the console.
2. Select **Tools -> Configure PXE Server**.
3. Connect to the server.

The available actions:

- **Configure PXE Server** – upload bootable components in the same way as when creating bootable media. The previously uploaded components will be deleted from the PXE server before uploading the newly selected ones.
- **Remove products from PXE Server** – remove any component from the PXE server.
- **Disable PXE Server** – disable the PXE server. The service does not stop, but no longer responds to incoming requests.
- **Enable PXE Server** – enable the disabled PXE server.

6.8. Managing System Restore

Microsoft Windows System Restore tool, available in Windows XP and Windows Vista operating systems, is used to undo harmful changes to the system without losing recently changed or created user data.

To run the System Restore tool or find out more about it, select **Start -> Programs -> Accessories -> System Tools -> System Restore**.

You might not need the System Restore tool if you regularly back up the system using Acronis Backup & Recovery 10. Turning off System Restore on any volume will free up to 12% of the volume space. You can do so using Windows System Properties or directly from Acronis Backup & Recovery 10.

To manage System Restore

1. Select **Tools > Manage System Restore**.
2. Select the appropriate parameters as follows:
   - **On for all drives** - turns on System Restore on all volumes at once.
   - **Off for all drives** - turns off System Restore on all volumes at once.

   Turning off the System Restore will delete all the existing restore points. Please make sure you do not need those restore points before proceeding.

   - **Custom** - enables you to choose the volumes to be protected with Windows System Restore.

   You cannot turn off System Restore on the system volume, yet keep it on other volumes.

   Turning off System Restore on a volume deletes all previously created restore points for that volume. Please make sure you do not need those restore points before proceeding.

3. Click **OK** to apply the changes.

6.9. Disk management

Acronis Disk Director Lite is the tool for preparing a computer 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 was placed into a safe storage, the computer 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 could be recovered exactly “as it was” or with any alteration of the disk or volume structure the user might consider necessary.

PLEASE NOTE: All operations on disks and volumes involve a certain risk of data damage. Operations on system, bootable or data volumes must be done very carefully to avoid potential problems with booting process or hard disk data storage.

Operations with hard disks and volumes take certain time and any power loss, unintentional turning off the computer or accidental pressing the Reset button during the procedure could result in volume damage and data loss.

Please take all necessary precautions (p. 232) to avoid possible data loss.

6.9.1. 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. Create a disk image of the disk on which volumes will be created or managed. Having your most important data backed up to another hard disk or CD will allow you to work on disk volumes ensured that your data is safe.

   Acronis Backup & Recovery 10 is an extremely effective comprehensive data backup and recovery solution. It creates a data or disk backup copy stored in a compressed archive file that can be restored in case of any accident.

2. Test your disk to make sure it's fully functional and does not contain bad sectors or file system errors.
3. Do not perform any volume operations while running other software that has low-level disk access, such as antivirus or backup tools. Close these programs before running Acronis Disk Director Lite.

With these simple precautions, you will protect yourself against incidental data loss.

6.9.2. 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, there is the icon "Disk Management" on its Main Window, with which you can start Acronis Disk Director Lite's Graphic User Interface.

Running Acronis Disk Director Lite from a bootable media

You can run Acronis Disk Director Lite on a bare metal or on a crashed computer that cannot boot normally. You can also manage disks and volumes on a non-Windows computer. To do so, you will need bootable media with the standalone Acronis Disk Director Lite version created from Acronis Backup & Recovery 10.

Acronis Disk Director Lite is available only as a part of the Acronis Backup & Recovery 10 package, so you will have to create bootable media using the Acronis Bootable Media Builder. For this, you will need a blank CD-R/RW, DVD+R/RW, several formatted diskettes (the wizard will tell you the exact number), or any other media your server can boot from, such as a flash drive or a Zip drive.

If you have other Acronis products, installed on your computer, you can include standalone versions of these programs on the same bootable disk as well.

6.9.3. Acronis Disk Director Lite main window

The program is controlled through the main window. It includes the menu, toolbar, and the drives tree that represents the computer disk and volume list, providing the necessary general information.

The main window also depicts all unallocated disk space that can be used in volume creation.

Disk and volume information

You might need to know the detailed disk and volume statistics to see whether the created configuration meets your requirements.

General disk and volume statistics are provided in the drives tree of the main window.

- If you select one of the volumes, the main window begins presenting the volume's disk number, assigned letter, label, type, status, size, free space size and file system. To start an action, you can summon the context menu for the selected volume by right-clicking the mouse on it; otherwise you can use the main menu.
- If you select one of the disks compiling a volume, you will be able to call its statistics, including it's number in the Disk Group, size, partition type and status. To start an action, you can summon the context menu for the selected disk by right-clicking the mouse on it; otherwise you can use the main menu.

The bottom part of the main window also graphically depicts the selected volume and its disks as rectangles with basic data on them (label, letter, size, status, type and file system).

Starting the operations

Any operation can be launched:

- From customizable toolbars
- From the volume or disk context menu (both in the main window and the graphic panel)
- From the Operations sidebar list

Note that the list of available operations in the Volume/Disk submenu and the Operations sidebar depends on the selected volume or disk type. The same is true for unallocated space as well.

Displaying operation results

Results of any disk or volume operation are immediately displayed in the Acronis Disk Director Lite main window.

For example, if you create a volume, it will be immediately shown in the hard disk volume list in the main window, as well as in graphical form on the bottom of the main windows.

Any volume changes, including changing the volume letter or label, are also immediately displayed in the main window.

6.9.4. Disk operations

Acronis Disk Director Lite includes the following operations that can be performed on disks:

- Disk Initialization (p. 234) - initializes the new hardware added to the system
- Clone Disk (p. 234) - transfers complete data from the source hard disk to the target
- Convert to GPT (p. 235) - converts an MBR partition table to GPT
- Convert to MBR (p. 236) - converts a GPT partition table to MBR
- Convert to Dynamic (p. 237) - converts a basic disk to dynamic
- Convert to Basic (p. 237) - converts a dynamic disk to basic

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 can not be blocked, close the disk management applications that use this disk and start over. If you can not determine which applications use the disk, close them all.

6.9.4.1. Disk initialization

If you add any new hardware to your computer, 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 some file structure unknown to the computer 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 will prompt you to initialize it:

The main window will show the detected new hardware as a gray block with a grayed miniature, thus indicating that the disk is unusable by the system. The prompt window will provide the basic hardware details such as the disk's number, model and capacity to aid you in the choice of your possible action.

Delayed or immediate disk initialization

If for any reason you want to keep the disk uninitialized, you can decline the initialization prompt so that it would not be triggered in the future, and the disk will remain unusable to the system. If in the future you decide to initialize the disk, you will be able to summon the initialization prompt by right-clicking on the gray block, depicting the disk and choosing the Disk Initialization action in the context menu.

If you want to initialize the new hardware immediately, you can choose whether you want the disk to have the MBR or the GPT partition scheme. You can also decide on the disk's immediate conversion to dynamic, or it can be left basic, depending on the tasks you plan for this new hardware.

Please note, that the new disk Initialization will not be postponed and added to the pending operations list to be committed later. After you accept the settings and click the OK button, the initialization will start immediately. It's a quick procedure and will be finished in a short time, but still please wait until it is over, do not try to interrupt it. If you decide to change the new disk settings it can be done later using the standard Acronis Disk Director Lite disk tools.

After the disk is initialized, it is included into the drive and volume tree, its miniature now is shown green instead of grayed and the necessary new disk info appears. The disk block remains gray though, because after the initialization all the disk space is still 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.

6.9.4.2. Disk clone

Sometimes it becomes necessary to transfer the whole 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 hardware fault; in any case, the reason for such operation can be summed up as necessity to transfer the whole source disk data to a target disk exactly as it is.

Selecting source and target disks

If the program finds two disks, one partitioned and another unpartitioned, it will automatically recognize the partitioned disk as the **source** and the unpartitioned disk as the **target**.

If the program finds several partitioned disks, it will ask the user to select the source, that is, the disk, which data will be transferred to another disk. If there is some data on the disk that was chosen as the target, the user will receive a warning: "**Target disk is not empty. Would you like to overwrite it?**", meaning that all the data currently located at the chosen target disk will be lost irrevocably. If the user agrees to overwrite the existing data, he will be able to proceed with the **Clone** operation.

If the user declines the permission to overwrite the existing data, he will be returned to the **Target disk selection** window.

Target disk size

The **Clone** 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 with the **source**. If the destination has a larger size, then it would be advisable to resize the **source** disk volumes to avoid leaving unallocated space on the **target** disk. The option to **Clone** "as is" remains, but the default method of cloning will be with proportional enlarging all the **source** disk volumes so that on the **target** disk no unallocated space remains.

If the destination is smaller, then the option of cloning "as is" 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 but without any data loss is possible, 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** operation will be impossible and the user will not be able to continue.

Summary page

The final window shows the complete list of the planned **Clone** operation with all the target volume specifications.

If you choose **Proceed**, the planned operations will turn to pending so you will not be able to edit the new volume options, but either **Commit** the pending operations to start the process of transferring the data from the **source** volume to the **target**, or **Cancel** the pending operations, thus annulling the **Clone Volume** operation.

6.9.4.3. Disk MBR to GPT conversion

You would want to convert an MBR basic disk to 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 from the list in the Acronis Disk Director Lite main window.
2. Select **Operations-> Convert to GPT** in the Operations sidebar list, or click **Convert to GPT** on the selected disk on the toolbar (or select it from disk or its graphical representation 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. 247) 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 operation of conversion of the MBR Disk to GPT will fail.

If you plan to install some 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 MBR to GPT conversion

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 dynamic to basic conversion (p. 237) using the **Convert to Basic** operation.
2. basic disk MBR to GPT conversion using the **Convert to GPT** operation.
3. GPT disk basic to dynamic conversion (p. 237) using the **Convert to Dynamic** operation.

6.9.4.4. Disk GPT to MBR conversion

If you plan to install some 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 from the list in the Acronis Disk Director Lite main window.
2. Select **Operations-> Convert to MBR** in the **Operations** sidebar list, or click **Convert to MBR** on the selected disk on the toolbar (or select it from disk or its graphical representation 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 a conversion will stop a disk from being accessed by the system, the operating system will stop loading after such a conversion or some volumes at 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 primary volume, belonging a GPT disk to convert, will be logical one after the operation that 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. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

6.9.4.5. Disk Basic to Dynamic conversion

You would want to convert a basic disk to dynamic in the following cases:

- If you plan to use the disk as a 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 from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Convert to Dynamic** or click **Convert to Dynamic** on the selected disk on the toolbar (or select it from volume or the graphical representation context menu). You receive the 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 computer will be restarted.

Please note: 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 on your computer some OS 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**.

6.9.4.6. Disk Dynamic to Basic conversion

You would want to convert dynamic disks back to basic ones, e.g. if you want to start using on your computer some OS 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 from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Convert to Basic** or click **Convert to Basic** on the selected disk on the toolbar (or select it from volume or the graphical representation context menu). You receive the 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 a 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 into 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.

6.9.5. Volume operations

Acronis Disk Director Lite includes the following operations that can be performed on volumes:

- Create Volume (p. 238) - Creates a new volume with the help of the Create Volume Wizard.
- Delete Volume (p. 242) - Deletes the selected volume.
- Set Active (p. 242) - Sets the selected volume Active so that computer will be able to boot with the OS installed there.
- Change Letter (p. 243) - Changes the selected volume letter
- Change Label (p. 243) - Changes the selected volume label
- Extend Simple/Spanned Volume (p. 246) - Extends the existing volume size, adding to is any unallocated disk space
- Add Mirror (p. 244) - Adds redundancy to a volume, turning it into a mirrored volume or enlarging a mirrored volume redundancy
- Remove Mirror (p. 245) - Gives up additional mirrored volume redundancy, removing one mirror from a volume, thus gaining additional disk space
- Break Mirror (p. 245) - Gives up additional mirrored volume redundancy, splitting a mirrored volume into separate disks, thus gaining additional disk space
- Format Volume (p. 246) - Formats a volume giving it the necessary file system

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 over. If you can not determine which applications use the volume, close them all.

6.9.5.1. Creating a volume

You might need a new volume to:

- Recover a previously saved backup copy at the “exactly as was” configuration;
- Store separately collections of similar files — for example, an MP3 collection or video files on a separate volume;
- Store backups (images) of other volumes/disks on special volume;
- Install a new operating system (or swap file) on a new volume;
- Add new hardware to a computer or server.

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 is 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 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 would be 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 the 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 disk. 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 inherited 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 always on the 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 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

Create Volume wizard lets you create a volume of any type (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 **Wizards -> Create Volume** or a similar **Wizards** list item on the sidebar, or by clicking Create Volume on the toolbar. You can also right-click any unallocated space and select Create Volume in the appearing Context Menu.

Select the type of the volume being created

At the first step you have to specify the volume type you want to create. The following volume types are available:

- Basic
- Simple/Spanned
- Striped
- Mirrored
- RAID-5

You will obtain a brief description of every volume type for better understanding of advantages and limitations of each possible volume architecture.

*If the current operating system, installed on this computer, does not support selected volume type, you will receive the appropriate warning. In this case the **Next** button will be disabled and you will have to select some other 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.

Choose destination disk and number of mirrors

The next wizard page, Choose Destination Disk and Number of Mirrors will prompt you to pick the disks, which space will be used for the volume creation. After you pick 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.

You will have to pick two or more disks to create a Spanned volume, two disks to create a Mirror volume and three or more disks to create RAID-5.

If you click the **Back** button, you will be returned to the previous page, Choose new volume requirements (p. 240).

If you click the **Next** button, you will proceed to the next page, Set volume options (p. 240).

Set volume options

On the next wizard page, Volume Options, 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 select the file system.

The wizard will prompt you to choose one of the Windows file systems: FAT16 (disabled, if the Volume Size has been set for more than 2 Gb), FAT32 (disabled, if the Volume Size has been set for more than 32 Gb), NTFS or to leave the volume **Unformatted**.

In setting the cluster size you can choose between any number in the amount preset for each file system. For **FAT16** file system you can choose between 1 Kb, 2 Kb, 4 Kb, 8 Kb, 16 Kb, 32 Kb and default 64 Kb cluster size presets. For **FAT 32** the range of presets will be 1 Kb, 2 Kb, 4 Kb, 8 Kb and default 16 Kb. For the **NTFS** disk the presents will be 512 bytes, 1 Kb, 2 Kb and the default 4 Kb.

If you are creating a basic volume, which can be made system, this Set Volume Type window, will be different, giving you an opportunity to select Volume type — **Primary** (**Active Primary**) or **Logical**.

Typically **Primary** is selected to install an operating system to a volume. Select **Active** (default) value if you want to install an operating system on this volume to boot at computer startup. If **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, or Active and Primary options will be disabled and you will be able to select only the Logical volume type. The warning message will advise you that an OS installed on this volume will not be bootable.

*If setting a new volume label you use characters, 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, Choose destination disk and number of mirrors (p. 240).

If you click the **Next** button, you will proceed to the next page, Set volume size (p. 241).

Set volume size

On the Volume Size wizard page, you will be able to define the size of the future volume, according to the previously made selections. Using the slider or entering the necessary values into the special windows between the minimum and the maximum values or clicking on the special handle, holding and dragging the borders of the disk's picture by the cursor, you will be able to choose the necessary size between the minimum and the maximum values.

The maximum value normally includes the most possible unallocated space. But in some cases the possible unallocated space and the proposed volume maximum 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 are reserved for the future conversion of the disk from basic to dynamic).

For basic volumes if some unallocated space is left on disk, you also will be able to choose the position of the new volume on disk.

Note that if you plan to create the FAT16 file system on this volume, the volume size over 2 GB, will be disabled despite the available amount of the unallocated space.

If you click the **Back** button, you will be returned to the previous page, Set volume options (p. 240).

If you click the **Next** button, you will proceed to the last, Summary page (p. 242).

Summary page

The final window shows the complete list of the planned operations with all the new volume specifications. It is still possible to change any of the options, if you click the Back button to return to the previous screens one by one.

If you click the **Back** button at this window, you will be returned to the previous page, Set volume size (p. 241).

If you choose **Proceed**, the planned operations will turn to pending so you will not be able to edit the new volume options, but either **Commit** the pending operations to start the volume creation, or **Cancel** the pending operations, thus annulling the new volume creation.

6.9.5.2. Delete volume

This version of Acronis Disk Director Lite has the reduced functionality because it is mostly the tool for preparing bare-metal systems for recovering the previously saved volume images. The features of resizing the existing volumes and creating the new volumes using free space from the existing ones were left with the software full version, 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 a new volume or to change some other 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 **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.

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. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

6.9.5.3. 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 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 active from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Set Active** or a similar item in the **Operations** sidebar list, or click **Set the selected volume as active** on the toolbar (or select it from volume or the graphical representation context menu).

If there is 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.

Please note: even if you have the Operating System on the new active volume, in some cases the computer will not be able to boot from it. You will have to confirm your decision to set the new volume active.

By clicking **OK** in the **Set Active Volume** window, you'll add the pending operation of setting active volume.

(To finish the added operation you will have to commit (p. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

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.

The new volume structure will be graphically represented in the Acronis Disk Director Lite main window immediately.

6.9.5.4. 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 at 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 letters 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 the hard disk and volume on it.
2. Select **Operations -> Change Letter** or a similar item in the **Operations** sidebar list, or click **Change Letter** on the toolbar.
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. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the Acronis Disk Director Lite main window immediately.

6.9.5.5. Change volume label

The volume label is an optional attribute. It’s 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 at such a volume.

In Windows, volume labels are shown in the Explorer disk and folder tree: WIN98(C:), WINXP(D:), DATA(E:), etc. WIN98, WINXP and DATA 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. Select **Operations -> Change Label** or a similar item in the **Operations** sidebar list, or click **Change** the selected volume label on the toolbar.
2. Enter a new label in the **Volume Label** window text field.
3. By clicking **OK** in the **Volume Label** window, you'll add the pending operation of volume label changing.

   *If setting a new volume label you use characters, 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. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new label will be graphically represented in the Acronis Disk Director Lite main window immediately.

6.9.5.6. Volume add mirror

Mirrored volumes provide additional data consistency by storing two or more exact copies of data on more than one disk. So, you may want to provide extra safety to your simple or striped volume, or add some extra fault-tolerance to a mirrored volume, byadding a mirror to it.

If you want to add mirror to a volume:

1. Select a basic, simple, striped or mirrored volume to add mirror to from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Add Mirror** or a similar item in the **Operations** sidebar list, or click **Add Mirror on the selected volume** on the toolbar (or select it from volume or the graphical representation context menu).
3. Select the target disk that will be added to a volume as the mirror. If there is not enough empty space on the selected disk to create the mirror, it will not be available for selection.
4. Click **OK** to proceed with adding mirror operation.

   If you are adding mirror to a basic volume, or the target disk you selected for adding mirror is basic, you will receive the warning that it will be **converted to dynamic**.

   If you do not accept this conversion and therefore click **Cancel**, the operation will be interrupted and no changes will be done to the disk configuration.

   If you accept the disk conversion and click **OK**, you'll add a pending operation of adding the target disk as the mirror to the source volume.

   (To finish the added operation you will have to commit (p. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

As a result of this operation the basic or simple volume will be converted into the mirrored volume, the striped volume will be converted into the striped-mirrored one and to the mirrored volume another mirror will be added. The new volume structure will be graphically represented in the Acronis Disk Director Lite main window immediately.

6.9.5.7. Volume remove mirror

Mirrored volumes provide additional data consistency by storing two or more exact copies of data on more than one disk but that means that some disk space is used for data redundancy. So, you may want to abstain from the volume additional fault-tolerance but gain extra disk space disconnecting a mirror from a mirrored or striped-mirrored volume.

If you want to remove a mirror from a mirrored volume:

1. Select a source striped-mirrored or mirrored volume to remove a mirror from out of the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Remove Mirror** or a similar item in the **Operations** sidebar list, or click **Remove Mirror on the selected volume** on the toolbar (or select it from volume or the graphical representation context menu).

   If you are removing a mirror from a mirrored volume, the selected subdisk will be separated from the disk group, and the volume will become either a simple volume or a mirrored volume with one less mirror, inheriting the source volume's filetype, volume letter and all its data.

   If you are removing a mirror from a striped-mirrored volume, the selected subdisk will be separated from the disk group, and the volume will become either a striped volume or a striped-mirrored volume with one less mirror, inheriting the source volume's filetype, volume letter and all its data.

   The separated subdisk will become an independent dynamic disk retaining its former type, but all its space will be unallocated.
3. If you click **OK** to proceed with the **Remove Mirror** operation, you'll add a pending operation of removing a mirror from a mirrored volume.

   (To finish the added operation you will have to commit (p. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the Acronis Disk Director Lite main window immediately.

6.9.5.8. Break mirrored volume

Mirrored volumes provide additional data consistency by storing two or more exact copies of data on more than one disk but that means that some disk space is used for data redundancy. So, you may want to abstain from the volume additional fault-tolerance but gain extra disk space disconnecting mirrors from a mirrored or striped-mirrored volume. This operation is similar to Remove Mirror operation with the difference that in this case all subdisks are disconnected from the diskgroup and they all retain their format and contain the exact copy of the volume data.

If you want to break mirrored volume:

1. Select a striped-mirrored or mirrored volume to break from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Break Mirrored Volume** or a similar item in the **Operations** sidebar list, or click **Break Mirrored Volume on the selected volume** on the toolbar (or select it from volume or the graphical representation context menu).

   If you are breaking *mirrored volume*, the subdisks of the disk group will be separated, and the volume will become two or more simple volumes of the same filetype, keeping the identical data. One of the new volumes will inherit the letter from the source volume; all other new volumes will be automatically assigned the new letters.

If you are breaking *striped-mirrored volume*, the subdisks of the disk group will be separated, and the volume will become two or more striped volumes of the same filetype, keeping the identical data. One of the new volumes will inherit the letter from the source volume; all other new volumes will be automatically assigned the new letters.

3. If you click **OK** to proceed with the breaking mirrored volume operation, you'll add a pending operation of breaking mirrored volume.

   (To finish the added operation you will have to commit (p. 247) it; Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the Acronis Disk Director Lite main window.

6.9.5.9. Extend Simple/Spanned volume

In some cases when there is some unallocated space at a disk, instead of creating a new volume a user might want to add this space to some volume that already exists in the system.

If you want to extend a simple/spanned volume:

1. Select a volume to extend from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Extend** or a similar item in the **Operations** sidebar list, or click **Extend Simple/Spanned Volume** on the selected volume on the toolbar (or select it from volume or the graphical representation context menu).

   You will be forwarded to the Extend Volume window, where you will be able to set the volume new size.

6.9.5.10. Format volume

You might want to format 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
- to remove the possible volume size limitations before Expand Volume operations
- 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 from the list in the Acronis Disk Director Lite main window.
2. Select **Operations -> Format** or a similar item in the **Operations** sidebar list, or click **Format Volume on the selected volume** on the toolbar (or select it from volume or the graphical representation 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 4 Gb), FAT32 (disabled, if the Volume Size is more than 32 Gb) 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 amount preset for each file system.

   For FAT16 file system you can choose between 1 Kb, 2 Kb, 4 Kb, 8 Kb, 16 Kb, 32 Kb and default 64 Kb cluster size presets.

For FAT 32 the range of cluster size presets will be 1 Kb, 2 Kb, 4 Kb, 8 Kb and default 16 Kb.

For the NTFS disk the cluster size presets will be 512 bytes, 1 Kb, 2 Kb and the default 4 Kb.

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. 247) it. Exiting the program without committing the pending operations will effectively cancel them.)

The new volume structure will be graphically represented in the Acronis Disk Director Lite main window.

6.9.6. Performing disk and volume operations

All operations, which were prepared by the user in the 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.

Therefore, you can view (p. 247) the graphical representation of the new volume structure first and then decide whether to perform (p. 247) the pending operations or cancel them without finalizing the changes.

6.9.6.1. Viewing postponed operations

All pending operations are added to the pending operations list. You can view it in the Pending Operations window by clicking **Operations -> Show**.

This approach enables you to control all planned operations, double-check the intended changes, and, if necessary, cancel operations before they are executed.

Quitting the program without committing the pending operations effectively cancels them, so if you try to exit the program without committing the pending operations, you will receive the appropriate warning.

6.9.6.2. Performing pending operations

To execute planned disk or volume operations, the user has to select one of the two following commands:

- Select **Operations -> Commit** in the main menu
- Click **Commit** toolbar button

To prevent you from performing any unintentional change on your disk, the program will first display the list of all pending operations. Clicking **Proceed** will launch their execution. You will not be able to undo any actions or operations after you chose to Proceed the operation.

You can also Cancel the planned operations, by clicking the **Cancel** button in the list of the pending operations.

Also, quitting the program without committing the pending operations effectively cancels them.

7. Centralized management

7.1. Administering Acronis Backup & Recovery 10 Management Server

7.1.1. Dashboard

7.1.2. Backup policies

To be able to manage and protect multiple machines and other entities as a single entity, the administrator creates a backup plan template called "backup policy". By applying this template to a group of managed entities, the administrator deploys multiple backup plans with a single action. Backup policies exist only on the Acronis Backup & Recovery 10 Management Server.

To find out whether a backup policy is currently being deployed, revoked, updated or no changes are being applied to the policy, the administrator checks the policy Deployment state (p. 248).

The administrator does not have to connect to each machine separately to check whether the data is successfully protected or what actions are currently performed on the machine. Instead, the administrator checks the policy state and status on each managed entity or the cumulative status of the policy (p. 249) on all managed entities the policy is applied to.

To check what actions are currently being performed on a machine after a policy is deployed, the administrator has to connect to the machine and look through the backup plans states (p. 156).

Way of work with backup policies view

- Use the **toolbar**'s operational buttons to create new policies, apply the existing policies to managed entities and perform other operations (see Operations with backup policies (p. 250).)
- Use the **Information** pane's tabs to view detailed information on 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 content of the pane is also duplicated in the **Policy details** window (see Policy details (p. 252).)
- Use filtering and sorting capabilities of the policy table for easy browsing and examination (see Filtering and sorting backup policies (p. 251).)

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 on the backup policy deployment state on a single machine, see the Understanding backup policies (p. 71) section.

7.1.2.2. Backup policy statuses

A backup policy status is the cumulative status of the policy statuses on all managed machines the policy is applied to. For example, if the policy is applied to three managed 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 managed machines is the cumulative status of the policy statuses on the managed machines included in the group.

The following table shows summary of possible backup policy statuses.

	Status	How it is determined	How to handle
1	**Error**	The policy status on at least one managed entity is "Error." Otherwise, see 2.	View the log or identify the failed tasks 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 backup policy to prevent the future failure
2	**Warning**	The policy status on at least one managed entity is "Warning." Otherwise, see 3.	View the log to read the warnings -> [optionally] Perform actions to prevent the future warnings or failure.
3	**OK**	The policy status on all managed entities is "OK."	No action is required. Note that if a backup policy is not applied to any managed entity, 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 any 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 the **Failed** 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** 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 the future failure.

*Use alerts on 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 any 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 the **Succeeded** with warnings 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 warnings reason: 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 failed tasks by backup plans or by machines.*

*In the **Log** view, apply the ⚠ 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 the future warnings or failure.

*Use alerts on the **Dashboard** to quickly access the "error" and "warning" log entries.*

What to do if a policy status is OK

No action is required.

What to do if a policy status is blank

To have state and status, the policy must be applied to at least one managed entity. The blank state and status means that the policy has not been applied or has been revoked from all managed entities.

7.1.2.3. Operations with backup policies

The following is a guidance for you to perform operations with backup policies.

To	Do
Create a backup policy	Click **Create backup policy**. The procedure of creating a backup policy is described in-depth in the Creating a backup policy (p. 291) section.
Apply policy to machines or groups	Click **Apply to**. In the Machines selection (p. 251) window, specify machines (groups) the backup policy
Edit a policy	Click **Edit**. Editing policies is performed in the same as creating (p. 291).

Delete a policy	Click **Delete**. As a result, the policy will be revoked from the managed entities (machines, groups) it was applied to and deleted from the management server.
View details of a policy	Click **View details**. In the Policy details (p. 252) window, examine information on the selected policy.
View tasks of a policy	Click **View tasks**. The Tasks (p. 272) view will display a list of the tasks related to the selected policy.
View log of a policy	Click **View log**. The Log (p. 275) view will display a list of the log entries related to the selected policy.
Refresh a list of machines	Click **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 be retrieved from the management server not immediately due to some latency. The manual refresh guarantees that the most recent data is displayed.

7.1.2.4. 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 deploys the policy to the machines, or groups. On each of the machines, the agent creates a backup plan or multiple plans if multiple managed entities reside on the machine.

7.1.2.5. Filtering and sorting backup policies

The following is a guidance for you to filter and sort backup policies.

To	Do
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 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

The table has seven columns that are displayed by default: **Name**, **Source type**, **Deployment state**, **Status**, **Schedule**, **Owner**, **Comments**. If required, you can hide the shown columns.

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 column headers presenting in the table.
2. Click the items you want to be displayed/hidden.

7.1.2.6. Policy details

Accumulates in five tabs all information on the selected backup policy. Lets 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 on the selected policy from the table of backup policies.

Source

The tab displays information on the type of source is backed up and backup selection rules.

Destination

The tab displays information on the target location.

Settings

The tab displays information on the backup scheme used and backup options modified.

Applied to

The tab displays a list of the machines and the groups the selected policy is applied to.

Operations

To	Do
View details of the machine/group.	Click **View details**. In the Machine details (p. 259)/Group details (p. 267) window, examine all information on the select machine/group.
View tasks of the machine/group	Click **View tasks**. The Tasks (p. 272) view will display a list of the tasks related to the selected machine/group
View log of the machine/group	Click **View log**. The Log (p. 275) view will display a list of the log entries related to the selected machine/group.
Revoke policy from the machine/group.	Click **Revoke**. The management server will revoke the policy from the selected machine or group of machines the policy was deployed to. The policy itself remains on the management server.

7.1.3. Machines

Machine is one of Acronis Backup & Recovery 10 Management Server managed entities, to which a backup policy can be applied. The management server administrator adds a machine, or imports machines from Active Directory, or from text files. Once a machine is added to the management server, it becomes available for grouping, applying the backup policies and monitoring the activities related to the data protection.

Managed machines appear on an management server as a result of registration (p. 325).

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. 156) (both local and centralized) existing on the machine and backup policies (p. 249) 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 and the like.

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 an management server. The group cannot be deleted or renamed. Built-in group cannot include nested groups. A backup policy can be applied to a built-in group. The example of built-in group is the **All 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 a managed machine appears on the management server, the server analyzes the machines's properties. 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. 63) section.

To learn more about how policies are applied to machines and groups, see the Policies on machines and groups (p. 66) section.

Way of working with machines

- Use the **toolbar**'s operational buttons to add machines to the management server and perform other operations—see Actions on machines (p. 264).
- Use the **Information** pane's tabs to view detailed information on the selected machine and perform additional operations, such as start/stop tasks, revoke policies, etc. The content of the panel is also duplicated in the **Machine details** (p. 259) window.
- Use filtering and sorting (p. 264) capabilities of the machines table for easy browsing and examination the machines in question.

Way of working with groups

- In the **Machines** view, select the group.
- Use the toolbar's operational buttons to perform actions on the selected group (p. 264).
- Use the **Information** panel tabs to perform additional operations on the group. The content of the panel is also duplicated in the **Group details** (p. 267) window.

7.1.3.1. Actions on machines

The following is a guidance for you to perform operations with machines.

Registering machines on the management server

Once the machine is added or imported to the **All 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 **Machines** view or the **All machines** group in the navigation tree.

To	Do
Add a new machine to the management server	Click **Add machine to AMS**. In the **Add machine** (p. 257) window, select the machine that needs to be added to the management server.
Import machines from Active Directory	Click **Import machines from AD**. In the **Import machines from Active Directory** (p. 258) window, specify machines or organizational units to import to the management server.
Import machines from a text file	Click **Import machines from file**. In the **Import machines from file** (p. 258) window, browse for a .txt or .csv file, containing names (or IP addresses) of machines to import to the management server.

Applying policies

To	Do
Apply a backup policy to a machine	Click **Apply backup policy**. In the **Policy selection** (p. 258) window, specify the backup policy you need to apply to the selected machine.

Grouping actions

To	Do
Move a machine to another static group	Click **Move to another group**. In the **Move to group** (p. 259) 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 member of, will be applied to the machine.
Add a machine to another group	Click **Add to another group**. In the **Add to group** (p. 259) 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.
Remove a machine from the current group	Click **Remove from group**. The backup policies applied to the group will be revoked from the machine automatically.
Add machines to the currently selected group	Click **Add machines to group**. In the **Add to group** (p. 259) window, select the machines that will be added to the group selected machine is a member of. The policies applied to the group will be applied to the machine.
Create a group in the	Click **Create group**.

currently selected group	In the **Create group** (p. 265) window, specify the required parameters of of the group. The new group will be created in the group, the selected machine is a member of.

Deleting the selected machine from the management server

To	Do
Delete a machine from the management server	Click **Delete from AMS**. As a result, all backup policies are revoked from the machine if the machine is online at the moment. If the machine 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 gets available to the server. The centralized plans will no longer affect the machine.

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 section.
Recover data	Click **Recover**. This operation is described in depth in the Recovering data (p. 188) 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. 259) window, examine information on the machine.
View tasks existing on a machine	Click **View tasks**. The **Tasks** (p. 272) view will display a list of the tasks, existing on the machine.
View log entries of the machine	Click **View log**. The **Log** (p. 275) view will display a list of the machine's log entries.
Refreshing and synchronizing	
Synchronize information on a machine with the management server	Click **Synchronize**. The management server will query the managed machine and updates the database with most recent information. Along with synchronizing, the refresh operation will be performed automatically in order to update a 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 be retrieved from the management server not immediately due to some latency. The 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. This can be done by adding the machine to the **All machines** group.

To add a machine

1. In the Navigation tree, select **Machines.**
2. Click **Add machine** on the toolbar.
3. In the **IP/Name** field, enter the machine's name or its IP address, or click **Browse...** and browse network for the machine.
4. To provide a valid account for the machine, click **Options>>**, and specify:
   - **User name**. When entering a name of an Active Directory user account, be sure to specify also 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**.

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 being installed on a registered machine becomes registered on the same management server automatically. Multiple agents are jointly registered and deregistered.

The administrator of the Acronis Backup & Recovery 10 Management Server can configure the server's IP/name (ID) in the agent's settings while

- installing the agent
- using the console-agent connection.

This will launch the standard registration procedure.

To register an agent during the agent installation, you have to be logged with the management server administrator's account or provide the server administrator's credentials on prompt.

To perform registration through a local or remote console-agent connection, you have to be connected with the management server administrator's credentials or provide the server administrator's credentials. On connecting, select from he menu **Options - Machine options - Machine management**, then opt for **Centralized management** and then enter the management server's ID and the server administrator's credentials.

Importing machines from Active Directory

To import machines from Active Directory

1. In the **Search for** field, type a machine or/and an organizational unit name, then click **Search**.
2. The left part of the window displays the machine/organizational unit names 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 items found, click **Add all>>**.
3. The right part of the window displays the items you selected for import. If required, remove the erroneously selected items by using the respective **Remove** and **Remove all** buttons.
4. Click **OK** to start import.

The management console addresses to the agents and initiates the registration procedure.

Multiple agents residing on the same machine are always registered on the same management server. An agent being installed on a machine that already has a registered agent becomes registered on the same management server automatically. Multiple agents are jointly registered and deregistered.

Because registration requires the agent's participation, it cannot take place when the agent is offline.

Importing machines from a file

To import machines from a file

1. In the **Path** field, enter a path to the .txt or .csv file, or click browse and browse to 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
   ```

2. Click **OK**.

The management console addresses to the agents and initiates the registration procedure.

Multiple agents residing on the same machine are always registered on the same management server. An agent being installed on a machine that already has a registered agent becomes registered on the same management server automatically. Multiple agents are jointly registered and deregistered.

Because registration requires the agent's participation, it cannot take place when the agent is offline.

Policy selection

By applying a backup policy, the administrator protects the registered machine.

To apply the backup policy to the selected machine

1. From the list, select the backup policy that you want to apply to the machine.
   Use filters to display the desired policies.
2. Click **OK**.

Adding a machine to another group

Select the group the machine will be added to.

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 group will be deployed to the machine.

Moving a machine to another group

Select the group the machine will be moved to.

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 the group

Select the machines to be added to the group.

Once the machine appears in the group, the policy applied to the group, if any, is deployed to the machine. Physically, the new tasks will be created on the machine. If the machine 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 gets available to the server.

Machine details

Aggregates in four tabs all information on the selected machine. Lets 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** pane.

Machine

The tab displays useful information on the managed machine.

Backup policies

Displays a list of backup policies ever been applied on the selected managed machine and lets perform the following operations:

To	Do
View details of a policy	Click **View details**. In the Policy details (p. 252) window, examine all information related to the selected backup policy.
View tasks of a policy	Click **View tasks**. The Tasks (p. 272) 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. 275) view will display a list of the log entries related to the selected backup policy.
Revoke policy from the machine.	Click **Revoke**. The management server will revoke the policy from the machine the policy was deployed to. 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 the single machine without removing the machine from the group at first.

Filtering and sorting

Filtering and sorting on the backup policies is performed in the same way as for the Backup policies view. See the Filtering and sorting backup policies (p. 251) section for details.

Plans and tasks

Displays a list of the plans (both local and centralized) and tasks existing on the selected managed machine.

Operations

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

To	Do
View details of a plan/task	**Backup plan** Click **View details**. The **Plan's Details** window will appear. **Task** Click **View details**. The **Task's Details** window will appear.
View plan's/task's log	**Backup plan** Click **View log**. You will be taken to the **Log** view containing the list of the plan-related log entries. **Task** Click **View log**.

	You will be taken to the **Log** view containing the list of the task-related log entries.
Run a plan/task	**Backup plan** Click **Run**  . In the **Start Backup Plan** 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. *Why cannot 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.
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 – 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 plan/task	**Backup plan** Click **Edit**. *Why cannot 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 the centralized origin. The direct editing of centralized backup plans, deployed by backup policies is impossible. You need to edit the original backup policy. **Task** Click **Edit**. *Why cannot 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 **Delete**. *What will happen if I delete the backup plan?* The plan's deletion deletes all its tasks. *Why cannot I delete the backup plan?* • The backup plan in the "Running" state A backup plan cannot be deleted, if at least one its task 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 the centralized origin. A centralized plan can be deleted by the management server administrator by revoking the backup policy that produced the plan.

	Task Click **Delete**. *Why cannot 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 **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 be retrieved from the managed machine not immediately due to some latency. The manual refresh guarantees that the most recent data is displayed.

Filtering and sorting

Filtering and sorting on 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. 163) section for details.

Member of

This tab appears only if the selected machine is added to a one or more custom groups and displays a list of the groups the machine is a member of.

Operations

To	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.

Filtering and sorting machines

To	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 backup policies 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, las 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: **Name**, **Status**, **Last connect**, **Last backup**, **Comments**. Other columns (**IP address**, **Operating system**, and **Availability**) are hidden. If required, you can hide the shown 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 column headers presenting in the table.
2. Click the items you want to be displayed/hidden.

7.1.3.2. Actions on groups

The following is a guidance for you to perform operations with selected groups or subgroups

To	Do
Create a custom static or a dynamic group	Click **Create group**. In the **Create group** (p. 265) window, specify the required parameters of of the group.
Apply a backup policy to a group	Click **Apply backup policy**. In the **Policy selection** (p. 258) 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 a group	Click **View details**. In the **Group details** (p. 267) window, examine information on the selected group.
View tasks related to a group	Click **View tasks**. The **Tasks** (p. 272) view will display a list of the tasks, related to the selected group.
View log entries of a group	Click **View log**.

	The **Log** (p. 275) view will display a list of the log entries related to the selected group.
Rename a custom group/subgroup	Click **Rename**. 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. 267) window, change the required parameters of the group.
Move one custom group to another	Click **Move**. In the **Move to group** (p. 265) window, specify a group that will be a parent the selected group.
Delete a custom group	Click **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.
Add machine to a group	Click **Add machine**. In the **Add to group** (p. 259) window, select the group the machine will be member of.
Refresh a list of groups	Click **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 be retrieved from the management server not immediately due to some latency. The manual Refresh guarantees that the most recent data is displayed.

Move one group to another

To move the selected group to another group or to the root

1. In the machines tree, click the group to move the selected group to. You can move a custom group of any type (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**.

Creating custom static and dynamic groups

To create a group

1. In the **Group name** field, enter a name for the created group.

2. Choose the type of the 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.
      - **Organization unit**

        All the machines belonging to the specified organization unit will be members of the dynamic group.
      - **IP range**

        All the machines whose IP addresses are within the specified IP range will be members of the dynamic group.

      Several criteria of the same pattern are combined by logical OR. Several criteria of different patterns are combined by logical AND.
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 all to the same group 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

   Organization unit: SERVERS

   IP range: 192.168.17.0 - 192.168.17.55

   will add to the same group machines whose operating system is Windows 2000 OR Windows 2003 AND belong to the SERVERS organization unit AND whose IP addresses are within the range 192.168.17.0 - 192.168.17.55.

How long a dynamic group member remains 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 does not meet the criteria anymore
- the administrator changes the criteria so that the member does not meet the criteria anymore

There is no way to remove a machine from a dynamic group manually except for deleting the machine from the Management Server.

Editing custom groups

Editing a custom group is performed is the same way as creating (p. 265).

Changing the type of group will result its conversion. Any custom group can be converted to dynamic group if it was static, and vice versa.

- When converting static group to dynamic, provide grouping criteria. All the members existed in the static group not matching the provided criteria will be removed from the dynamic group.
- When converting 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. Lets perform operations with the policies applied to the group.

This information is also duplicated in the Information pane.

Group

Displays information on the group (as in the row of the groups table).

Backup policies

Displays a list of backup policies related to the selected group and lets perform the following operations:

To	Do
View details of a policy	Click **View details**. In the Policy details (p. 252) window, examine all information related to the selected backup policy.
View tasks of a policy	Click **View tasks**. The Tasks (p. 272) 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. 275) view will display a list of the log entries related to the selected backup policy.
Revoke a policy from the group.	Click **Revoke**. 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 on the group is **Revoking**.The policy itself remains on the management server.

Filtering and sorting

Filtering and sorting on the backup policies is performed in the same way as for the Backup policies view. See the Filtering and sorting backup policies (p. 251) section for details.

Sorting groups

To	Do
Sort groups 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.

Configuring the groups table

By default the table has four columns: Name, Machines, Type, Comments. By default all of them are displayed. If required, hide the shown ones.

To show new or hide shown columns

1. Right-click any table header to open the context menu. Items corresponding to headers used in the table by default are ticked.
2. Click the items you want to be displayed/hidden.

7.1.4. Storage nodes

Acronis Backup & Recovery 10 Storage Node is aimed to optimize usage of various resources required for the enterprise data protection. This goal is achieved through organizing managed locations that serve as dedicated storages of the enterprise backup archives.

Storage node enables the administrator to:

- relieve managed machines of unnecessary CPU load by using the storage node-side cleanup and storage node-side validation.
- drastically reduce backup traffic and storage space taken by the archives by using deduplication (p. 74).
- prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted locations.

The administrator controls the storage nodes centrally from the management server (p. 324). The direct console-storage node connection is not possible. Up to 20 storage nodes can be registered on a management server, each server being able to manage up to 20 locations.

To learn more about Acronis Backup & Recovery 10 Storage Node, see the Acronis Backup & Recovery 10 Storage Node (p. 24) section.

The key elements of the Storage nodes view

- **Storage nodes list with toolbar**

  The toolbar contains operational buttons and lets the administrator perform operations with the selected storage node.

  The storage nodes list displays a list of online and offline storage nodes added to the management server. It also informs the administrator on the total number of backups and archives on the storage node.

- **Information panel**

  Contains the detailed information on the selected storage node—which is also duplicated in the Storage node details (p. 271) window—and the location managed by it; and lets storage node administrators manage the compacting task.

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 and its managed locations.

   Note: *You can skip this step if both the storage node and the users' machines are in one Active Directory domain.*

   For information about user rights on a storage node and in its managed locations, see User rights on a storage node (p. 82).
3. Add the storage node to the Acronis Backup & Recovery 10 Management Server—see Adding a storage node (p. 271).
4. Create one or more centralized managed locations: for each location, specify path to the location, the storage node that will manage the location, and deduplication and encryption options—see Creating a managed centralized location (p. 129).
5. Create backup policies (p. 291) or backup plans that will use the centralized location.

- *[optional]* Use the **Centralized location** view (p. 126) to examine the location, perform operations with location and its content (see Operations with centralized locations (p. 128).)

7.1.4.1. Operations with storage nodes

To perform any operation (except for adding and refreshing) with a Acronis Backup & Recovery 10 Storage Node, you must select it first.

All the operations described here, are performed by clicking the corresponding buttons on the toolbar. These operations can be also accessed from the **[Storage node name] Actions** bar (on the **Actions and Tools** pane) and from the **Storage node** item of the main menu.

The following is a brief guidance for you to perform operations with storage nodes.

To	Do
Add a storage node to the management	Click **Add**.

server	In the Add storage node (p. 271) window, specify the machine the storage node resides on. Once the storage node is added to the management server, it becomes registered and available for managing centralized locations. Registration provides a trusted relationship between the storage node, residing on the machine, and the management server.
Remove a storage node from the management server	Click **Remove**. Once the storage node is removed from the management server, the locations being managed by the storage node disappear from the locations list (p. 125) and become unavailable for performing operations: all the plans and tasks that use these location will fail. The remove does not delete the storage node physically, it just removes the storage node from the management server. All the storage node's databases and locations remain untouched. It is possible to add the previously removed storage node. In this case, all the locations managed by the storage node, appear in the locations list and become available over again for all the plans and tasks using them.
Create a centralized managed location on the selected storage node	Click **Create location**. The Create managed location page (p. 129) will be opened with the storage node, pre-selected as a storage node for managing the location being created. Perform the steps remained to create the location.
Change the compacting task schedule	After deleting backups either manually or during cleanup from deduplicating locations, unreferenced data may appear in deduplicating locations 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. 141), weekly (p. 143), and monthly (p. 145) 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 storage node and manage the compacting task manually	Click **View details**. In the Storage node details (p. 271) window (the content of which is duplicated on the **Information** panel), examine information on the storage node, locations residing on it. You can manually start and stop the compacting task.
Refresh a list of storage nodes	Click **Refresh**. 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 be retrieved from the management server not immediately due to some latency. The 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 network for the machine.
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 specify also the domain name (DOMAIN\Username). 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**.

The management console addresses to the machine the storage node resides on and initiates the registration procedure. Because registration requires the storage noder's participation, it cannot take place when the machine is offline.

7.1.4.2. Storage node details

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 on the storage node

- Name - the name of the selected storage node
- IP - the IP address of the machine the storage node resides on
- Availability - whether the storage node is online of offline
- Archives - the total number of backups stored in all the locations managed by the storage node
- Backups - the total number of backups stored within the archives in all the locations managed by the storage node.

Locations

Displays a list of the locations, managed by the storage node.

To open a managed location for detailed examination and performing operations on it, select the location, then click **Open location** (on the tab's toolbar). In the Centralized location (p. 126) view, perform the required actions.

Services

Displays the compacting task scheduling parameters.

Service tasks

Lets the management server administrator to review and manage the compacting task. Only one compacting task is available per storage node.

7.1.5. 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. Status of a task helps you to estimate whether the task is successfully accomplished.

To learn more about task states and statuses, see the Understanding task states and statuses (p. 158) section.

Way of working with tasks

- Use filtering and sorting (p. 275) capabilities to display the desired tasks.
- Select a task to perform operation (p. 272) on it.

7.1.5.1. Operations with tasks

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

To	Do
View details of a task	Click **View details**. In the Tasks details (p. 274) window, examine all information related to the selected task.
View a task's log	Click **View log**. The Log (p. 275) 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 cannot 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 task	Click **Delete**. *Why cannot 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 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 be retrieved from the managed machine not immediately due to some latency. The manual refresh guarantees that the most recent data is displayed.

7.1.5.2. Task details

The **Task Details** window gathers in four tabs all information on the selected task and lets you perform operations with the machines and/or the groups of machines the policy is applied to.

This information is also duplicated on the **Information** pane.

1. Backup task
   - Backup disk
   - Backup file
2. Restore task
   - Restore disk
   - Restore files
   - Restore volumes
   - Restore volumes as files
3. Validation task
   - Archive validation
   - Backup validation
   - Location validation
4. Cleanup task
5. Export task
6. Convert to VD
7. Acronis Secure Zone (ASZ) tasks
   - Acronis Secure Zone creation
   - Acronis Secure Zone deletion
   - Acronis Secure Zone management (resizing)

Task

The **Task** tab is common for all types tasks.

If the task requires user interaction, at the top of the tab the respective message appears. It contains a brief description of a problem and action buttons, letting you to retry or cancel the task.

Below the message, the general information on the task is displayed. It is collected from single row of the task table.

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, backups' 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 location the backup stored in.

Settings

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

Progress

The **Task** tab is common for all types of tasks. Provides information on the task's progress (if it is running currently running), elapsed time, backup speed, etc.

7.1.5.3. Filtering and sorting tasks

The following is a guidance for you to filter and sort tasks.

To	Do
Set a number of tasks to display	Options -> ... and set the desired number. Further filtering and sorting will performed with the given number of tasks only.
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's name, or the backup plan's name) in the field below the corresponding column header. As a result you will see the list of tasks, whose names (owner's names, or backup plan names) fully or just partly coincide with the entered value.
Filter tasks by type, execution state, status, type, origin, last result, schedule.	In a field below the corresponding header, select the required value from the list.

Configuring tasks table

The table has eight columns that are displayed by default. Other columns are hidden. If required, you can hide the shown 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 column headers presenting in the table.
2. Click the items you want to be displayed/hidden.

7.1.6. Log

The Acronis Backup & Recovery 10 log stores the history of actions the program does on the machine or a user does on the machine using the program. For example, when a user edits an Acronis task, an entry is added to the log. When the program 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 locations, 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 Centralized Storage Server has its own local event log. This log events are accessible through the centralized log only.

Local log entry details (p. 166)

Centralized event log

Centralized event log is accessible when the console connected to the management server. With centralized log, you can examine the history of centralized management operations, such as creating a managed entities group, applying a policy, managing a centralized location; as well as the history of operations logged in the local logs of the registered machines and the storage nodes.

Physically, centralized event log is a table in the dedicated Microsoft SQL database. The table contains log entries for events occurred on the management server and the local logs' entries extended to the centralized log entry format. You can select the types of entries to be collected from the local logs to the centralized database or disable collecting local entries at all (see Collecting log.) The management server events logging cannot be adjusted or disabled.

Note that there are limitations on log entries number in the centralized event log because the SQL Express database has 4 Gb limitation for the database size.

Centralized log entry details (p. 279)

The key elements of the Log view

- **Toolbar** with operational and filtering buttons, lets you to perform actions with the selected log entry.
- **Log table** - displays the list of log entries according to applied filters.
- **Information** pane (at the foot of the page) - displays the detailed information on the currently selected log entry. The panel is collapsed by default. To expand the panel, click the  chevron. This information can be also accessed by clicking on **View Details** item in the **Log Actions** bar.
- **Log Actions** bar (on the **Actions and Tools** pane) - contains actions you can perform with selected log entries. All actions are also duplicated in the menu and toolbar. Thus, you can choose the way optimal for your work.

Way of working with log entries

- Use filters to display the desired log entries only (see Filtering and sorting log entries (p. 278).)

- Select a log entry (or log entries) to perform operation on it (see Operations with log entries (p. 277).)
- Configure the log table appearance (see Configuring log table (p. 278).)

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 selected machine or group.
Backup policies	Select a backup policy, then click 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.

7.1.6.1. Operations with log entries

All the operations described below are performed by clicking the corresponding items on the Log Actions bar. In addition, the operations can be also performed with the context menu (by right-clicking the log), or with the toolbar (by clicking the corresponding buttons).

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

To	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 other 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	1 Select a single log entry or multiple log entries. 2 Click **Save Selected to File**. 3 In the opened window, specify a path and a name for the file.
Save all the log entries to a file	1 Make sure, that the filters are not set. 2 Click **Save All to File**.

	3 In the opened window, specify a path and a name for the file.
Save all the filtered log entries to a file	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

7.1.6.2. Filtering and sorting log entries

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

To	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's buttons: to filter error messages to filter warning messages to filter information messages
Filter log entries by the original backup plan or task	Under **Backup plan** column header, select the backup plan from the list.
Filter log entries by owner	Type an owner's name in the field below the **Owner** header. As a result, in the drop-down list you will see the owner names 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.

7.1.6.3. Configuring log table

By default the table has five columns: Type, Date and time, Backup plan, Task, Message. If required, you can display other columns (Code, Module, and Owner) and/or hide the shown.

To show new or hide shown columns

1. Right-click any table header to open the context menu. Items corresponding to headers used in the table by default are ticked.

2. Click the items you want to be displayed/hidden.

7.1.6.4. Centralized log entry details

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

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

Log entry data fields

A centralized log entry contains the following data fields:

- Type - Type of the 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 the managed entity the event was occurred on (if any)
- Managed entity - The name of the managed entity the event was occurred on (if any)
- Machine - The name of the machine the event was occurred on (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 (only under operating system)
- Message - The event text description.

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

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

7.1.7. Configuring Acronis Backup & Recovery 10 components

There are three ways to configure various parameters of Acronis Backup & Recovery 10 components in Windows:

- By using the graphical user interface (GUI)
- By using Group Policy in Windows
- By modifying the Windows registry

In Linux, instead of Group Policy and registry, parameters are configured by modifying the correspondent configuration files.

The following subtopics describe each way of configuration and the parameters that can be configured through it.

7.1.7.1. Parameters set through Group Policy

The following are the parameters of Acronis Backup & Recovery 10 components that can be set via Group Policy.

You can set one or more parameters via Group Policy by using Acronis Administrative Template. For information on how to apply the administrative template, see How to apply Acronis Administrative Template (p. 88).

The administrative template contains 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.

Acronis Backup & Recovery 10 Storage Node

The following are the parameters of Acronis Backup & Recovery 10 Storage Node that can be set via Group Policy by using Acronis Administrative Template.

ClientConnectionLimit

Description: Specifies the maximum number of simultaneous connections of Acronis Backup & Recovery 10 Agents to Acronis Backup & Recovery 10 Storage Node.

Possible values: Any integer number between **1** and **2147483647**

Default value: **20**

BackupQueueLimit

Description: Specifies the maximum number of items in the backup queue.

Possible values: Any integer number between **1** and **2147483647**

Default value: **200**

CompactingTriggerThreshold

Description: Specifies the percentage a backup's free space for starting a compacting procedure.

Possible values: Any integer number between **0** and **100**

Default value: **80**

When the amount of free space in a backup within an archive falls below this percentage, a compacting procedure is performed on that backup.

FreeSpaceWarningLimit

Description: Specifies the amount of a managed location's free space, 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 location's free space is the amount of free space on the medium—such as a disk volume—that stores the location.

When the amount of free space in a location equals the value in **FreeSpaceWarningLimit** or less, a warning is recorded in the storage node's log, indicating the location in question. You can view storage node warnings in the Dashboard.

FreeSpaceWarningPercentage

Description: Specifies the amount of a managed location's free space, 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 location is the location's free space plus the size of all archives that are contained in the location.

For example, suppose that two locations, Location A and Location B, are both stored on a disk volume. Suppose further that the size of archives in Location A is 20 GB and the one in Location B is 45 GB. If the volume has 5 GB of free space, then the total size of Location A is 20 GB + 5 GB = 25 GB, and that of Location B is 45 GB + 5 GB = 50 GB, regardless of the size of the volume.

The percentage of free space in a location is the location's free space divided by the location's total size. In the previous example, Location A has 5 GB / 25 GB = 20% of free space, and Location B has 5 GB / 50 GB = 10% of free space.

When the percentage of free space in a location equals the value in **FreeSpaceWarningPercentage** or less, a warning is recorded in the storage node's log, indicating the location in question. You can view storage node warnings in the Dashboard.

Note: *The parameters* ***FreeSpaceWarningLimit*** *and* ***FreeSpaceWarningPercentage*** *are independent of each other: a warning will be recorded every time that either of the thresholds is reached.*

FreeSpaceErrorLimit

Description: Specifies the amount of a managed location's free space, in megabytes, below which an error is recorded in the storage node's log and any backup to the location becomes prohibited.

Possible values: Any integer number between **0** and **2147483647**

Default value: **50**

When the amount of free space in a location equals the value in **FreeSpaceErrorLimit** or less, an error is recorded in the storage node's log. Backups performed to the location will keep failing until the location's free space is above the limit.

AmsHost

Description: Specifies the name or IP address of the machine where Acronis Backup & Recovery 10 Management Server is installed.

Possible values: Any string 0 to 255 characters long

Default value: **%AMSMachineID%**

The environment variable AMSMachineID is defined if both Acronis Backup & Recovery 10 Management Server and Acronis Backup & Recovery 10 Storage Node are installed on the same machine.

Acronis Backup & Recovery 10 Management Server

The following are the parameters of Acronis Backup & Recovery 10 Management Server that can be set via Group Policy by using Acronis Administrative Template.

CollectingLogs

Specifies when to collect log entries from machines managed by Acronis Backup & Recovery 10 Management Server.

This parameter contains two settings:

TraceState

Description: Specifies whether to collect the log entries about Acronis Backup & Recovery 10 components' events from the registered machines.

Possible values: **True** or **False**

Default value: **True**

TraceLevel

Description: Specifies the minimum level of severity of collected entries. Only entries of levels greater than or equal to the value in **TraceLevel** 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)

WindowsEventLog

Specifies when to record Acronis Backup & Recovery 10 Management Server's events into the Application event log in Windows.

This parameter has two settings:

TraceState

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**

TraceLevel

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 **TraceLevel** 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 **TraceState** is set to **True**)

Snmp

Specifies the parameters for sending notifications about Acronis Backup & Recovery 10 Management Server's events by means of Simple Network Management Protocol (SNMP).

This parameter contains three settings:

CommonTrace

Specifies when to send the SNMP notifications.

This setting has two sub-settings:

TraceState

Description: Specifies whether to send the SNMP notifications.

Possible values: **True** or **False**

Default value: **False**

TraceLevel

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 **TraceLevel** 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 **TraceState** is set to **True**)

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

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 six settings:

MaxConnections

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 **MaxConnections**, 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 **PriorityHigh** 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.

MaxWorkers

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 **MaxConnections** described previously).

Period

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 **Period** seconds, by using an available worker thread (see **MaxWorkers** described previously).

If there are fewer worker threads than normal-priority machines, the actual interval between synchronizations may be longer than the value in **Period**.

PeriodHigh

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 setting is analogous to the **Period** setting described previously.

RealTimeMonitoring

Description: Specifies whether to perform a 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 is known as polling mechanism.

If **RealTimeMonitoring** is set to **True**, the management server sends requests to machines to provide new data whenever it appears, and then enters a listening mode. This 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 **MaxConnections** earlier in this topic).

SecondAttemptConnect

Description: Specifies whether to try to connect to a registered machine by using its latest-known IP address after an attempt to connect to it by using its host name 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 **SecondAttemptConnect** is set to **True** and the 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 **SecondAttemptConnect** to **True** only in networks which often experience problems with their DNS servers and provided that the machines' IP addresses change infrequently—as in case 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.

Acronis Backup & Recovery 10 Agent for Windows

The following are the parameters of Acronis Backup & Recovery 10 Agent that can be set via Group Policy by using Acronis Administrative Template.

LicenseCheckInterval

Description: Specifies how often, in days, to check for license availablility on Acronis License Server.

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 specified by **LicenseCheckInterval**.

When the agent could not 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 specified in **MaxTimeWithoutLicenseServer** (see the next parameter).

See also **LicenseServerConnectionRetryInterval** later in this topic.

MaxTimeWithoutLicenseServer

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 **MaxTimeWithoutLicense**, as counted from the moment of installation or of the last successful check.

LicenseServerConnectionRetryInterval

Description: Specifies the interval 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 **LicenseCheckInterval** 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 specified in **LicenseServerConnectionRetryInterval**.

If the value is **0**, no reconnection attempts will be performed; the agent will only check for the license as determined by **LicenseCheckInterval**.

LicenseServerAddress

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

AmsHost

Description: Specifies the network name or IP address of Acronis Backup & Recovery 10 Management Server.

Possible values: Any string 0 to 32765 characters long

Default value: Empty string

CentralizedManaged

Description: Specifies whether the machine is managed by Acronis Backup & Recovery 10 Management Server.

Possible values: **True** or **False**

Default value: **False**

If the value is **False**, the machine acts as a stand-alone machine.

If the value is **True**, the machine acts as a registered machine managed by the management server whose name is specified in **AmsHost**. When Acronis Backup & Recovery 10 Agent first starts on such machine, it connects to the management server and adds the machine to it.

If you set **CentralizedManaged** to **True**, you must specify the name of the management server in **AmsHost**.

Caution: *When applying Acronis Administrative Template to a domain or organizational unit, keep in mind that all parameters, including* ***CentralizedManaged****, will affect all machines in that domain or organizational unit. In particular, all its currently stand-alone machines will become managed machines or vice versa, depending on the parameter value.*

WindowsEventLog

Specifies when to record Acronis Backup & Recovery 10 Agent's events into the Application event log in Windows.

This parameter has two settings:

TraceState

Description: Specifies whether to record the agent's events into the event log.

Possible values: **True** or **False**

Default value: **False**

TraceLevel

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 **TraceLevel** 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 **TraceState** is set to **True**)

Snmp

Specifies the parameters for sending notifications about Acronis Backup & Recovery 10 Agent's events by means of Simple Network Management Protocol (SNMP).

This parameter has three settings:

CommonTrace

Specifies when to send the SNMP notifications.

This setting has two sub-settings:

TraceState

Description: Specifies whether to send the SNMP notifications.

Possible values: **True** or **False**

Default value: **False**

TraceLevel

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 **TraceLevel** 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 **TraceState** is set to **True**)

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

Community

Description: Specifies the community name for the SNMP notifications.

Possible values: Any string 0 to 32765 characters long

Default value: **public**

Event tracing parameters

Parameters called tracing parameters specify when to record notifications about events from Acronis Backup & Recovery 10 components—Acronis Backup & Recovery 10 Management Server, Acronis Backup & Recovery 10 Agent, and Acronis Backup & Recovery 10 Storage Node—into the event log or a file.

Severity levels

Event notifications are rated on a scale from 0 to 5 based on the event's severity, as shown in the following table:

Level	Name	Description
0	Unknown	Notification about an event whose level of severity is unknown or not applicable
1	Debug	Notification used for debug purposes
2	Information	Informational notification, such as one about a successful completion of an operation or startup of a service
3	Warning	Notification about a possible impending problem, such as low free space in a location
4	Error	Notification about an event that resulted in loss of data or functionality

5	Critical error	Notification about an event that resulted in the termination of a process such as the agent's process

The following is how to set up tracing parameters for machines running Windows and those running Linux.

Parameters

Windows

In Windows, the notifications are recorded into the event log, a file, or both.

Tracing parameters are specified as the following parameters in Acronis Administrative Template:

FileTraceMinLevel

Description: Specifies the minimum severity level of notifications to be recorded in the file . Only notifications of levels greater than or equal to **FileTraceMinLevel** will be recorded.

Possible values: Any severity level between **0** and **5**, or **6** to not record any notifications

Default value: **2** (notifications with severity levels two through five will be recorded)

Win32TraceMinLevel

Description: Specifies the minimum severity level of notifications to be recorded in the System event log. Only notifications of levels greater than or equal to **Win32TraceMinLevel** will be recorded.

Possible values: Any severity level between **0** and **5**, or **6** to not record any notifications

Default value: **4** (notifications about errors and critical errors will be recorded)

Linux

In Linux, notifications are recorded into a file whose name is specified in the configuration file **/etc/Acronis/Trace.config**

The configuration file is a text file containing two parameters, **FileTraceMinLevel** and **FileTraceFileName**, with one or more spaces or tab characters separating the name of a parameter and its value:

```
FileTraceMinLevel    2

FileTraceFileName    /etc/Acronis/Trace.log
```

Both parameters are analogous to those in Windows, as described earlier in this topic.

7.1.7.2. Parameters set through GUI

The following parameters of can be set through the graphical user interface:

- For Acronis Backup & Recovery 10 Management Server: **CollectingLogs**, **WindowsEventLog**, and **Snmp**
- For Acronis Backup & Recovery 10 Agent: **WindowsEventLog** and **Snmp**

You will find the description of these parameters in the correspondent topic about configuration through Group Policy.

If the values of any of these parameters set through Group Policy differ from those set through graphical user interface, the Group-Policy-based parameters take precedence and are effective immediately; the parameters shown in GUI will be changed accordingly.

7.1.7.3. Parameters set through Windows registry

The following two parameters determine paths to Acronis Backup & Recovery 10 Storage Node's internal databases, which contain information about managed locations. They can be modified only by editing the registry.

When to modify

Normally, you do not need to modify the default paths, which are subfolders of the folder determined by the ALLUSERSPROFILE environment variable—such as C:\Documents and Settings\All Users\Acronis\StorageNode.

While the database located in the folder determined by **DatabasePath** is typically small, the tape database, located at **TapeDatabasePath**, may be large if the tape library contains thousands of archives, and you may want to store the tape database on a volume other that the system volume.

Parameters

Important: *We do not recommend modifying these parameters. If you do need to modify either of them, you should do this before creating any correspondent (tape or non-tape) managed locations. Otherwise, the storage node will lose access to those locations until you re-attach them, and re-attaching a location—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 locations database.

This database contains a list of locations that are managed by the storage node, other than tape locations (see the next parameter). Its typical size does not exceed a few kilobytes.

Possible values: Any string 0 to 32765 characters long

Default value: **%ALLUSERSPROFILE%\Application Data\Acronis\StorageServer**

Registry key: HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\StorageNode\Configuration\StorageServer\DatabasePath

TapesDatabasePath

Description: Specifies the folder where Acronis Backup & Recovery 10 Storage Node stores its tape locations database.

This database contains a list of tape locations 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: **%ALLUSERSPROFILE%\Application Data\Acronis\StorageServer\TapesLocation**

Registry key: HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\StorageNode\Configuration\StorageServer\TapesDatabasePath

7.2. 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. 292)

[Optional] You can change the policy account credentials if necessary. To access this option, select the check box for **Advanced view**.

Policy comments

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

What to back up

Items to back up (p. 293)

Specify what 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. 297)

[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 check box for **Advanced view**.

Exclusions (p. 297)

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

Where to back up

Archive (p. 298)

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 location must be available at the time when the management server starts to deploy the policy.

Access credentials (p. 299)

[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 check box for **Advanced view**.

Archive comments

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

How to back up

Backup scheme (p. 300)

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

[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 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 defaults**.

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

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

To learn more about Acronis service credentials, see the Rights for Acronis services (p. 83) section.

To learn more about operations available depending on the user privileges, see the Users' privileges on a managed machine (p. 34) section.

7.2.2. 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. 293)
- Files to back up selection rules
- Virtual machines to back up selection rules

7.2.2.1. 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.

Below is a list of default rules that can be selected. 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.

To include	Do: In the Volumes column	Comments
Windows volumes		
Volume C:	Type **C:** or select it from the list	
System volume	Type or select **[System]**	[System] refers to the volume on which Windows is located.
Boot volume	Type or select **[Boot]**	[Boot] refers to the volume from which

		the machine starts.
Linux volumes		
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 disk drive. For more details, see Note on Linux machines.
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 disk drive. For more details, see Note on Linux machines.
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.
Windows and Linux volumes		
All volumes	Type or select **[All volumes]** from the list	**[All volumes]** substitutes for all volumes both Windows and 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 the volume **C:** on Windows machines and the 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

7.2.2.2. 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 default ones.

Windows

Full path

Point to the folders and files to be backed up. If a path to a file or folder is specified 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

Point to Windows folders. Using 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 folder where program files are located (for example, C:\Program Files)
Windows folder	**%WINDIR%**	Points to the folder where Windows is located (for example, C:\Windows)
• Common folders for all user profiles (for Windows XP) • All user profiles for Windows Vista	**%ALLUSERSPROFILE%**	• *Windows XP*: Points to the folder where the common data of all user profiles are located (for example, C:\Documents and Settings\All Users) • *Windows Vista*: Points to the folder where all user profiles are located (for example, C:\ProgramData)

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**

Placeholders

Placeholders are similar to environment variables, but already pre-customized.

To include	In the Files and folders column, type or select:	Comments
All files on all partitions on a machine	**[All Files]**	Points to all files on all partitions of the machine.
All user profiles existing on a machine	**[All Profiles Folder]**	Points to the folder where all user profiles are located (for example, C:\Documents and Settings\ - for Windows XP, and C:\ProgramData - for Windows Vista).

Linux

To include	In the Files and folders column, type or select:	Comments
Text file on the partition /dev/hda3 mounted to /home/usr/docs	**/dev/hda3/file.txt** or **/home/usr/docs/file.txt**	
Home folder of the common users	**/home**	
The administrative user's home folder	**/root**	
Folder for all user-related programs	**/usr**	

Folder for system configuration files	**/etc**	

7.2.3. Access credentials for source

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

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

7.2.4. Exclusions

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

Predefined exclusions

To use the predefined sets of exclusions, check any of the following:

- **Exclude all hidden files and folders** – all files and folders marked with a "hidden" attribute will be excluded. Hidden files and folders contain user preferences, program and operating system-related data.
- **Exclude all system files and folders** - all files and folders marked with a "system" attribute will be excluded. System files have .sys extension in Windows.

Custom exclusions

If required, you can create your own exclusions for the specific files and/or file types.

To do so, proceed as follows

1. Select the **Exclude files matching the following criteria** check box

2. Click **Add…**
3. In the **Add file exclusion criterion** window specify an exclusion criterion, and then click **OK**. Both explicit rules (by entering an exact file name, or path) and common Windows masking rules (with using wildcard characters) are supported. See the table below for examples of exclusions.
4. The created exclusion appears in the bottom area of the window.
5. Click **OK** to save your settings.

*Note: *.bak, *.~, *.tmp file types are excluded by default.*

Exclusion examples

Criterion	Example	Description
By name	File1.log	Excludes all files named File1.log.
By path	C:\Finance\test.log	Excludes file named test.log and located on disk C:.
Mask (*)	*.log	Excludes all files with .log extension.
Mask (?)	my???.log	Excludes all .log files with names consisting of five symbols and starting with “my”.

Editing custom exclusions

Since you create a custom exclusion, you are able to edit and/or delete it. To do so, click the respective buttons.

7.2.5. Location

Specify path to the archive location and define a name for the new backup archives.

1. Choose where to store machines' archives:
   - **Store all machines' archives in a single location**
     - To store archives in a centralized or personal location, expand the **Centralized locations** or **Personal locations** group and click the location.
     - 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 it.
     - 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.

An FTP server must allow passive mode for file transfers. It is recommended that you change the managed machine firewall settings to open ports 20 and 21 for both TCP and UDP protocols and disable the Routing and Remote Access Windows service.

As appears from the original FTP specification, credentials required for access to FTP servers are transferred through a network as plaintext. This means that 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. 217) section.

2. 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 opted **Store all machines' archives in a single location**, you have to use variables in order to provide the unique archive names within the location.

   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, FINDEPT2, FINDEPT3*), the following three archives will be created in the location:

   FINDEPT1_SYSTEM_BACKUP_Archive1

   FINDEPT2_SYSTEM_BACKUP_Archive1

   FINDEPT3_SYSTEM_BACKUP_Archive1

3. 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.*

7.2.6. 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 that 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 location.

     Specify:

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

2. Click **OK**.

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

7.2.7. 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 right after its creation.
- **Back up later** – to create a backup task for manual start OR schedule the 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 per 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 lifetimes for the daily (referred as "sons"), weekly (referred as "fathers") and monthly (referred as "grandfathers") backups. The backups whose lifetimes have expired 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 number of backup levels (up to 16). In this scheme, the data can be backed up more than once per 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 in the archive 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.

7.2.7.1. Back up now scheme

With the **Back up now** scheme, the backup tasks will be executed immediately (right after you click the **OK** button).

In the **Backup type** field, select whether you want to create a full, incremental or differential backup. If you have not backed up the selected data yet, or the full archive seems too old to append incremental changes to it, choose full backup. Otherwise it is recommended that you create an incremental or differential backup.

7.2.7.2. Back up later scheme

With the Back up later scheme, the backup plan will be executed 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 you have not backed up the selected data yet, or the full archive seems too old to append incremental changes to it, choose full backup. Otherwise it is recommended that you create an incremental or differential backup.
Date and time	Specify when to start the backup plan.
The task will be started manually	Select this check box, if you do not need to put the backup plan on a schedule and wish to start it manually afterwards (by clicking Run on the Tasks view).

7.2.7.3. Simple scheme

With the simple backup scheme you just schedule when and how often to back up data and set the retention rule. At first time the full backup will be created. Next backups will be incremental.

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

Backup	Set up the backup schedule - when and how often to backup the data. To learn more about setting up schedule, see the Scheduling (p. 140) section.
Retention rule	With the simple scheme, only one retention rule is available. Set the backups' lifetime.

7.2.7.4. 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.

	Mo	Tu	We	Th	Fr	Sa	Su
Jan 1—Jan 7	D	D	D	D	W	-	-
Jan 8—Jan 14	D	D	D	D	W	-	-
Jan 15—Jan 21	D	D	D	D	W	-	-

Jan 22—Jan 28	D	D	D	D	M	-	-
Jan 29—Feb 4	D	D	D	D	W	-	-
Feb 5—Feb 11	D	D	D	D	W	-	-
Feb 12—Feb 18	D	D	D	D	W	-	-
Feb 19—Feb 25	D	D	D	D	M	-	-
Feb 26—Mar 4	D	D	D	D	W	-	-

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

- Monthly ("Grandfather") backups are full;
- Weekly ("Father") backups are differential;
- Daily ("Son") backups are incremental.

Parameters

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

Start backup at:	Specifies when to start a backup. The default value is 12:00 PM.
Back up on:	Specifies the days on which to perform a backup. The default value is Workdays.
Weekly/Monthly:	Specifies which of the days selected in the Back up on field you want to reserve for weekly and monthly backups. The default value is Friday. A monthly backup will be performed every fourth such day.
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: 1 week (recommended minimum) Weekly: 1 month Monthly: indefinite The retention period for weekly backups must exceed that for daily backups; monthly backups' retention period must be greater than 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

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

- Back up files every day, including weekends;
- Be able to recover files as of any date over the past seven days;
- Have access to weekly backups of the past month;
- Keep monthly backups indefinitely.

Backup scheme parameters can then be set up as follows.

- Start backup at: **11:00 PM**
- Back up on: **All days**
- Weekly/monthly: **Saturday** (for example)
- Keep backups:
  - Daily: **1 week**
  - Weekly: **1 month**
  - Monthly: **indefinite**

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.

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

Let the plan allow for

- Performing backups at the end of each working day;
- A possibility to recover an accidentally deleted or inadvertently modified file if this has been relatively quickly discovered;
- Having access to a weekly backup for 10 days after it was created;
- Keeping monthly backups for half a year.

Backup scheme parameters can then be set up as follows.

- Start backup at: **6:00 PM**
- Back up on: **Workdays**
- Weekly/monthly: **Friday**
- Keep backups:

- Daily: **1 week**
- Weekly: **10 days**
- Monthly: **6 months**

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

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, 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, over the six months, and keep the monthly backups for five years.

The following GFS scheme would suit such purposes:

- Start backup at: **11:30 PM**
- Back up on: **Tuesday, Thursday, Friday**
- Weekly/monthly: **Friday**
- Keep backups:
  - Daily: **16 weeks**
  - Weekly: **6 months**
  - Monthly: **5 years**

Here, daily incremental backups will be created on Tuesdays and Thursdays, with weekly and monthly backups performed on Fridays. Note that, in order to choose Friday in the Weekly/monthly field, you should first enlist 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, have a five-year history of all documents, etc.

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: **indefinite**

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.

7.2.7.5. 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, weekly, or monthly schedule. Setting up schedule parameters allows create simple schedules (example of a simple daily schedule: a backup task will be run every 1 day on 10 AM) as well as more complex schedules (example of a complex daily schedule: a task will be run on every 3 days, starting form 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 towards recent 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 7 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 differential backup of level 3 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. This enables us to go as far back as 11 days, which is the best-case scenario.

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

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

4	1	2	1	3	1	2	1	4	1	2	1	3

This gives us only a four days' 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 falling 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 equals four days.

7.2.7.6. 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.
Retention rules	Specifies what retention rules will be applied to the archive. For example, the cleanup procedure can be set up to delete all backups older than six months.
Apply the rules (only if the retention rules are set)	Specifies when to apply the retention rules (p. 42). 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 (only if **On schedule** is selected)	Specifies a schedule for archive cleanup. 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

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).

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. 42).

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**: **Machine 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 the machine is 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. 42).

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.

7.2.8. 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 system resources.
3. **Validation schedule** (appears only if you have selected the on schedule in the step 1) - set the schedule of validation. For more information see the Scheduling section.

7.2.9. Backup options

Customize your backup policy by configuring the backup options, such as pre/post backup commands, maximum network bandwidth allocated for the backup stream or the backup archive compression level.

Default options and presets

Each of Acronis agents has its own default backup options with preset values. To access the default options, select **Options - Agent options - Default backup options** from the menu. Once you have changed a preset value, it becomes default for this agent and will be used as default in further backup operations. Nevertheless, while backing up you can override the default settings with the custom settings that will be specific for this plan only.

The policy-specific settings and the option values set by default

Before any setting is changed against default value the **Backup options** section on the **Create backup policy** page contains the only line with the settings status Default.

Clicking the **Change…** link opens **Backup Options** window, where you can change settings of the current backup operation.

After one or more settings are changed against default values, a new line that contains the text description of modified backup options appears in the section. The settings status is changed from the **Default** to **Custom**.

Clicking the **Reset to defaults** resets all backup settings to the default values.

Availability of the backup options

The set of available backup options differs depending on:

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

The following table shows the summary of the backup options availability.

	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. 100) (password + encryption)	+	+	+	+	+	+
Source files exclusion (p. 100)	+	+	+	+	+	+
Pre-post backup commands (p. 101)	+	+	+	+	PE only	PE only
Pre-post data capture commands (p. 103)	+	+	+	+	-	-
Multi-volume snapshot (p. 106)	+	+	-	-	-	-
File-level backup snapshot (p. 105)	-	+	-	+	-	-
Use VSS (p. 106)	+	+	-	-	-	-
Compression level (p. 106)	+	+	+	+	+	+
Backup performance:						
Backup priority (p. 107)	+	+	+	+	-	-
HDD writing speed (p. 107)	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD	Dest: HDD
Network connection speed (p. 107)	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share	Dest: network share
Fast incremental/differential backup (p. 110)	+	-	+	-	-	-
Backup splitting (p. 110)	+	+	+	+	+	+
File-level security (p. 111):						
Preserve files' security settings in archives	-	+	-	-	-	-
In archives, store encrypted files in decrypted state	-	+	-	-	-	-
Media components (p. 112)	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	-	-
Error handling (p. 112):						

Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+
Re-attempt if an error occurs	+	+	+	+	+	+
Ignore bad sectors	+	+	+	+	+	+
Dual destination (p. 113)	+	+	+	+	-	-
Task start conditions (p. 114)	+	+	+	+	-	-
Task failure handling (p. 115):						
Stop executing the backup plan	+	+	+	+	-	-
Restart the failed task	+	+	+	+	-	-
Additional settings (p. 115):						
Overwrite data on a tape without prompting user for confirmation	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape	Dest: Tape
Dismount media after backup is finished						
Ask for first media while creating backup archives on removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media	Dest: removable media
Validate backup after creation	-	-	-	-	+	+
Reset archive bit	-	+	-	-	-	+
Reboot after the backup	-	-	-	-	+	+
Create full backups as synthetic backups	+	+	+	+	+	+
Notifications:						
E-mail (p. 108)	+	+	+	+	-	-
Win Pop-up (p. 108)	+	+	+	+	-	-
Event tracing:						
Windows events log (p. 109)	+	+	-	-	-	-
SNMP (p. 109)	+	+	+	+	-	-

8. Glossary

A

Acronis Active Restore

The Acronis proprietary technology that brings a system online immediately after the system recovery is started. The system boots from the backup (p. 320) 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. 327) image using Acronis WinPE ISO Builder. The resulting bootable media (p. 317) can be used to boot any PC-compatible machine and perform, with certain limitations, most of the direct management (p. 319) operations without help of an operating system. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 318).

Acronis Secure Zone

A hidden volume for storing backup archives (p. 314) within a managed machine (p. 324). 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 handy for mobile users
- enables using Acronis Startup Recovery Manager
- can serve as primary location for the dual destination backup.

In terms of archive locations management, Acronis Secure Zone is considered as personal location (p. 324). Limitation: Acronis Secure Zone cannot be organized on a dynamic disk (p. 320) or a disk using the GPT partitioning style.

Acronis Startup Recovery Manager (ASRM)

The bootable rescue utility (in fact, modification of the bootable agent (p. 316)) residing in the Acronis Secure Zone (p. 313) and configured to start at boot time on pressing F11. Acronis Startup Recovery Manager eliminates the need for a rescue media or network connection to start the rescue utility.

Acronis Startup Recovery Manager is especially handy 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 the ordinary bootable media.

Limitations: cannot be organized on a dynamic disk (p. 320); requires manual configuration of boot loaders, such as LILO and GRUB; requires re-activation of third-party loaders.

Agent (Acronis Backup & Recovery 10 Agent)

An application that performs data backup and recovery and enables other management operations on the machine (p. 323), 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 disks and files and the agents for virtual machines.

Agent-side cleanup

Cleanup (p. 318) performed by an agent (p. 314) according to the backup plan (p. 315) that owns the archive (p. 314). Agent-side cleanup is performed in unmanaged locations (p. 327).

Agent-side validation

Validation (p. 327) performed by an agent (p. 314) according to the backup plan (p. 315) that owns the archive (p. 314). Agent-side validation is performed in unmanaged locations (p. 327).

Archive

See Backup archive (p. 315).

Archive location

A place for storing backup archives (p. 315). A location can be organized on a local or networked drive or detachable media, such as external USB drive. There are no settings for limiting a location size or number of backups in a location. You can limit each archive size using the cleanup (p. 318), but the total size of archives stored in a location is limited by the storage size only.

Archive location (advanced version)

A place for storing backup archives (p. 315).

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

A location 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 location size or number of backups in a location. You can limit each archive size using the cleanup (p. 318), but the total size of archives stored in a location is limited by the storage size only.

B

Backup

The result of a single backup operation (p. 315). Physically, it is a file or a tape record that contains a copy of the backed up data as of specific date and time. Backup files created by Acronis Backup & Recovery 10 have the TIB extension. The TIB files resulted from backups export (p. 322) or consolidation (p. 319) are also called backups.

Backup archive (Archive)

A set of backups (p. 314) created and managed by a backup plan (p. 315). An archive can contain multiple full backups (p. 322) as well as incremental (p. 323) and differential backups (p. 319). Backups belonging to the same archive are always stored in the same location (p. 314). 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 normally managed in their entirety by the backup plan. Manual operations with archives (validation (p. 327), viewing contents, mounting and deleting backups) should be performed by means of Acronis Backup & Recovery 10. The Locations page provides the necessary controls. 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 (p. 323)'s hard disk for the purpose of recovery or reverting the data to a specified date and time.

Backup options

Configuration parameters of a backup operation (p. 315), such as pre/post backup commands, maximum network bandwidth allocated for the backup stream or the backup archive (p. 314) compression level. Backup options are a part of backup plan (p. 315).

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. 315) (archive location and name)
- backup scheme (p. 316), that includes the backup schedule and [optionally] the retention rules
- [optionally] archive validation rules (p. 327)
- backup options (p. 315).

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 right after its creation **(this is a validation rule)**
- protect the archive with the password **(this is an option)**.

Physically, a backup plan is a bundle of tasks (p. 326) configured for execution on a managed machine (p. 324).

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. 316) deployment (centralized plan (p. 317)).

Backup policy (Policy)

A backup plan template created by the management server (p. 324) administrator and stored on the management server. Backup policy contains the same rules as a backup plan, but might not explicitly specify what data items to back up. Instead, the selection rules (p. 325), 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 (like /dev/sda or C:\Windows), the policy will back up this item on each machine where this exact path will be 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. 323).
3. The management server deploys the policy to the machines.
4. On each of the machines, the agent (p. 314) finds the 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 of the machines, the agent creates a backup plan (p. 315) using other rules specified by the policy. Such backup plan is called a centralized plan (p. 317).
6. On each of the machines, the agent creates a set of centralized tasks (p. 318) that will carry out the plan.

Backup scheme

A part of backup plan (p. 315) that includes the backup schedule and [optionally] the retention rules and the cleanup (p. 318) schedule. For example: perform full backup (p. 322) monthly on the last day of the month at 10:00AM and incremental backup (p. 323) 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 (p. 322) and Tower of Hanoi (p. 326), create a custom backup scheme or back up data once.

Bootable agent

A bootable rescue utility that includes most of functionality of the Acronis Backup & Recovery 10 Agent (p. 314). Bootable agent is based on Linux kernel. A machine (p. 323) can be booted into a bootable agent using either bootable media (p. 317) or Acronis PXE Server. Operations can be configured and controlled either locally through the GUI or remotely using the console (p. 318).

Bootable media

A physical media (CD, DVD, USB flash drive or other media supported by a machine (p. 323) BIOS as a boot device) that contains the bootable agent (p. 316) or Windows Preinstallation Environment (WinPE) (p. 327) with the Acronis Plug-in for WinPE (p. 313). 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 be thought of as a kind of bootable media too.

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. 322) on bare metal
- back up sector-by-sector a disk with 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. 324).

The management server has the built-in groups that contain all machines of each type (All physical machines, All virtual machines.) You might not see some of the groups because empty built-in groups are hidden. A backup policy (p. 316) can be applied to any built-in group that contains at least one member.

Built-in groups cannot be deleted, moved to other groups or manually modified. Custom groups cannot be created inside built-in groups. There is no way to remove a machine from a built-in group except for deleting the machine from the management server.

C

Centralized backup plan

A backup plan (p. 315) that appears on the managed machine (p. 324) as a result of deploying a backup policy (p. 316) from the management server (p. 324). Such plan can be modified only through editing the backup policy.

Centralized location

A networked location allotted by the management server (p. 324) administrator to serve as storage for the backup archives (p. 315). A centralized location can be managed by a storage node (p. 325) (managed location (p. 323)) or be unmanaged. The total number and size of archives stored in a centralized location is limited by the storage size only.

As soon as the management server administrator commits creating a centralized location, the location path and name are distributed to all machines registered (p. 324) on the server. The shortcut to the location appears on the machines in the Centralized locations list. Any backup plan (p. 315) existing on the machines, including local plans, can use the centralized location.

On a machine that is not registered on the management server, a user having the privilege to back up to the centralized location can do so by specifying the full path to the location. In case the location is managed, user's archives will be managed by the storage node as well as other archives stored in the location.

Centralized management

Management of the Acronis Backup & Recovery 10 infrastructure through a central management unit known as Acronis Backup & Recovery 10 Management Server (p. 324). The centralized management operations include:

- creating, applying and managing backup policies (p. 316)
- creating and managing static (p. 325) and dynamic groups (p. 321) of machines (p. 323)
- managing the tasks (p. 326) existing on the machines
- creating and managing centralized archive storage locations (p. 317)
- managing storage nodes (p. 325)
- monitoring activities of the Acronis Backup & Recovery 10 components; creating reports; viewing the centralized log and more.

Centralized task

A task (p. 326) belonging to a centralized backup plan (p. 317). Such task appears on the managed machine (p. 324) as a result of deploying a backup policy (p. 316) from the management server (p. 324) and can be modified only through editing the backup policy.

Cleanup

Deleting backups (p. 314) from a backup archive (p. 315) 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. 315) that owns the archive. This operation checks the archive for expired backups and/or exceeding the maximum size. This may result or not result in deleting backups depending on whether the retention rules are violated or not.

For more information please refer to Retention rules (p. 42).

Console (Acronis Backup & Recovery 10 Management Console)

A tool for remote or local access to Acronis agents (p. 314) and Acronis Backup & Recovery 10 Management Server (p. 324).

Having connected the console to the management server, the administrator sets up and manages backup policies (p. 316) and accesses other management server functionality, that is, performs centralized management (p. 318). Using the direct console-agent connection, the administrator performs direct management (p. 319).

Consolidation

Combining two or more subsequent backups (p. 314) belonging to the same archive (p. 315) into a single backup.

Consolidation is used to obtain synthetic backups (p. 326). Consolidation might be needed when deleting backups, either manually or during cleanup (p. 318). For example, the retention rules require to delete a full backup (p. 322) whose life time has expired but retain the next incremental (p. 323) one. The backups will be combined into a single full backup which will be dated the incremental backup date. Since consolidation may take a lot of time and system resources, retention rules provide an option not to delete backups with dependencies. In our example, the full backup will be retained until the incremental one also becomes obsolete. Then both backups will be deleted.

D

Deduplicating location

A managed location (p. 323) in which deduplication (p. 319) 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. 315) stored on storage nodes (p. 325). 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. 322). You need access to this full backup to recover the data from a differential backup.

Direct management

Any management operation that is performed on a managed machine (p. 324) using the direct console (p. 318)-agent (p. 314) connection (as opposed to the centralized management (p. 318) when the operations are configured on the management server (p. 324) and propagated by the server to the managed machines.)

The direct management operations include:

- creating and managing local backup plans (p. 323)

- creating and managing local tasks (p. 323), such as recovery tasks
- creating and managing personal locations (p. 324)
- viewing state, progress and properties of the centralized tasks (p. 318) existing on the machine
- disk management operations, such as clone a disk, create volume, convert volume.

A kind of direct management is performed when using bootable media (p. 317). Some of the direct management operations can be performed via the management server GUI as well. This presumes, however, either an explicit or a background direct connection to the selected machine.

Disk backup (Image)

A backup (p. 314) 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. 320) 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. 323) 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 member 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 be created or connected again, a disk group with 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 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.)

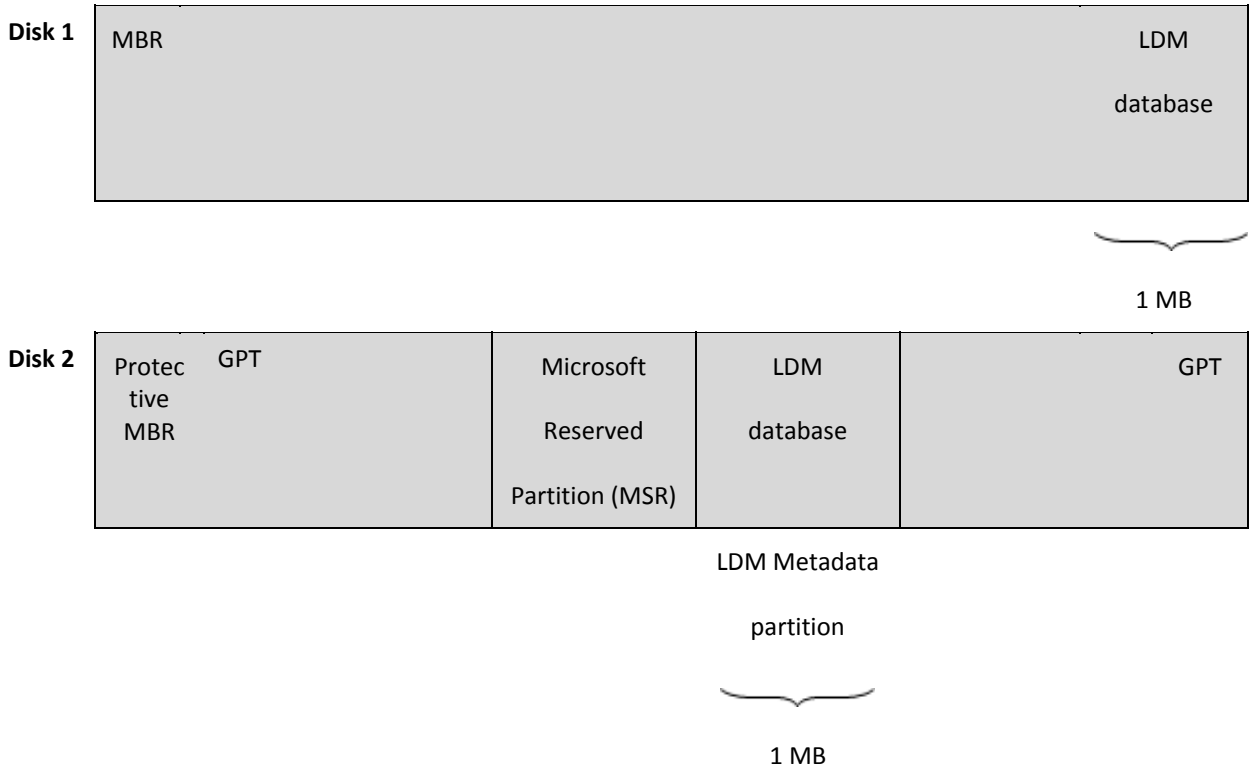

Dynamic disks organized on MBR (Disk 1) and GPT (Disk 2) disks.

For more information about dynamic disks please refer to Microsoft knowledge base articles:

Disk Management (Windows XP Professional Resource Kit) http://technet.microsoft.com/en-us/library/bb457110.aspx 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 computers http://support.microsoft.com/kb/816307

Dynamic group

A group of machines (p. 323) which is populated automatically by a management server (p. 324) 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 the criteria 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. 320), or more exactly, on a disk group (p. 320). Dynamic volumes can span multiple disks. Dynamic volumes are usually configured depending on the desired goal:

- to increase the volume size (a spanned volume)
- to reduce the access time (a striped volume)
- to achieve fault tolerance by introducing redundancy (mirrored and RAID-5 volumes.)

E

Encrypted archive

A backup archive (p. 315) encrypted according to Advanced Encryption Standard (AES). When the encryption option and a password for the archive are set in the backup options (p. 315), each backup belonging to the archive is encrypted by the agent (p. 314) before saving the backup to a location (p. 314).

AES cryptographic algorithm operates in CBC mode and uses the randomly generated key with a user-defined size of 128, 192 or 256 bits. The encryption key is in turn encrypted with AES-256 using SHA-256 hash of the password as a key. The password itself is not stored anywhere on the disk or in the backup, the password hash is used for verification purposes. With this two-level security, the backup data is protected from any unauthorized access, but the lost password recovery is not possible.

Encrypted location

A managed location (p. 323) every write to which is encrypted and every read from which is decrypted transparently by the storage node (p. 325), using a location-specific encryption key stored on the node. In case the storage medium is stolen or accessed by unauthorized person, the malefactor will not be able to decrypt the location contents without access to the storage node. Encrypted archives (p. 322) will be encrypted over the encryption performed by the agent (p. 314).

Export

An operation that creates a self-sufficient duplicate of a backup (p. 314) in the same or another location (p. 314). The operation will create a new archive (p. 315) consisting of one full backup (p. 322) dated the original backup date.

F

Full backup

A self-sufficient backup (p. 314) 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. 316) aimed to keep up the optimal balance between a backup archive (p. 315) size and the number of recovery points (p. 324) available from the archive. GFS enables recovering with daily resolution for last several days, weekly resolution for last several weeks and monthly resolution for any time in the past.

I

Image

The same as Disk backup (p. 320).

Incremental backup

A backup (p. 314) that stores changes to the data against the latest backup. You need access to other backups from the same archive (p. 315) to restore data from an incremental backup.

L

Local backup plan

A backup plan (p. 315) created on the managed machine (p. 324) using direct management (p. 319).

Local task

A task (p. 326) belonging to a local backup plan (p. 323) 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 through 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 location

A centralized location (p. 317) managed by a storage node (p. 325). Archives (p. 315) in a managed location can be accessed as follows:

bsp://ss_name/location_name/archive_name/

Physically, managed locations can reside on a network share, SAN, NAS, or on a hard drive local to the storage node. The storage node performs server-side cleanup (p. 325) and server-side validation (p. 326) for each archive stored in the managed location. An administrator can specify additional operations that the storage node will perform (deduplication (p. 319), encryption.)

Any managed location is self-contained, that is, contains all metadata the storage node needs to manage the location. 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 location is attached to another storage node, the same procedure takes place.

Managed machine

A machine (p. 323) where at least one Acronis Backup & Recovery 10 Agent (p. 314) is installed.

Management server (Acronis Backup & Recovery 10 Management Server)

A central server that manages data protection within the local network. Acronis Backup & Recovery 10 Management Server provides the administrator with:

- a single entry point to the Acronis Backup & Recovery 10 infrastructure
- easy way to protect data on numerous machines (p. 323) using backup policies (see "Backup policy (Policy)" on page 316) and grouping
- enterprise-wide monitoring and reporting functionality
- ability to create centralized locations (p. 317) for storing enterprise backup archives (p. 315)
- ability to manage storage nodes (p. 325).

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

Media builder

A dedicated tool for creating bootable media (p. 317).

P

Personal location

A local or networked archive location (p. 314) created using direct console (p. 318) connection to a managed machine (p. 324). Once a personal location is created, a shortcut to it appears in the Personal locations section of the Navigation pane. Multiple machines can refer to the same physical location, but each of the machines has its own shortcut.

Plan

See Backup plan (p. 315).

Policy

See Backup policy (p. 316).

R

Recovery point

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

Registered machine

A machine (p. 323) managed by a management server (p. 324). A machine can be registered on only one management server at a time. A machine becomes registered as a result of the registration (p. 325) procedure.

Registration

A procedure that adds a managed machine (p. 324) to a management server (p. 324). Registration sets up a trusted relationship between the agent (p. 314) residing on the machine and the server. Once the registration is accomplished, the mutual server-agent authentication takes place each time either of them tries to initiate connection. This helps prevent any attempts of network attackers from establishing the faked connection on behalf of a trusted principal (either the management server or the agent.) It also makes impossible the threat of the type “man-in-the-middle”.

S

Selection rule

A part of backup policy (see "Backup policy (Policy)" on page 316). Enables the management server (p. 324) administrator to select the data to back up within a machine.

Static group

A group of machines which a management server (p. 324) 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 an enterprise data protection. This goal is achieved through organizing managed locations (p. 323). Storage node enables the administrator to:

- relieve managed machines (p. 324) of unnecessary CPU load by using the storage node-side cleanup (p. 325) and storage node-side validation (p. 326)
- drastically reduce backup traffic and storage space taken by the archives (p. 315) by using deduplication (p. 319)
- prevent access to the backup archives, even in case the storage medium is stolen or accessed by a malefactor, by using encrypted locations (p. 322).

Storage node-side cleanup

Cleanup (p. 318) performed by a storage node (p. 325) according to the backup plans (p. 315) that own the archives (p. 315) stored in a managed location (p. 323). Being an alternative to the agent-side cleanup (p. 314), the cleanup on the storage node side relieves the production servers of unnecessary CPU load.

Since the cleanup schedule exists on the machine (p. 323) the agent (p. 314) 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 shows the summary of all cleanup types used in Acronis Backup & Recovery 10.

Cleanup is:	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. 327) performed by a storage node (p. 325) according to the backup plans (p. 315) that own the archives (p. 315) stored in a managed location (p. 323). Being an alternative to the agent-side validation (p. 314), the validation on the storage node side relieves the production servers of unnecessary CPU load.

Synthetic backup

A full backup (p. 322) obtained by consolidating the newly backed up data changes (incremental (p. 323) or differential backup (p. 319)) with the latest backups up to the latest full backup. Thus the full backup is created without transferring all backed up data through the network. The maximum resource saving is achieved when backing up to a managed location (p. 323), since the consolidation (p. 319) will be performed by the storage node (p. 325) without taking resources from the managed machine (p. 324).

This option enables you to have an archive (p. 315) consisting of full backups which makes for faster recovery and faster cleanup because full backups do not depend on each other.

T

Task

In Acronis Backup & Recovery 10, a task is a set of sequential actions to be performed on a managed machine (p. 324) when certain time comes or a certain event occurs. The actions are described in an xml script file that has the .TIS extension. The start condition (schedule) exists in the protected registry keys.

Tower of Hanoi

A popular backup scheme (p. 316) aimed to keep up the optimal balance between a backup archive (p. 315) size and the number of recovery points (p. 324) available from the archive. Unlike the GFS (p. 322) scheme that has only three levels of the recovery points resolution (daily, weekly, monthly resolution), the Tower of Hanoi scheme continuously reduces the resolution as the backup age grows. This allows for very efficient usage of the backup storage.

U

Universal Restore (Acronis Backup & Recovery 10 Universal Restore)

The Acronis proprietary technology that helps boot up Windows on dissimilar hardware or a virtual machine. The Universal Restore handles differences in devices that are critical for the operating system start-up, such as storage controllers, motherboard or chipset.

The Universal Restore is not available:

- when the machine is booted with Acronis Startup Recovery Manager (p. 313) (using F11) or
- the image being recovered is located in Acronis Secure Zone (p. 313) or
- when using Acronis Active Restore (p. 313),

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

The Universal Restore is not available when recovering Linux.

Unmanaged location

Any location (p. 314) that is not a managed location (p. 323).

V

Validation

An operation that checks the possibility of data recovery from a backup (p. 314).

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 in case 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 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 the successful recovery in the future. At least ensure that the backup can be successfully validated using the bootable media (p. 317).

Validation rules

A part of backup plan (p. 315). Rules that define when and how often to perform validation (p. 327) and whether to validate the entire archive (p. 315) or the latest backup in the archive.

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. 313) enables running the Acronis Backup & Recovery 10 Agent (p. 314) in the preinstallation environment.

9. Index

A

B

C

D

E

N

O

P

R

S

T

U

V

W