

Acronis



Acronis Backup 12 Update 1

USER GUIDE

Table of contents

1	What's new in Acronis Backup 12	6
2	Installation	6
2.1	Installation overview	6
2.2	Components	8
2.3	Software requirements	10
2.3.1	Supported web browsers	10
2.3.2	Supported operating systems and environments	10
2.3.3	Supported Microsoft SQL Server versions	12
2.3.4	Supported Microsoft Exchange Server versions	13
2.3.5	Supported Microsoft SharePoint versions	13
2.3.6	Linux packages	13
2.3.7	Compatibility with encryption software	15
2.4	System requirements	16
2.5	Supported file systems	17
2.6	On-premise deployment	18
2.6.1	Installing the management server	18
2.6.2	Adding machines via the web interface	20
2.6.3	Installing agents locally	23
2.6.4	Managing licenses	26
2.7	Cloud deployment	27
2.7.1	Preparation	27
2.7.2	Proxy server settings	28
2.7.3	Installing agents	30
2.7.4	Activating the account	30
2.8	Deploying agents through Group Policy	31
2.9	Updating agents	32
2.10	Uninstalling the product	32
3	Accessing the backup console	34
4	Backup console views	34
5	Backup	35
5.1	Backup plan cheat sheet	37
5.2	Selecting data to back up	38
5.2.1	Selecting disks/volumes	38
5.2.2	Selecting files/folders	39
5.2.3	Selecting system state	40
5.2.4	Selecting ESXi configuration	41
5.3	Selecting a destination	41
5.3.1	About Secure Zone	42
5.4	Schedule	43
5.5	Retention rules	45
5.6	Replication	45
5.7	Encryption	46

5.8	Starting a backup manually	47
5.9	Backup options	47
5.9.1	Backup consolidation	49
5.9.2	Backup validation	49
5.9.3	Changed block tracking (CBT).....	50
5.9.4	Compression level	50
5.9.5	Email notifications.....	50
5.9.6	Error handling	51
5.9.7	Fast incremental/differential backup	52
5.9.8	File filters.....	52
5.9.9	File-level backup snapshot	53
5.9.10	File-level security	54
5.9.11	Log truncation.....	54
5.9.12	LVM snapshotting	54
5.9.13	Mount points.....	54
5.9.14	Multi-volume snapshot.....	55
5.9.15	Performance	56
5.9.16	Pre/Post commands.....	57
5.9.17	Pre/Post data capture commands	58
5.9.18	Scheduling.....	60
5.9.19	Sector-by-sector backup	60
5.9.20	Splitting	61
5.9.21	Task failure handling	61
5.9.22	Volume Shadow Copy Service (VSS)	61
5.9.23	Volume Shadow Copy Service (VSS) for virtual machines	62
5.9.24	Weekly backup	62
5.9.25	Windows event log	62
6	Recovery	63
6.1	Recovery cheat sheet.....	63
6.2	Creating bootable media	63
6.3	Recovering a machine.....	64
6.3.1	Physical machine	64
6.3.2	Physical machine to virtual	65
6.3.3	Virtual machine	66
6.3.4	Recovering disks by using bootable media	68
6.3.5	Using Universal Restore	68
6.4	Recovering files.....	71
6.4.1	Recovering files by using the web interface.....	71
6.4.2	Downloading files from the cloud storage	72
6.4.3	Recovering files by using bootable media	72
6.4.4	Extracting files from local backups.....	73
6.5	Recovering system state	74
6.6	Recovering ESXi configuration	74
6.7	Recovery options	75
6.7.1	Backup validation	76
6.7.2	Date and time for files	76
6.7.3	Error handling	76
6.7.4	File exclusions.....	77
6.7.5	File-level security	77
6.7.6	Flashback.....	77
6.7.7	Full path recovery.....	77
6.7.8	Mount points.....	77

6.7.9	Performance	78
6.7.10	Pre/Post commands.....	78
6.7.11	SID changing	79
6.7.12	VM power management	80
6.7.13	Windows event log	80
7	Operations with backups	80
7.1	The Backups tab.....	80
7.2	Mounting volumes from a backup.....	81
7.3	Deleting backups.....	82
8	Operations with backup plans	82
9	Bootable Media Builder	83
9.1	Linux-based bootable media	83
9.1.1	Kernel parameters.....	84
9.1.2	Network settings	86
9.1.3	Network port	87
9.1.4	Drivers for Universal Restore	87
9.2	WinPE-based bootable media	88
9.2.1	Preparation: WinPE 2.x and 3.x.....	88
9.2.2	Preparation: WinPE 4.0 and later.....	89
9.2.3	Adding Acronis Plug-in to WinPE	89
10	Protecting mobile devices	90
11	Protecting applications	92
11.1	Prerequisites	93
11.2	Database backup.....	94
11.2.1	Selecting SQL databases	94
11.2.2	Selecting Exchange Server data	95
11.3	Application-aware backup	95
11.3.1	Required user rights.....	96
11.4	Recovering SQL databases.....	96
11.4.1	Recovering system databases	98
11.4.2	Attaching SQL Server databases	98
11.5	Recovering Exchange databases.....	99
11.5.1	Mounting Exchange Server databases	100
11.6	Recovering Exchange mailboxes and mailbox items	100
11.6.1	Recovering mailboxes	101
11.6.2	Recovering mailbox items	102
12	Advanced operations with virtual machines	103
12.1	Running a virtual machine from a backup (Instant Restore).....	103
12.1.1	Running the machine	104
12.1.2	Deleting the machine.....	105
12.1.3	Finalizing the machine	106
12.2	Replication of virtual machines	106
12.2.1	Creating a replication plan.....	107
12.2.2	Testing a replica.....	108
12.2.3	Failing over to a replica	108
12.2.4	Replication options	110

12.2.5	Failback options.....	110
12.2.6	Seeding an initial replica	110
12.3	Managing virtualization environments.....	111
12.4	Machine migration.....	112
12.5	Agent for VMware - necessary privileges	112
12.6	Windows Azure and Amazon EC2 virtual machines	114
13	Management server settings	115
13.1	Email server	115
13.2	Email notifications	116
14	Managing groups and accounts	116
14.1	Accounts and groups	116
14.2	Creating a group	117
14.3	Creating an account.....	118
14.4	Creating a report about the service usage	119
14.5	Limiting access to the web interface	120
15	Troubleshooting	121
16	Glossary	123

1 What's new in Acronis Backup 12

- Brand new, modern-looking web interface (p. 34)
- Management Server — the central point for managing all of your backups

Installation

- Cloud or on-premise deployment (p. 6)
- Remote installation of agents (p. 20)

Application backup

- Backup of Exchange databases (p. 95)
- Backup of SQL databases (p. 94)
- Application-aware backup of physical and virtual machines (p. 95)
- Granular recovery of Exchange data (p. 100)

Virtualization

- Backup of Hyper-V virtual machines
- Running a virtual machine from a backup (Instant Restore) (p. 103)
- Replication of ESXi virtual machines (p. 106)
- WAN optimization for replication of virtual machines (replica seeding) (p. 110) — in on-premise deployments only
- File recovery into ESXi virtual machines (p. 71)
- Bare metal recovery of ESXi (p. 41)
- Agent-based backup of virtual machines hosted on Windows Azure and Amazon EC2 (p. 114)
- Flashback — incremental recovery of virtual machines (p. 77)

Common

- Backup of Mac
- Backup of mobile devices (p. 90) — in cloud deployments only
- Managing groups and accounts (p. 116) — in cloud deployments only
- Search for files and folders in a backup (p. 71)
- Powerful file filters (p. 52)

2 Installation

2.1 Installation overview

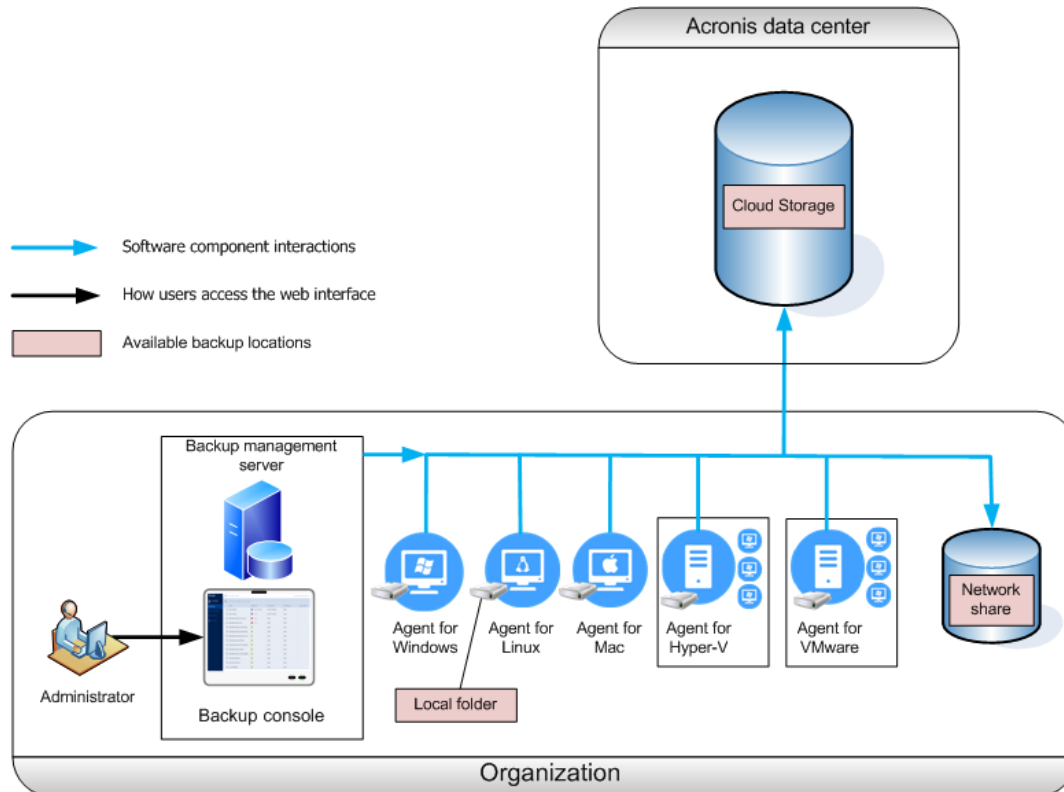
Acronis Backup supports two methods of deployment: on-premise and cloud. The main difference between them is the location of Acronis Backup Management Server.

Acronis Backup Management Server is the central point for managing all of your backups. With the on-premise deployment, it is installed in your local network; with the cloud deployment, it is located in one of the Acronis data centers. The web interface to this server is named a backup console.

Both types of deployment require that a backup agent is installed on each machine that you want to back up. The supported types of storage are also the same: local folders, network shares, and Acronis Cloud Storage. The cloud storage space is sold separately from the Acronis Backup licenses.

On-premise deployment

On-premise deployment means that all of the product components are installed in your local network. This is the only deployment method available with a perpetual license. Also, you have to use this method if your machines are not connected to the Internet.



Management server location

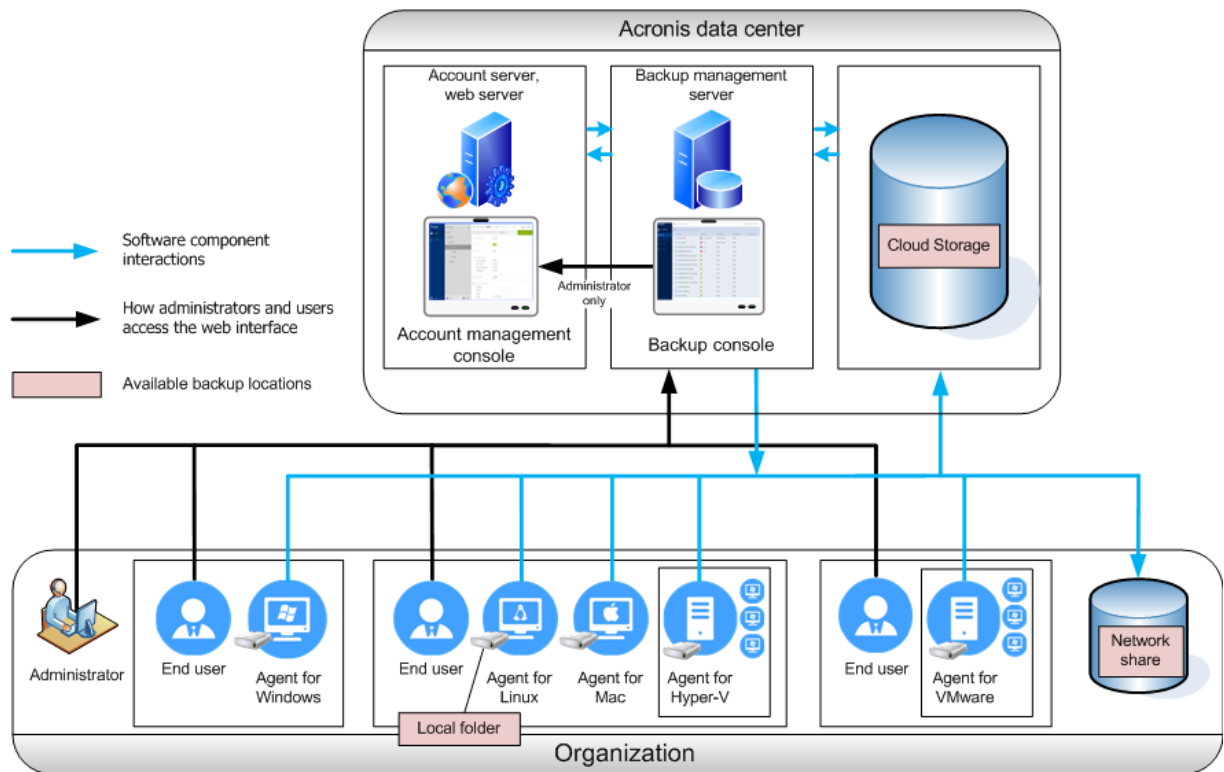
You can install the management server on a machine running either Windows or Linux. Installation in Windows is recommended because you will be able to deploy agents to other machines from the management server. Installation in Linux is recommended in a Linux-only environment. You will need to install an agent locally on the machines that you want to back up.

Cloud deployment

Cloud deployment means that the management server is located in one of the Acronis data centers. The benefit of this approach is that you do not need to maintain the management server in your local network. You can think of Acronis Backup as of a backup service provided to you by Acronis.

Access to the account server enables you to create user accounts, set service usage quotas for them, and create groups of users (units) to reflect the structure of your organization. Every user can access the backup console, download the required agent, and install it on their machines in minutes.

Administrator accounts can be created at the unit or organization level. Each account has a view scoped to their area of control. Users have access only to their own backups.



The following table summarizes differences between the on-premise and cloud deployments.

On-premise deployment	Cloud deployment
<ul style="list-style-type: none"> ▪ On-premise management server ▪ Both subscription and perpetual licenses can be used ▪ Agent for VMware (Virtual Appliance) and Agent for VMware (Windows) ▪ WAN optimization for replication of virtual machines (replica seeding) ▪ Bootable Media Builder ▪ Backup and disk management in bootable media ▪ Upgrade from previous versions of Acronis Backup, including Acronis Backup for VMware ▪ Participation in the Acronis Customer Experience Program 	<ul style="list-style-type: none"> ▪ Group and account management ▪ A subscription license is required ▪ No Agent for VMware (Virtual Appliance) ▪ Mobile backup to cloud

2.2 Components

Agents

Agents are applications that perform data backup, recovery, and other operations on the machines managed by Acronis Backup.

Choose an agent, depending on what you are going to back up. The following table summarizes the information, to help you decide.

Note that Agent for Windows is installed along with Agent for Exchange, Agent for SQL, and Agent for Active Directory. If you install, for example, Agent for SQL, you also will be able to back up the entire machine where the agent is installed.

What are you going to back up?	Which agent to install?	Where to install it?	Agent availability	
			On-prem	Cloud
Disks, volumes, and files on physical machines running Windows	Agent for Windows	On the machine that will be backed up.	+	+
Disks, volumes, and files on physical machines running Linux	Agent for Linux	On the machine that will be backed up.	+	+
Disks, volumes, and files on physical machines running OS X	Agent for Mac	On the machine that will be backed up.	+	+
VMware ESXi virtual machines	Agent for VMware (Windows)	On a Windows machine that has network access to vCenter Server and to the virtual machine storage.*	+	+
	Agent for VMware (Virtual Appliance)	On the ESXi host.	+	-
Hyper-V virtual machines	Agent for Hyper-V	On the Hyper-V host.	+	+
SQL databases	Agent for SQL	On the machine running Microsoft SQL Server.	+	+
Exchange databases	Agent for Exchange	On the machine running the Mailbox role of Microsoft Exchange Server.	+	+
Machines running Active Directory Domain Services	Agent for Active Directory	On the domain controller.	+	+
Virtual machines hosted on Windows Azure	The same as for physical machines**	On the machine that will be backed up.	+	+
Virtual machines hosted on Amazon EC2	The same as for physical machines**	On the machine that will be backed up.	+	+
Mobile devices running Android	Mobile app for Android	On the mobile device that will be backed up.	-	+
Mobile devices running iOS	Mobile app for iOS	On the mobile device that will be backed up.	-	+

*If your ESXi uses a SAN attached storage, install the agent on a machine connected to the same SAN. The agent will back up the virtual machines directly from the storage rather than via the ESXi host and LAN.

**A virtual machine is considered virtual if it is backed up by an external agent. If an agent is installed in the guest system, the backup and recovery operations are the same as with a physical machine. Nevertheless, the machine is counted as virtual when you set quotas for the number of machines in a cloud deployment.

Other components

Component	Function	Where to install it?	Availability	
			On-prem	Cloud

Component	Function	Where to install it?	Availability	
			On-prem	Cloud
Management Server	Manages the agents. Provides the web interface to users.	On a machine running Windows or Linux.	+	–
Bootable Media Builder	Creates bootable media.	On a machine running Windows or Linux.	+	–
Backup Monitor	Enables users to monitor backups outside the web interface.	On a machine running Windows or OS X.	+	+
Command-Line Tool	Provides the command-line interface.	On a machine running Windows or Linux.	+	+

2.3 Software requirements

2.3.1 Supported web browsers

The web interface supports the following web browsers:

- Google Chrome 29 or later
- Mozilla Firefox 23 or later
- Opera 16 or later
- Windows Internet Explorer 10 or later
- Safari 5.1.7 or later running in the OS X and iOS operating systems

In other web browsers (including Safari browsers running in other operating systems), the user interface might be displayed incorrectly, or some functions may be unavailable.

2.3.2 Supported operating systems and environments

2.3.2.1 Agents

Agent for Windows

Windows XP Professional SP2+ (x86, x64)

Windows Server 2003/2003 R2 – Standard and Enterprise editions (x86, x64)

Windows Small Business Server 2003/2003 R2

Windows Vista

in cloud deployments – all editions

in on-premise deployments – all editions except for Vista Home Basic and Vista Home Premium (x86, x64)

Windows Server 2008 – Standard, Enterprise, Datacenter, and Web editions (x86, x64)

Windows Small Business Server 2008

Windows 7

in cloud deployments – all editions

in on-premise deployments – all editions except for the Starter and Home editions (x86, x64)

Windows Server 2008 R2 – Standard, Enterprise, Datacenter, Foundation, and Web editions

Windows MultiPoint Server 2010/2011/2012
Windows Small Business Server 2011 – all editions
Windows 8/8.1 – all editions except for the Windows RT editions (x86, x64)
Windows Server 2012/2012 R2 – all editions
Windows Storage Server 2003/2008/2008 R2/2012/2012 R2
Windows 10 – Home, Pro, Education, and Enterprise editions
Windows Server 2016 – Technical Preview 4

Agent for SQL, Agent for Exchange, and Agent for Active Directory

Each of these agents can be installed on a machine running any operating system listed above and a supported version of the respective application.

Agent for Linux

Linux with kernel from 2.6.9 to 4.5 and glibc 2.3.4 or later
Various x86 and x86_64 Linux distributions, including:
Red Hat Enterprise Linux 4.x, 5.x, 6.x, 7.0, 7.1, 7.2
Ubuntu 9.10, 10.04, 10.10, 11.04, 11.10, 12.04, 12.10, 13.04, 13.10, 14.04, 14.10, 15.04, 15.10, 16.04
Fedora 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23
SUSE Linux Enterprise Server 10 and 11
SUSE Linux Enterprise Server 12 – supported on file systems, except for Btrfs
Debian 4, 5, 6, 7.0, 7.2, 7.4, 7.5, 7.6, 7.7, 8.0, 8.1, 8.2, 8.3, 8.4, 8.5
CentOS 5.x, 6.x, 7, and 7.1
Oracle Linux 5.x, 6.x, 7.0, 7.1, and 7.2 – both Unbreakable Enterprise Kernel and Red Hat Compatible Kernel
CloudLinux 5.x, 6.x, 7, and 7.1
ClearOS 5.x, 6.x, 7, and 7.1

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

Agent for Mac

OS X Mountain Lion 10.8
OS X Mavericks 10.9
OS X Yosemite 10.10
OS X El Capitan 10.11

Agent for VMware (Virtual Appliance)

This agent is delivered as a virtual appliance for running on an ESXi host.
VMware ESXi 4.1, 5.0, 5.1, 5.5, and 6.0

Agent for VMware (Windows)

This agent is delivered as a Windows application for running in any operating system listed above for Agent for Windows with the following exceptions:

- 32-bit operating systems are not supported.
- Windows XP, Windows Server 2003/2003 R2, and Windows Small Business Server 2003/2003 R2 are not supported.

Agent for Hyper-V

Windows Server 2008 (x64) with Hyper-V
Windows Server 2008 R2 with Hyper-V
Microsoft Hyper-V Server 2008/2008 R2
Windows Server 2012/2012 R2 with Hyper-V
Microsoft Hyper-V Server 2012/2012 R2
Windows 8, 8.1 (x64) with Hyper-V
Windows 10 – Pro, Education, and Enterprise editions with Hyper-V
Windows Server 2016 with Hyper-V

2.3.2.2 Management Server (for on-premise deployment only)

In Windows

Windows Server 2008 – Standard, Enterprise, and Datacenter editions (x86, x64)
Windows Small Business Server 2008
Windows 7 – all editions except for the Starter and Home editions (x86, x64)
Windows Server 2008 R2 – Standard, Enterprise, Datacenter, and Foundation editions
Windows MultiPoint Server 2010/2011/2012
Windows Small Business Server 2011 – all editions
Windows 8/8.1 – all editions except for the Windows RT editions (x86, x64)
Windows Server 2012/2012 R2 – all editions
Windows Storage Server 2008/2008 R2/2012/2012 R2
Windows 10 – Home, Pro, Education, and Enterprise editions
Windows Server 2016 – Technical Preview 4

In Linux

Linux with kernel from 2.6.18 to 4.5 and glibc 2.3.4 or later
Various x86_64 Linux distributions, including:
Red Hat Enterprise Linux 4.x, 5.x, 6.x, 7.0, 7.1, 7.2
Ubuntu 9.10, 10.04, 10.10, 11.04, 11.10, 12.04, 12.10, 13.04, 13.10, 14.04, 14.10, 15.04, 15.10, 16.04
Fedora 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23
SUSE Linux Enterprise Server 10, 11, and 12
Debian 4, 5, 6, 7.0, 7.2, 7.4, 7.5, 7.6, 7.7, 8.0, 8.1, 8.2, 8.3, 8.4, 8.5
CentOS 5.x, 6.x, 7, and 7.1
Oracle Linux 5.x, 6.x, 7.0, 7.1, and 7.2 – both Unbreakable Enterprise Kernel and Red Hat Compatible Kernel
CloudLinux 5.x, 6.x, 7, and 7.1

2.3.3 Supported Microsoft SQL Server versions

- Microsoft SQL Server 2016
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012
- Microsoft SQL Server 2008 R2

- Microsoft SQL Server 2008
- Microsoft SQL Server 2005

2.3.4 Supported Microsoft Exchange Server versions

- **Microsoft Exchange Server 2016** – all editions.
- **Microsoft Exchange Server 2013** – all editions, Cumulative Update 1 (CU1) and later.
- **Microsoft Exchange Server 2010** – all editions, all service packs.
- **Microsoft Exchange Server 2007** – all editions, all service packs. Recovery of mailboxes and mailbox items is not supported.

2.3.5 Supported Microsoft SharePoint versions

Acronis Backup 12 supports the following Microsoft SharePoint versions:

- Microsoft SharePoint 2013
- Microsoft SharePoint Server 2010 SP1
- Microsoft SharePoint Foundation 2010 SP1
- Microsoft Office SharePoint Server 2007 SP2*
- Microsoft Windows SharePoint Services 3.0 SP2*

*In order to use SharePoint Explorer with these versions, you need a SharePoint recovery farm to attach the databases to.

The backups or databases from which you extract data must originate from the same SharePoint version as the one where SharePoint Explorer is installed.

2.3.6 Linux packages

To add the necessary modules to the Linux kernel, the setup program needs the following Linux packages:

- The package with kernel headers or sources. The package version must match the kernel version.
- The GNU Compiler Collection (GCC) compiler system. The GCC version must be the one with which the kernel was compiled.
- The Make tool.
- The Perl interpreter.

The names of these packages vary depending on your Linux distribution.

In Red Hat Enterprise Linux, CentOS, and Fedora, the packages normally will be installed by the setup program. In other distributions, you need to install the packages if they are not installed or do not have the required versions.

Are the required packages already installed?

To check whether the packages are already installed, perform these steps:

1. Run the following command to find out the kernel version and the required GCC version:

```
cat /proc/version
```

This command returns lines similar to the following: **Linux version 2.6.35.6** and **gcc version 4.5.1**

2. Run the following command to check whether the Make tool and the GCC compiler are installed:

```
make -v
gcc -v
```

For **gcc**, ensure that the version returned by the command is the same as in the **gcc version** in step 1. For **make**, just ensure that the command runs.

3. Check whether the appropriate version of the packages for building kernel modules is installed:

- In Red Hat Enterprise Linux, CentOS, and Fedora, run the following command:

```
yum list installed | grep kernel-devel
```

- In Ubuntu, run the following commands:

```
dpkg --get-selections | grep linux-headers
dpkg --get-selections | grep linux-image
```

In either case, ensure that the package versions are the same as in **Linux version** in step 1.

4. Run the following command to check whether the Perl interpreter is installed:

```
perl --version
```

If you see the information about the Perl version, the interpreter is installed.

Installing the packages from the repository

The following table lists how to install the required packages in various Linux distributions.

Linux distribution	Package names	How to install
Red Hat Enterprise Linux	kernel-devel gcc make	The setup program will download and install the packages automatically by using your Red Hat subscription.
	perl	Run the following command: <pre>yum install perl</pre>
CentOS Fedora	kernel-devel gcc make	The setup program will download and install the packages automatically.
	perl	Run the following command: <pre>yum install perl</pre>
Ubuntu	linux-headers linux-image gcc make perl	Run the following commands: <pre>sudo apt-get update sudo apt-get install linux-headers-`uname -r` sudo apt-get install linux-image-`uname -r` sudo apt-get install gcc-<package version> sudo apt-get install make sudo apt-get install perl</pre>

The packages will be downloaded from the distribution's repository and installed.

For other Linux distributions, please refer to the distribution's documentation regarding the exact names of the required packages and the ways to install them.

Installing the packages manually

You may need to install the packages **manually** if:

- The machine does not have an active Red Hat subscription or Internet connection.

- The setup program cannot find the **kernel-devel** or **gcc** version corresponding to the kernel version. If the available **kernel-devel** is more recent than your kernel, you need to either update the kernel or install the matching **kernel-devel** version manually.
- You have the required packages on the local network and do not want to spend time for automatic search and downloading.

Obtain the packages from your local network or a trusted third-party website, and install them as follows:

- In Red Hat Enterprise Linux, CentOS, or Fedora, run the following command as the root user:

```
rpm -ivh PACKAGE_FILE1 PACKAGE_FILE2 PACKAGE_FILE3
```

- In Ubuntu, run the following command:

```
sudo dpkg -i PACKAGE_FILE1 PACKAGE_FILE2 PACKAGE_FILE3
```

Example: Installing the packages manually in Fedora 14

Follow these steps to install the required packages in Fedora 14 on a 32-bit machine:

1. Run the following command to determine the kernel version and the required GCC version:

```
cat /proc/version
```

The output of this command includes the following:

```
Linux version 2.6.35.6-45.fc14.i686
gcc version 4.5.1
```

2. Obtain the **kernel-devel** and **gcc** packages that correspond to this kernel version:

```
kernel-devel-2.6.35.6-45.fc14.i686.rpm
gcc-4.5.1-4.fc14.i686.rpm
```

3. Obtain the **make** package for Fedora 14:

```
make-3.82-3.fc14.i686
```

4. Install the packages by running the following commands as the root user:

```
rpm -ivh kernel-devel-2.6.35.6-45.fc14.i686.rpm
rpm -ivh gcc-4.5.1.fc14.i686.rpm
rpm -ivh make-3.82-3.fc14.i686
```

You can specify all these packages in a single **rpm** command. Installing any of these packages may require installing additional packages to resolve dependencies.

2.3.7 Compatibility with encryption software

There are no limitations on backing up and recovering data that is encrypted by *file-level* encryption software.

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

You can back up the data encrypted by the following disk-level encryption software:

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

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

Common installation rule

The strong recommendation is to install the encryption software before installing the backup agents.

The way of using Secure Zone

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

1. Install the encryption software; then, install the agent.
2. Create Secure Zone.
3. Exclude Secure Zone when encrypting the disk or its volumes.

Common backup rule

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

Software-specific recovery procedures

Microsoft BitLocker Drive Encryption

To recover a system that was encrypted by BitLocker:

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

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

McAfee Endpoint Encryption and PGP Whole Disk Encryption

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

If the recovered system fails to boot, rebuild Master Boot Record as described in the following Microsoft knowledge base article: <https://support.microsoft.com/kb/2622803>

2.4 System requirements

The following table summarizes disk space and memory requirements for typical installation cases. The installation is performed with the default settings.

Components to be installed	Occupied disk space	Minimum memory consumption
Agent for Windows	850 MB	150 MB
Agent for Windows and one of the following agents: <ul style="list-style-type: none">▪ Agent for SQL▪ Agent for Exchange	950 MB	170 MB

Agent for Windows and one of the following agents: <ul style="list-style-type: none"> ▪ Agent for VMware (Windows) ▪ Agent for Hyper-V 	1170 MB	180 MB
Agent for Linux	720 MB	130 MB
Agent for Mac	500 MB	150 MB
For on-premise deployments only		
Management Server in Windows	1.7 GB	200 MB
Management Server in Linux	0.6 GB	200 MB
Management Server and Agent for Windows	2.4 GB	360 MB
Management Server and agents on a machine running Windows, Microsoft SQL Server, Microsoft Exchange Server, and Active Directory Domain Services	3.35 GB	400 MB
Management Server and Agent for Linux	1.2 GB	340 MB

While backing up, an agent typically consumes about 350 MB of memory (measured during a 500-GB volume backup). The peak consumption may reach 2 GB, depending on the amount and type of data being processed.

Bootable media or a disk recovery with a reboot requires at least 1 GB of memory.

A management server with one registered machine consumes 200 MB of memory. Each of the newly registered machines adds about 4 MB. Thus, a server with 100 registered machines consumes approximately 600 MB above the operating system and running applications. The maximum number of registered machines is 900-1000. This limitation originates from the management server's embedded SQLite.

2.5 Supported file systems

A backup agent can back up any file system that is accessible from the operating system where the agent is installed. For example, Agent for Windows can back up and recover an ext4 file system if the corresponding driver is installed in Windows.

The following table summarizes the file systems that can be backed up and recovered. The limitations apply to both the agents and bootable media.

File system	Supported by				Limitations
	Agents	Win-PE bootable media	Linux-based bootable media	Mac bootable media	
FAT16/32	All agents	+	+	+	No limitations
NTFS		+	+	+	
ext2/ext3/ext4		+	+	-	
HFS+	Agent for Mac	-	-	+	

File system	Supported by				Limitations
	Agents	Win-PE bootable media	Linux-based bootable media	Mac bootable media	
JFS	Agent for Linux	-	+	-	Files cannot be excluded from a disk backup
ReiserFS3		-	+	-	
ReiserFS4		-	+	-	
ReFS	All agents	+	+	+	<ul style="list-style-type: none"> ▪ Files cannot be excluded from a disk backup ▪ Volumes cannot be resized during a recovery
XFS		+	+	+	
Linux swap	Agent for Linux	-	+	-	No limitations

The software automatically switches to the sector-by-sector mode when backing up drives with unrecognized or unsupported file systems. A sector-by-sector backup is possible for any file system that:

- is block-based
- spans a single disk
- has a standard MBR/GPT partitioning scheme

If the file system does not meet these requirements, the backup fails.

2.6 On-premise deployment

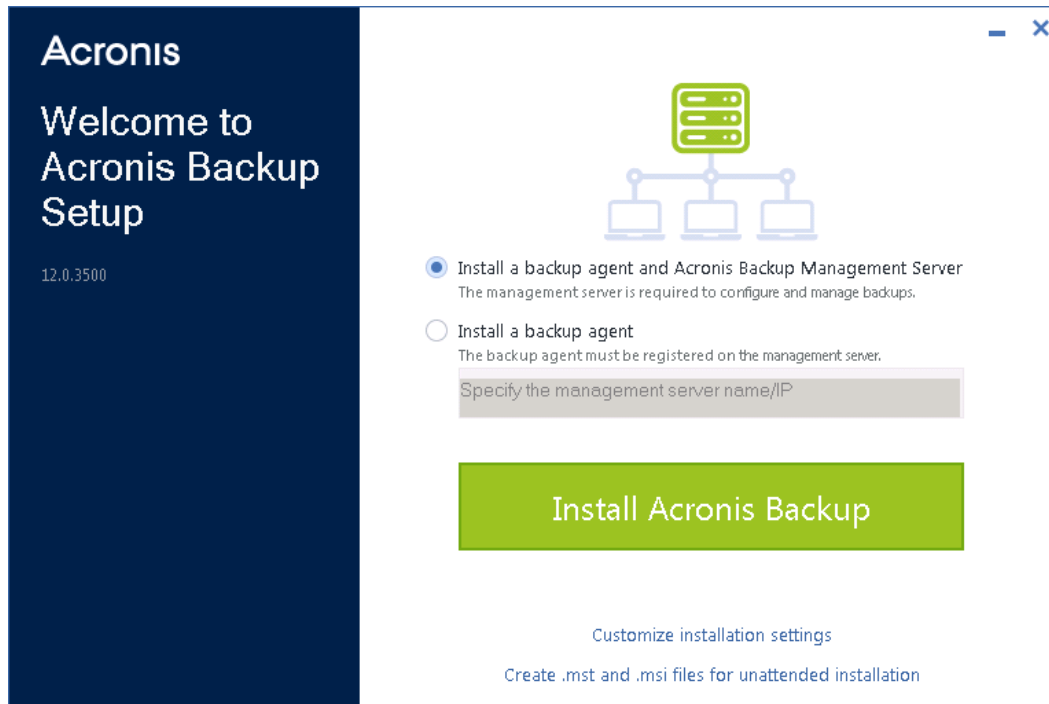
2.6.1 Installing the management server

2.6.1.1 Installation in Windows

To install the management server

1. Log on as an administrator and start the Acronis Backup setup program.
2. [Optional] To change the language the setup program is displayed in, click **Setup language**.
3. Accept the terms of the license agreement and select whether the machine will participate in the Acronis Customer Experience Program (CEP).

4. Leave the default setting **Install a backup agent and Acronis Backup Management Server**.



By default, the following components will be installed:

- Management Server
- Agent for Windows
- Other agents (Agent for Hyper-V, Agent for Exchange, Agent for SQL, and Agent for Active Directory), if the respective hypervisor or application is detected on the machine
- Bootable Media Builder
- Command-Line Tool
- Backup Monitor

You can configure the setup by clicking **Customize installation settings**.

5. Proceed with the installation.
6. After the installation completes, click **Close**. The backup console will open in your default web browser.

2.6.1.2 Installation in Linux

Preparation

1. Before installing the product on a system that does not use RPM Package Manager, such as an Ubuntu system, you need to install this manager manually; for example, by running the following command (as the root user): **apt-get install rpm**.
2. If you want to install Agent for Linux along with the management server, ensure that the necessary Linux packages (p. 13) are installed on the machine.

Installation

To install the management server

1. As the root user, run the installation file.
2. Accept the terms of the license agreement.
3. [Optional] Select the components that you want to install.

By default, the following components will be installed:

- Management Server
 - Agent for Linux
 - Bootable Media Builder
4. Specify the port that will be used by a web browser to access the management server. The default value is 9877.
 5. Specify the port that will be used for communication between the product components. The default value is 7780.
 6. Click **Next** to proceed with the installation.
 7. After the installation completes, select **Open web console**, and then click **Exit**. The backup console will open in your default web browser.

2.6.2 Adding machines via the web interface

To start adding a machine to the management server, click **All devices > Add**.

If the management server is installed in Linux, you will be asked to select the setup program based on the type of the machine that you want to add. Once the setup program is downloaded, run it locally on that machine.

The operations described later in this section are possible if the management server is installed in Windows. In most cases, the agent will be silently deployed to the selected machine.

2.6.2.1 Adding a machine running Windows

Preparation

1. For successful installation on a remote machine running Windows XP, the option **Control panel > Folder options > View > Use simple file sharing** must be *disabled* on that machine.
For successful installation on a remote machine running Windows Vista or later, the option **Control panel > Folder options > View > Use Sharing Wizard** must be *disabled* on that machine.
2. For successful installation on a remote machine that is *not* a member of an Active Directory domain, User Account Control (UAC) must be *disabled* (p. 21).
3. File and Printer Sharing must be *enabled* on the remote machine. To access this option:
 - On a machine running Windows XP with Service Pack 2 or Windows 2003 Server: go to **Control panel > Windows Firewall > Exceptions > File and Printer Sharing**.
 - On a machine running Windows Vista, Windows Server 2008, Windows 7, or later: go to **Control panel > Windows Firewall > Network and Sharing Center > Change advanced sharing settings**.
4. Acronis Backup uses TCP ports 445 and 25001 for remote installation. Also, it uses TCP port 9876 for remote installation and for communication between the components.
Port 445 is automatically opened when you enable File and Printer Sharing. Ports 9876 and 25001 are automatically opened through Windows Firewall. If you use a different firewall, make sure that these three ports are open (added to exceptions) for both incoming and outgoing requests.
After the remote installation is complete, you can remove ports 445 and 25001 from exceptions. Port 25001 is automatically closed through Windows Firewall. Port 9876 needs to remain open.

Adding the machine

1. Click **All devices > Add**.

2. Click **Windows** or the button that corresponds to the application that you want to protect. Depending on the button you click, one of the following options is selected:
 - Agent for Windows
 - Agent for Hyper-V
 - Agent for SQL + Agent for Windows
 - Agent for Exchange + Agent for Windows
 - Agent for Active Directory + Agent for Windows
3. Specify the host name or IP address of the machine, and the credentials of an account with administrative privileges on that machine.
4. Click **Add**.

Requirements on User Account Control (UAC)

On a machine that is running Windows Vista or later and is not a member of an Active Directory domain, centralized management operations (including remote installation) require that UAC be disabled.

To disable UAC

Do one of the following depending on the operating system:

- **In a Windows operating system prior to Windows 8:**
Go to **Control panel > View by: Small icons > User Accounts > Change User Account Control Settings**, and then move the slider to **Never notify**. Then, restart the machine.
- **In any Windows operating system:**
 1. Open Registry Editor.
 2. Locate the following registry key:
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System
 3. For the **EnableLUA** value, change the setting to **0**.
 4. Restart the machine.

2.6.2.2 Adding a machine running Linux

1. Click **All devices > Add**.
2. Click **Linux**. This will download the installation file.
3. On the machine that you want to protect, run the setup program locally (p. 24).

2.6.2.3 Adding a machine running OS X

1. Click **All devices > Add**.
2. Click **Mac**. This will download the installation file.
3. On the machine that you want to protect, run the setup program locally (p. 25).

2.6.2.4 Adding a vCenter or an ESXi host

There are three methods of adding a vCenter or a stand-alone ESXi host to the management server:

- **Deploying Agent for VMware (Virtual Appliance) (p. 22)**
This method is recommended in most cases. The virtual appliance will be automatically deployed to every host managed by the vCenter you specify. You can select the hosts and customize the virtual appliance settings.

- **Installing Agent for VMware (Windows) (p. 22)**
You may want to install Agent for VMware on a physical machine running Windows for the purpose of an offloaded or LAN-free backup. The agent will be automatically deployed to the machine you specify.
 - **Offloaded backup**
Use if your production ESXi hosts are so heavily loaded that running the virtual appliances is not desirable.
 - **LAN-free backup**
If your ESXi uses a SAN attached storage, install the agent on a machine connected to the same SAN. The agent will back up the virtual machines directly from the storage rather than via the ESXi host and LAN.
- **Registering an already installed Agent for VMware (p. 23)**
Use this method if you installed Agent for VMware (Windows) manually, deployed Agent for VMware (Virtual Appliance) from an OVF template (p. 25), or had to re-install the management server.

Deploying Agent for VMware (Virtual Appliance) via the web interface

1. Click **All devices > Add**.
2. Click **VMware ESXi**.
3. Select **Deploy as a virtual appliance to each host of a vCenter**.
4. Specify the address and access credentials for the vCenter Server or stand-alone ESXi host. We recommend using an account that has the **Administrator** role assigned. Otherwise, provide an account with the necessary privileges (p. 112) on the vCenter Server or ESXi.
5. [Optional] Click **Settings** to customize the deployment settings:
 - ESXi hosts that you want to deploy the agent to (only if a vCenter Server was specified in the previous step).
 - The virtual appliance name.
 - The datastore where the appliance will be located.
 - The resource pool or vApp that will contain the appliance.
 - The network that the virtual appliance's network adapter will be connected to.
 - Network settings of the virtual appliance. You can choose DHCP auto configuration or specify the values manually, including a static IP address.
6. Click **Deploy**.

Installing Agent for VMware (Windows)

Preparation

Follow the preparatory steps described in the "Adding a machine running Windows" (p. 20) section.

Installation

1. Click **All devices > Add**.
2. Click **VMware ESXi**.
3. Select **Remotely install on a machine running Windows**.
4. Specify the host name or IP address of the machine, and the credentials of an account with administrative privileges on that machine. Click **Connect**.

5. Specify the address and credentials for the vCenter Server or stand-alone ESXi host, and then click **Connect**. We recommend using an account that has the **Administrator** role assigned. Otherwise, provide an account with the necessary privileges (p. 112) on the vCenter Server or ESXi.
6. Click **Install** to install the agent.

Registering an already installed Agent for VMware

This section describes registering Agent for VMware via the web interface.

Alternative registration methods:

- You can register Agent for VMware (Virtual Appliance) by specifying the management server in the virtual appliance UI. See step 3 under "Configuring the virtual appliance" in the "Deploying Agent for VMware (Virtual Appliance) from an OVF template" (p. 25) section.
- Agent for VMware (Windows) is registered during its local installation (p. 23).

To register Agent for VMware

1. Click **All devices > Add**.
2. Click **VMware ESXi**.
3. Select **Register an already installed agent**.
4. If you register *Agent for VMware (Windows)*, specify the host name or IP address of the machine where the agent is installed, and credentials of an account with administrative privileges on that machine.

If you register *Agent for VMware (Virtual Appliance)*, specify the host name or IP address of the virtual appliance, and credentials for the vCenter Server or the stand-alone ESXi host where the appliance is running.

Click **Connect**.

5. Specify the host name or IP address of the vCenter Server or the ESXi host, and credentials to access it, and then click **Connect**. We recommend using an account that has the **Administrator** role assigned. Otherwise, provide an account with the necessary privileges (p. 112) on the vCenter Server or ESXi.
6. Click **Register** to register the agent.

2.6.3 Installing agents locally

2.6.3.1 Installation in Windows

To install Acronis Backup agents in Windows, except for Agent for VMware (Windows)

1. Log on as an administrator and start the Acronis Backup setup program.
2. [Optional] To change the language the setup program is displayed in, click **Setup language**.
3. Accept the terms of the license agreement and select whether the machine will participate in the Acronis Customer Experience Program (CEP).
4. Select **Install a backup agent**, and then specify the host name or IP address of the machine where the management server is installed.

By default, the following components will be installed:

- Agent for Windows
- Other agents (Agent for Hyper-V, Agent for Exchange, Agent for SQL, and Agent for Active Directory), if the respective hypervisor or application is detected on the machine

- Bootable Media Builder
- Command-Line Tool
- Backup Monitor

You can configure the setup by clicking **Customize installation settings**.

5. Proceed with the installation.
6. After the installation completes, click **Close**.

To install Agent for VMware (Windows)

1. Log on as an administrator and start the Acronis Backup setup program.
2. [Optional] To change the language the setup program is displayed in, click **Setup language**.
3. Accept the terms of the license agreement and select whether the machine will participate in the Acronis Customer Experience Program (CEP).
4. Click **Customize installation settings**.
5. Next to **What to install**, click **Change**.
6. Select the **Agent for VMware (Windows)** check box. If you do not want to install other components on this machine, clear the corresponding check boxes. Click **Done** to continue.
7. Next to **Acronis Backup Management Server**, click **Change**, and then specify the host name or IP address of the machine where the management server is installed. Click **Done** to continue.
8. [Optional] Change other installation settings.
9. Click **Install** to proceed with the installation.
10. After the installation completes, click **Close**.
11. Perform the procedure described in the "Registering an already installed Agent for VMware" (p. 23) section.

2.6.3.2 Installation in Linux

Preparation

1. Before installing the product on a system that does not use RPM Package Manager, such as an Ubuntu system, you need to install this manager manually; for example, by running the following command (as the root user): **apt-get install rpm**.
2. Ensure that the necessary Linux packages (p. 13) are installed on the machine.

Installation

To install Agent for Linux

1. As the root user, run the appropriate installation file (an .i686 or an .x86_64 file).
2. Accept the terms of the license agreement.
3. Clear the **Acronis Backup Management Server** check box, and then click **Next**.
4. Specify the host name or IP address of the machine where the management server is installed.
5. Click **Next** to proceed with the installation.
6. After the installation completes, click **Exit**.

Troubleshooting information is provided in the file:

/usr/lib/Acronis/BackupAndRecovery/HOWTO.INSTALL

2.6.3.3 Installation in OS X

To install Agent for Mac

1. Double-click the installation file (.dmg).
2. Wait while the operating system mounts the installation disk image.
3. Double-click **Install**.
4. Follow the on-screen instructions.

2.6.3.4 Deploying Agent for VMware (Virtual Appliance) from an OVF template

After the management server is installed, the virtual appliance's OVF package is located in the folder %ProgramFiles%\Acronis\ESXAppliance. The folder contains one .ovf file and two .vmdk files. Ensure that these files can be accessed from the machine running the vSphere Client.

Deploying the OVF template

1. Start the vSphere Client and log on to the vCenter Server.
2. On the **File** menu, click **Deploy OVF Template**.
3. In **Source**, specify the path to the virtual appliance's OVF package.
4. Review the **OVF Template Details** and click **Next**.
5. In **Name and Location**, type the name for the appliance or leave the default name of **AcronisESXAppliance**.
6. In **Host / Cluster**, select the ESXi host that the appliance will be deployed to.
7. [Optional] In **Resource Pool**, select the resource pool that will contain the appliance.
8. In **Storage**, leave the default datastore unless it does not have enough space for the virtual appliance. In this case, select another datastore. Skip this step if there is only one datastore on the server.
9. In **Disk Format**, select any required value. The disk format does not affect the appliance performance.
10. In **Network mapping**, select the bridged mode for the network adapter.
11. Review the summary, and then click **Finish**. After the successful deployment is reported, close the progress window.

Configuring the virtual appliance

1. Starting the virtual appliance

In the vSphere Client, display the **Inventory**, right-click the virtual appliance's name, and then select **Power > Power On**. Select the **Console** tab. On the welcome screen, click **Close**.

2. vCenter/ESX(i)

Under **Agent options**, in **vCenter/ESX(i)**, click **Change** and specify the vCenter Server name or IP address. The agent will be able to back up and recover any virtual machine managed by the vCenter Server.

If you do not use a vCenter Server, specify the name or IP address of the ESXi host whose virtual machines you want to back up and recover. Normally, backups run faster when the agent backs up virtual machines hosted on its own host.

Specify the credentials that the agent will use to connect to the vCenter Server or ESXi. We recommend using an account that has the **Administrator** role assigned. Otherwise, provide an account with the necessary privileges (p. 112) on the vCenter Server or ESXi.

You can click **Check connection** to ensure the access credentials are correct.

3. Acronis Backup Management Server

Under **Agent options**, in **Acronis Backup Management Server**, click **Change**.

Specify the host name or IP address of the machine where the management server is installed, and the access credentials for that machine.

4. Time zone

Under **Virtual machine**, in **Time zone**, click **Change**. Select the time zone of your location to ensure that the scheduled operations run at the appropriate time.

The virtual appliance is ready to work. In addition, you can change the following settings:

▪ Network settings

The agent's network connection is configured automatically by using Dynamic Host Configuration Protocol (DHCP). To change the default configuration, under **Agent options**, in **eth0**, click **Change** and specify the desired network settings.

▪ Local storages

You can attach an additional disk to the virtual appliance so the Agent for VMware can back up to this locally attached storage. This kind of backup is normally faster than a backup via LAN and it does not consume the network bandwidth.

The virtual disk size must be at least 10 GB. Add the disk by editing the settings of the virtual machine and click **Refresh**. The **Create storage** link becomes available. Click this link, select the disk, and then specify a label for it.

Be careful when adding an already existing disk. Once the storage is created, all data previously contained on this disk will be lost.

2.6.4 Managing licenses

Licensing of Acronis Backup is based on the number of the backed-up physical machines and virtualization hosts. Both subscription and perpetual licenses can be used. A subscription expiration period starts when you register it on the Acronis site.

To start using Acronis Backup, you need to add at least one license key to the management server. A license is automatically assigned to a machine when a backup plan is applied. You can also assign and revoke licenses manually.

To add a license key

1. Click **Settings > Licenses**.
2. Click **Add keys**.
3. Enter the license keys.
4. Click **Add**.
5. To activate a subscription, you must be signed in. If you entered at least one subscription key, enter the email address and password of your Acronis account, and then click **Sign in**. If you entered only perpetual keys, skip this step.
6. Click **Done**.

Tip *If you have already registered the subscription keys, the management server can import them from your Acronis account. To synchronize the subscription keys, click **Sync** and sign in.*

Managing perpetual licenses

To assign a perpetual license to a machine

1. Click **Settings > Licenses**.

2. Select a perpetual license.
The software displays the license keys that correspond to the selected license.
3. Select the key to assign.
4. Click **Assign**.
The software displays the machines that the selected key can be assigned to.
5. Select the machine, and then click **Done**.

To revoke a perpetual license from a machine

1. Click **Settings > Licenses**.
2. Select a perpetual license.
The software displays the license keys that correspond to the selected license. The machine that the key is assigned to is shown in the **Assigned to** column.
3. Select the license key to revoke.
4. Click **Revoke**.
5. Confirm your decision.
The revoked key will remain in the license keys list. It can be assigned to another machine.

Managing subscription licenses

To assign a subscription license to a machine

1. Click **Settings > Licenses**.
2. Select a subscription license.
The software displays the machines that the selected license is already assigned to.
3. Click **Assign**.
The software displays the machines that the selected license can be assigned to.
4. Select the machine, and then click **Done**.

To revoke a subscription license from a machine

1. Click **Settings > Licenses**.
2. Select a subscription license.
The software displays machines that the selected license is already assigned to.
3. Select the machine to revoke the license from.
4. Click **Revoke license**.
5. Confirm your decision.

2.7 Cloud deployment

2.7.1 Preparation

Step 1

Choose the agent, depending on what you are going to back up. For the information about the agents, refer to the "Components" (p. 8) section.

Step 2

Download the setup program. To find the download links, click **All devices > Add**.

The **Add devices** page provides web installers for each agent that is installed in Windows. A web installer is a small executable file that downloads the main setup program from the Internet and saves it as a temporary file. This file is deleted immediately after the installation.

If you want to store the setup programs locally, download a package containing all agents for installation in Windows by using the link at the bottom of the **Add devices** page. Both 32-bit and 64-bit packages are available. These packages enable you to customize the list of components to install. These packages also enable unattended installation, for example, via Group Policy. This advanced scenario is described in "Deploying agents through Group Policy" (p. 31).

Installation in Linux and OS X is performed from ordinary setup programs.

All setup programs require an Internet connection to register the machine in the backup service. If there is no Internet connection, the installation will fail.

Step 3

Before the installation, ensure that your firewalls and other components of your network security system (such as a proxy sever) allow both inbound and outbound connections through the following TCP ports:

- **443** and **8443** These ports are used for accessing the backup console, registering the agents, downloading the certificates, user authorization, and downloading files from the cloud storage.
- **7770...7800** The agents use these ports to communicate with the backup management server.
- **44445** The agents use this port for data transfer during backup and recovery.

If a proxy server is enabled in your network, refer to the "Proxy server settings" (p. 28) section to understand whether you need to configure these settings on each machine that runs a backup agent.

2.7.2 Proxy server settings

The backup agents can transfer data through an HTTP proxy server.

The agent installation requires an Internet connection. If a proxy server is configured in Windows (**Control panel > Internet Options > Connections**), the setup program reads the proxy server settings from the registry and uses them automatically. In Linux and OS X, you must specify the proxy settings before the installation.

Use the procedures below to specify the proxy settings before the agent installation or to change them at a later time.

In Linux

1. Create the file `/etc/Acronis/Global.config` and open it in a text editor.
2. Copy and paste the following lines into the file:

```
<?xml version="1.0" ?>
<registry name="Global">
  <key name="HttpProxy">
    <value name="Enabled" type="Tdword">"1"</value>
    <value name="Host" type="TString">"proxy.company.com"</value>
    <value name="Port" type="Tdword">"443"</value>
  </key>
</registry>
```

3. Replace `proxy.company.com` with your proxy server host name/IP address, and 443 with the decimal value of the port number.

4. Save the file.
5. If the backup agent is not installed yet, you can now install it. Otherwise, restart the agent by executing the following command in any directory:

```
sudo service acronis_mms restart
```

In OS X

1. Create the file **/Library/Application Support/Acronis/Registry/Global.config** and open it in a text editor, such as Text Edit.
2. Copy and paste the following lines into the file:

```
<?xml version="1.0" ?>
<registry name="Global">
  <key name="HttpProxy">
    <value name="Enabled" type="TdworD">"1"</value>
    <value name="Host" type="TString">"proxy.company.com"</value>
    <value name="Port" type="TdworD">"443"</value>
  </key>
</registry>
```

3. Replace `proxy.company.com` with your proxy server host name/IP address, and `443` with the decimal value of the port number.
4. Save the file.
5. If the backup agent is not installed yet, you can now install it. Otherwise, do the following to restart the agent:

- a. Go to **Applications > Utilities > Terminal**
- b. Run the following commands:

```
sudo launchctl stop acronis_mms
sudo launchctl start acronis_mms
```

In Windows

1. Create a new text document and open it in a text editor, such as Notepad.
2. Copy and paste the following lines into the file:

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\Global\HttpProxy]
"Enabled"=dword:00000001
"Host"="proxy.company.com"
"Port"=dword:000001bb
```

3. Replace `proxy.company.com` with your proxy server host name/IP address, and `000001bb` with the hexadecimal value of the port number. For example, `000001bb` is port 443.
4. Save the document as **proxy.reg**.
5. Run the file as an administrator.
6. Confirm that you want to edit the Windows registry.
7. If the backup agent is not installed yet, you can now install it. Otherwise, do the following to restart the agent:

- a. In the **Start** menu, click **Run**, and then type: **cmd**
- b. Click **OK**.
- c. Run the following commands:

```
net stop mms
net start mms
```

2.7.3 Installing agents

In Windows

1. Ensure that the machine is connected to the Internet.
2. Log on as an administrator and start the setup program.
3. Click **Install**.
4. Specify the credentials of the account to which the machine should be assigned.
5. Click **Show proxy settings** if you want to verify or change the proxy server host name/IP address and port. Otherwise, skip this step. If a proxy server is enabled in Windows, it is detected and used automatically.
6. [Only when installing Agent for VMware] Specify the address and access credentials for the vCenter Server or stand-alone ESXi host whose virtual machines the agent will back up. We recommend using an account that has the **Administrator** role assigned. Otherwise, provide an account with the necessary privileges (p. 112) on the vCenter Server or ESXi.
7. [Only when installing on a domain controller] Specify the user account under which the agent service will run. For security reasons, the setup program does not automatically create new accounts on a domain controller.
8. Click **Start installation**.

You can change the installation path and the account for the agent service by clicking **Customize installation settings** on the first step of the installation wizard.

In Linux

1. Ensure that the machine is connected to the Internet.
2. As the root user, run the installation file.
3. Specify the credentials of the account to which the machine should be assigned.
4. Complete the installation procedure.

Troubleshooting information is provided in the file:

`/usr/lib/Acronis/BackupAndRecovery/HOWTO.INSTALL`

In OS X

1. Ensure that the machine is connected to the Internet.
2. Double-click the installation file (.dmg).
3. Wait while the operating system mounts the installation disk image.
4. Double-click **Install**.
5. If prompted, provide administrator credentials.
6. Specify the credentials of the account to which the machine should be assigned.
7. Complete the installation procedure.

2.7.4 Activating the account

When an administrator creates an account for you, an email message is sent to your email address. The message contains the following information:

- **An account activation link.** Click the link and set the password for the account. Remember your login that is shown on the account activation page.
- **A link to the backup console login page.** Use this link to access the console in future. The login and password are the same as in the previous step.

2.8 Deploying agents through Group Policy

You can centrally install (or deploy) Agent for Windows onto machines that are members of an Active Directory domain, by using Group Policy.

In this section, you will find out how to set up a Group Policy object to deploy agents onto machines in an entire domain or in its organizational unit.

Every time a machine logs on to the domain, the resulting Group Policy object will ensure that the agent is installed and registered.

Prerequisites

Before proceeding with agent deployment, ensure that:

- You have an Active Directory domain with a domain controller running Microsoft Windows Server 2003 or later.
- You are a member of the **Domain Admins** group in the domain.
- You have downloaded the **All agents for installation in Windows** setup program. The download link is available on the **Add devices** page in the backup console.

Step 1: Creating the .mst transform and extracting the installation package

1. Log on as an administrator on any machine in the domain.
2. Create a shared folder that will contain the installation packages. Ensure that domain users can access the shared folder—for example, by leaving the default sharing settings for **Everyone**.
3. Copy the setup program to the folder you created.
4. Start the setup program.
5. Click **Create .mst and .msi files for unattended installation**.
6. If prompted, specify the credentials of the account to which the machines should be assigned.
7. Review or modify the installation settings that will be added to the .mst file.
8. Click **Generate**.

As a result, the .mst transform is generated and the .msi and .cab installation packages are extracted to the folder you created. You can now move or delete the setup program .exe file.

Step 2: Setting up the Group Policy objects

1. Log on to the domain controller as a domain administrator; if the domain has more than one domain controller, log on to any of them as a domain administrator.
2. If you are planning to deploy the agent in an organizational unit, ensure that the organizational unit exists in the domain. Otherwise, skip this step.
3. In the **Start** menu, point to **Administrative Tools**, and then click **Active Directory Users and Computers** (in Windows Server 2003) or **Group Policy Management** (in Windows Server 2008 and Windows Server 2012).
4. In Windows Server 2003:
 - Right-click the name of the domain or organizational unit, and then click **Properties**. In the dialog box, click the **Group Policy** tab, and then click **New**.In Windows Server 2008 and Windows Server 2012:
 - Right-click the name of the domain or organizational unit, and then click **Create a GPO in this domain, and Link it here**.
5. Name the new Group Policy object **Agent for Windows**.

6. Open the **Agent for Windows** Group Policy object for editing, as follows:
 - In Windows Server 2003, click the Group Policy object, and then click **Edit**.
 - In Windows Server 2008 and Windows Server 2012, under **Group Policy Objects**, right-click the Group Policy object, and then click **Edit**.
7. In the Group Policy object editor snap-in, expand **Computer Configuration**.
8. In Windows Server 2003 and Windows Server 2008:
 - Expand **Software Settings**.In Windows Server 2012:
 - Expand **Policies > Software Settings**.
9. Right-click **Software installation**, then point to **New**, and then click **Package**.
10. Select the agent's .msi installation package in the shared folder that you previously created, and then click **Open**.
11. In the **Deploy Software** dialog box, click **Advanced**, and then click **OK**.
12. On the **Modifications** tab, click **Add**, and then select the .mst transform that you previously created.
13. Click **OK** to close the **Deploy Software** dialog box.

2.9 Updating agents

- On-premise deployment: to update the agents, first update the management server, and then repeat the agent installation locally or by using the web interface.
- Cloud deployment: the agents are updated automatically as soon as a new version is released. If an automatic update fails for any reason, use the procedure described below.

To find the agent version, select the machine, and then click **Overview**.

To update an agent in a cloud deployment

1. Click **Settings > Agents**.

The software displays the list of machines. The machines with outdated agent versions are marked with an orange exclamation mark.
2. Select the machines that you want to update the agents on. The machines must be online.
3. Click **Update agent**.

The update progress is shown in the status column for each machine.

2.10 Uninstalling the product

If you want to remove individual product components from a machine, run the setup program, choose to modify the product, and clear the selection of the components that you want to remove. The links to the setup programs are present on the **Downloads** page (click the account icon in the top-right corner > **Downloads**).

If you want to remove all of the product components from a machine, follow the steps described below.

Warning *In on-premise deployments, please do not uninstall the management server by mistake. The backup console will become unavailable. You will no longer be able to back up and recover all machines that are registered on the management server.*

In Windows

1. Log on as an administrator.
2. Go to **Control panel**, and then select **Programs and Features (Add or Remove Programs in Windows XP) > Acronis Backup > Uninstall**.
3. [Optional] Select the **Remove the logs and configuration settings** check box.
Keep this check box cleared if you are uninstalling an agent and are planning to install it again. If you select the check box, the machine may be duplicated in the backup console and the backups of the old machine may not be associated with the new machine.
4. Confirm your decision.
5. If you are planning to install the agent again, skip this step. Otherwise, in the backup console, click **Settings > Agents**, select the machine where the agent was installed, and then click **Delete**.

In Linux

1. As the root user, run `/usr/lib/Acronis/BackupAndRecovery/uninstall/uninstall`.
2. [Optional] Select the **Clean up all product traces (Remove the product's logs, tasks, vaults, and configuration settings)** check box.
Keep this check box cleared if you are uninstalling an agent and are planning to install it again. If you select the check box, the machine may be duplicated in the backup console and the backups of the old machine may not be associated with the new machine.
3. Confirm your decision.
4. If you are planning to install the agent again, skip this step. Otherwise, in the backup console, click **Settings > Agents**, select the machine where the agent was installed, and then click **Delete**.

In OS X

1. Double-click the installation file (.dmg).
2. Wait while the operating system mounts the installation disk image.
3. Inside the image, double-click **Uninstall**.
4. If prompted, provide administrator credentials.
5. Confirm your decision.
6. If you are planning to install the agent again, skip this step. Otherwise, in the backup console, click **Settings > Agents**, select the machine where the agent was installed, and then click **Delete**.

Removing Agent for VMware (Virtual Appliance)

1. Start the vSphere Client and log on to the vCenter Server.
2. If the virtual appliance (VA) is powered on, right-click it, and then click **Power > Power Off**. Confirm your decision.
3. If the VA uses a locally attached storage on a virtual disk and you want to preserve data on that disk, do the following:
 - a. Right-click the VA, and then click **Edit Settings**.
 - b. Select the disk with the storage, and then click **Remove**. Under **Removal Options**, click **Remove from virtual machine**.
 - c. Click **OK**.As a result, the disk remains in the datastore. You can attach the disk to another VA.
4. Right-click the VA, and then click **Delete from Disk**. Confirm your decision.
5. If you are planning to install the agent again, skip this step. Otherwise, in the backup console, click **Settings > Agents**, select the virtual appliance, and then click **Delete**.

3 Accessing the backup console

Enter the login page address to the web browser address bar, and then specify the user name and password.

On-premise deployment

The login page address is the IP address or name of the machine where the management server is installed.

To be able to log in, you must be a member of the **Acronis Centralized Admins** group on the machine running the management server. By default, this group includes all members of the **Administrators** group. If the management server is installed in Linux, only the root user is allowed to log in.

Cloud deployment

The login page address is <https://backup.acronis.com/>. The user name and password are those of your Acronis account.

If your account was created by the backup administrator, you need to activate the account and set the password by clicking the link in your activation email.


Changing the language

When logged in, you can change the language of the web interface by clicking the human-figure icon in the top-right corner.

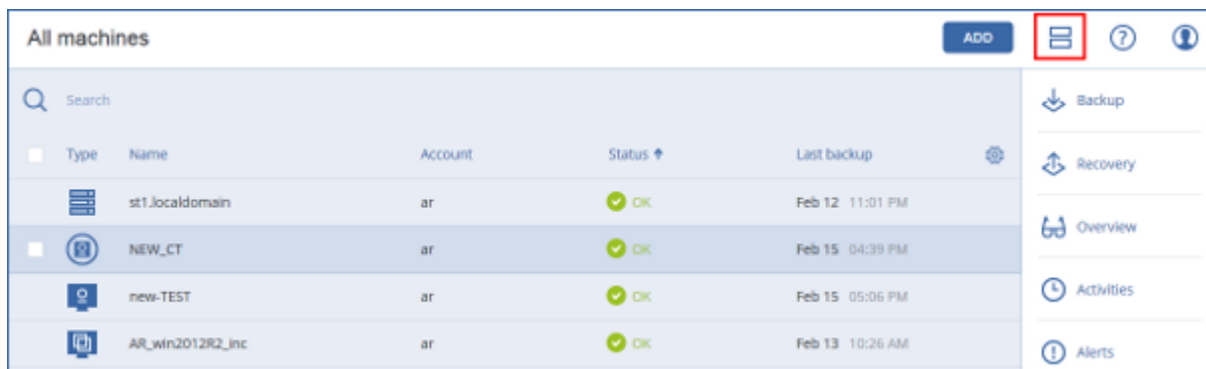
4 Backup console views

The backup console has two views: a simple view and a table view. To switch between the views, click the corresponding icon in the top right corner.

The simple view supports a small number of machines.



The table view is enabled automatically when the number of machines becomes large.



Both views provide access to the same features and operations. This document describes access to operations from the table view.

5 Backup

A backup plan is a set of rules that specify how the given data will be protected on a given machine.

A backup plan can be applied to multiple machines at the time of its creation, or later.

To create the first backup plan

1. Select the machines that you want to back up.
2. Click **Backup**.

The software displays a new backup plan template.

The screenshot shows a configuration panel for a backup plan. At the top, the title is 'Entire machine to Cloud Storage' with a pencil icon for editing and a gear icon for settings. Below the title are six rows of configuration options:


WHAT TO BACK UP	Entire machine ▼
WHERE TO BACK UP	Cloud Storage
SCHEDULE	Monday to Friday at 11:00 PM
HOW LONG TO KEEP	Monthly: 6 months Weekly: 4 weeks Daily: 7 days
ENCRYPTION	<input type="checkbox"/> Off ⓘ
APPLY	

3. [Optional] To modify the backup plan name, click the default name.
4. [Optional] To modify the plan parameters, click the corresponding section of the backup plan panel.
5. [Optional] To modify the backup options, click the gear icon.
6. Click **Apply**.

To apply an existing backup plan

1. Select the machines that you want to back up.
2. Click **Backup**. If a common backup plan is already applied to the selected machines, click **Add backup plan**.

The software displays previously created backup plans.



3. Select a backup plan to apply.
4. Click **Apply**.

5.1 Backup plan cheat sheet

The following table summarizes the available backup plan parameters. Use the table to create a backup plan that best fits your needs.

WHAT TO BACK UP	ITEMS TO BACK UP Selection methods	WHERE TO BACK UP	SCHEDULE Backup schemes (not for Cloud)	HOW LONG TO KEEP
Disks/volumes (physical machines)	Direct selection (p. 38) Policy rules (p. 38) File filters (p. 52)	Cloud (p. 41) Local folder (p. 41) Network folder (p. 41) NFS (p. 41)* Secure Zone (p. 41)**	Always incremental (Single-file) (p. 43) Always full (p. 43) Weekly full, daily incremental (p. 43) Custom (F-D-I) (p. 43)	By backup age (single rule/per backup set) (p. 45) By number of backups (p. 45) Keep indefinitely (p. 45)
Disks/volumes (virtual machines)	Policy rules (p. 38) File filters (p. 52)	Cloud (p. 41) Local folder (p. 41) Network folder (p. 41) NFS (p. 41)*		
Files (physical machines only)	Direct selection (p. 39) Policy rules (p. 39) File filters (p. 52)	Cloud (p. 41) Local folder (p. 41) Network folder (p. 41) NFS (p. 41)* Secure Zone (p. 41)**	Always full (p. 43) Weekly full, daily incremental (p. 43) Custom (F-D-I) (p. 43)	

ESXi configuration	Direct selection (p. 41)	Local folder (p. 41) Network folder (p. 41) NFS (p. 41)*		
System state	Direct selection (p. 40)	Cloud (p. 41) Local folder (p. 41) Network folder (p. 41)	Always full (p. 43) Weekly full, daily incremental (p. 43) Custom (F-I) (p. 43)	
SQL databases	Direct selection (p. 94)			
Exchange databases	Direct selection (p. 95)			

* Backup to NFS shares is not available in Windows.

** Secure Zone cannot be created on a Mac.

5.2 Selecting data to back up

5.2.1 Selecting disks/volumes

A disk-level backup contains a copy of a disk or a volume in a packaged form. You can recover individual disks, volumes, or files from a disk-level backup. A backup of an entire machine is a backup of all its disks.

There are two ways of selecting disks/volumes: directly on each machine or by using policy rules. You can exclude files from a disk backup by setting the file filters (p. 52).

Direct selection

Direct selection is available only for physical machines.

1. In **What to back up**, select **Disks/volumes**.
2. Click **Items to back up**.
3. In **Select items for backup**, select **Directly**.
4. For each of the machines included in the backup plan, select the check boxes next to the disks or volumes to back up.
5. Click **Done**.

Using policy rules

1. In **What to back up**, select **Disks/volumes**.
2. Click **Items to back up**.
3. In **Select items for backup**, select **Using policy rules**.
4. Select any of the predefined rules, type your own rules, or combine both.

The policy rules will be applied to all of the machines included in the backup plan. If no data meeting at least one of the rules is found on a machine when the backup starts, the backup will fail on that machine.

5. Click **Done**.

Rules for Windows, Linux, and OS X

- **[All volumes]** selects all volumes on machines running Windows and all mounted volumes on machines running Linux or OS X.

Rules for Windows

- **Drive letter (for example C:\)** selects the volume with the specified drive letter.

- **[Fixed Volumes (Physical machines)]** selects all volumes of physical machines, other than removable media. Fixed volumes include volumes on SCSI, ATAPI, ATA, SSA, SAS, and SATA devices, and on RAID arrays.
- **[BOOT+SYSTEM]** selects the system and boot volumes. This combination is the minimal set of data that ensures recovery of the operating system from the backup.
- **[Disk 1]** selects the first disk of the machine, including all volumes on that disk. To select another disk, type the corresponding number.

Rules for Linux

- **/dev/hda1** selects the first volume on the first IDE hard disk.
- **/dev/sda1** selects the first volume on the first SCSI hard disk.
- **/dev/md1** selects the first software RAID hard disk.

To select other basic volumes, specify **/dev/xdyN**, where:

- "x" corresponds to the disk type
- "y" corresponds to the disk number (a for the first disk, b for the second disk, and so on)
- "N" is the volume number.

To select a logical volume, specify its name along with the volume group name. For example, to back up two logical volumes, **lv_root** and **lv_bin**, both of which belong to the volume group **vg_mymachine**, specify:

```
/dev/vg_mymachine/lv_root
/dev/vg_mymachine/lv_bin
```

Rules for OS X

- **[Disk 1]** Selects the first disk of the machine, including all volumes on that disk. To select another disk, type the corresponding number.

5.2.2 Selecting files/folders

File-level backup is available only for physical machines.

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

There are two ways of selecting files: directly on each machine or by using policy rules. Either method allows you to further refine the selection by setting the file filters (p. 52).

Direct selection

1. In **What to back up**, select **Files/folders**.
2. Click **Items to back up**.
3. In **Select items for backup**, select **Directly**.
4. For each of the machines included in the backup plan:
 - a. Click **Select files and folders**.
 - b. Click **Local folder** or **Network folder**.
The share must be accessible from the selected machine.
 - c. Browse to the required files/folders or enter the path and click the arrow button. If prompted, specify the user name and password for the shared folder.

- d. Select the required files/folders.
- e. Click **Done**.

Using policy rules

1. In **What to back up**, select **Files/folders**.
2. Click **Items to back up**.
3. In **Select items for backup**, select **Using policy rules**.
4. Select any of the predefined rules, type your own rules, or combine both.

The policy rules will be applied to all of the machines included in the backup plan. If no data meeting at least one of the rules is found on a machine when the backup starts, the backup will fail on that machine.

5. Click **Done**.

Selection rules for Windows

- Full path to a file or folder, for example **D:\Work\Text.doc** or **C:\Windows**.
- Templates:
 - **[All Files]** selects all files on all volumes of the machine.
 - **[All Profiles Folder]** selects the folder where all user profiles are located (typically, **C:\Users** or **C:\Documents and Settings**).
- Environment variables:
 - **%ALLUSERSPROFILE%** selects the folder where the common data of all user profiles is located (typically, **C:\ProgramData** or **C:\Documents and Settings\All Users**).
 - **%PROGRAMFILES%** selects the Program Files folder (for example, **C:\Program Files**).
 - **%WINDIR%** selects the folder where Windows is located (for example, **C:\Windows**).

You can use other environment variables or a combination of environment variables and text. For example, to select the Java folder in the Program Files folder, type: **%PROGRAMFILES%\Java**.

Selection rules for Linux

- Full path to a file or directory. For example, to back up **file.txt** on the volume **/dev/hda3** mounted on **/home/usr/docs**, specify **/dev/hda3/file.txt** or **/home/usr/docs/file.txt**.
- **/home** selects the home directory of the common users.
- **/root** selects the root user's home directory.
- **/usr** selects the directory for all user-related programs.
- **/etc** selects the directory for system configuration files.

Selection rules for OS X

Full path to a file or directory.

Examples:

- To back up **file.txt** on your desktop, specify **/Users/<username>/Desktop/file.txt**, where **<username>** is your user name.
- To back up all users' home directories, specify **/Users**.
- To back up the directory where the applications are installed, specify **/Applications**.

5.2.3 Selecting system state

System state backup is available for machines running Windows Vista and later.

To back up system state, in **What to back up**, select **System state**.

A system state backup is comprised of the following files:

- Task scheduler configuration
- VSS Metadata Store
- Performance counter configuration information
- MSSearch Service
- Background Intelligent Transfer Service (BITS)
- The registry
- Windows Management Instrumentation (WMI)
- Component Services Class registration database

5.2.4 Selecting ESXi configuration

A backup of an ESXi host configuration enables you to recover an ESXi host to bare metal. The recovery is performed under bootable media.

The virtual machines running on the host are not included in the backup. They can be backed up and recovered separately.

A backup of an ESXi host configuration includes:

- The bootloader and boot bank partitions of the host.
- The host state (configuration of virtual networking and storage, SSL keys, server network settings, and local user information).
- Extensions and patches installed or staged on the host.
- Log files.

Prerequisites

- SSH must be enabled in the **Security Profile** of the ESXi host configuration.
- You must know the password for the 'root' account on the ESXi host.

To select an ESXi configuration

1. Go to **VMware > Host and clusters**.
2. Browse to the ESXi hosts that you want to back up.
3. Select the ESXi hosts and click **Backup**.
4. In **What to back up**, select **ESXi configuration**.
5. In **ESXi 'root' password**, specify a password for the 'root' account on each of the selected hosts or apply the same password to all of the hosts.

5.3 Selecting a destination

Click **Where to back up**, and then select one of the following:

- **Cloud storage**
Backups will be stored in the cloud data center.
- **Local folders**
If a single machine is selected, browse to a folder on the selected machine or type the folder path.

If multiple machines are selected, type the folder path. Backups will be stored in this folder on each of the selected physical machines or on the machine where the agent for virtual machines is installed. If the folder does not exist, it will be created.

- **Network folder**

This is a folder shared via SMB/CIFS.

Browse to the required shared folder or enter the path and click the arrow button. If prompted, specify the user name and password for the shared folder.

- **NFS folder** (available for machines running Linux or OS X)

Browse to the required NFS folder or enter the path in the following format:

`nfs://<host name>/<exported folder>/<subfolder>`

Then, click the arrow button.

It is not possible to back up to an NFS folder protected with a password.

- **Secure Zone** (available if it is present on each of the selected machines)

Secure Zone is a secure partition on a disk of the backed-up machine. This partition has to be created manually prior to configuring a backup. For information about how to create Secure Zone, its advantages and limitations, refer to "About Secure Zone" (p. 42).

5.3.1 About Secure Zone

Secure Zone is a secure partition on a disk of the backed-up machine. It can store backups of disks or files of this machine.

Should the disk experience a physical failure, the backups located in the Secure Zone may be lost. That's why Secure Zone should not be the only location where a backup is stored. In enterprise environments, Secure Zone can be thought of as an intermediate location used for backup when an ordinary location is temporarily unavailable or connected through a slow or busy channel.

Why use Secure Zone?

Secure Zone:

- Enables recovery of a disk to the same disk where the disk's backup resides.
- Offers a cost-effective and handy method for protecting data from software malfunction, virus attack, human error.
- Eliminates the need for a separate media or network connection to back up or recover the data. This is especially useful for roaming users.
- Can serve as a primary destination when using replication of backups.

Limitations

- Secure Zone cannot be organized on a Mac.
- Secure Zone is a partition on a basic disk. It cannot be organized on a dynamic disk or created as a logical volume (managed by LVM).
- Secure Zone is formatted with the FAT32 file system. Because FAT32 has a 4-GB file size limit, larger backups are split when saved to Secure Zone. This does not affect the recovery procedure and speed.
- Secure Zone does not support the single-file backup format (p. 123). When you change the destination to Secure Zone in a backup plan that has the **Always incremental (Single-file)** backup scheme, the scheme is changed to **Weekly full, daily incremental**.

How to create Secure Zone

1. Decide which disk you want to create Secure Zone on.
2. Start the command-line interface and type **acrocmd list disks** to see the disk's number.
3. Use the **create asz** command of the **acrocmd** utility. The command first uses the unallocated space on that disk and then, if the unallocated space is insufficient, takes free space from the specified volumes. For details, refer to "How creating Secure Zone transforms the disk" below.

Examples:

- Creating Secure Zone on disk 1 of the local machine. Secure Zone will be created with a default size that is the average between the maximum (all the unallocated space) and minimum (about 50 MB) values.

```
acrocmd create asz --disk=1
```

- Creating a password-protected Secure Zone of size 100 GB on disk 2 of the local machine. If the unallocated space is not enough, the space will be taken from the second volume of that disk.

```
acrocmd create asz --disk=2 --volume=2-2 --asz_size=100gb --password=abc12345
```

- Creating Secure Zone of size 20 GB on disk 1 of a remote machine.

```
acrocmd create asz --host=192.168.1.2 --credentials=john,pass1 --disk=1  
--asz_size=20gb
```

For the detailed description of the **create asz** command, see the command-line reference.

How creating Secure Zone transforms the disk

- Secure Zone is always created at the end of the hard disk. When calculating the final layout of the volumes, the program will first use unallocated space at the end.
- If there is no or not enough unallocated space at the end of the disk, but there is unallocated space between volumes, the volumes will be moved to add more unallocated space to the end.
- When all unallocated space is collected but it is still not enough, the program will take free space from the volumes you select, proportionally reducing the volumes' size. Resizing of locked volumes requires a reboot.
- However, there should be free space on a volume, so that the operating system and applications can operate; for example, for creating temporary files. The program will not decrease a volume where free space is or becomes less than 25 percent of the total volume size. Only when all volumes on the disk have 25 percent or less free space, will the program continue decreasing the volumes proportionally.

As is apparent from the above, specifying the maximum possible Secure Zone size is not advisable. You will end up with no free space on any volume, which might cause the operating system or applications to work unstably and even fail to start.

5.4 Schedule

The scheduling parameters depend on the backup destination.

When backing up to cloud storage

By default, backups are performed on a daily basis, Monday to Friday. You can select the time to run the backup.

If you want to change the backup frequency, move the slider, and then specify the backup schedule.

Important *The first backup is full, which means that it is the most time-consuming. All subsequent backups are incremental and take significantly less time.*

When backing up to a local or network folder

You can choose one of the predefined backup schemes or create a custom scheme. A backup scheme is a part of the backup plan that includes the backup schedule and the backup methods.

In **Backup scheme**, select one of the following:

- [Only for disk-level backups] **Always incremental (single-file)**
By default, backups are performed on a daily basis, Monday to Friday. You can select the time to run the backup.
If you want to change the backup frequency, move the slider, and then specify the backup schedule.
The backups use the new single-file backup format (p. 123).
- **Always full**
By default, backups are performed on a daily basis, Monday to Friday. You can select the time to run the backup.
If you want to change the backup frequency, move the slider, and then specify the backup schedule.
All backups are full.
- **Weekly full, Daily incremental**
By default, backups are performed on a daily basis, Monday to Friday. You can modify the days of the week and the time to run the backup.
A full backup is created once a week. All other backups are incremental. The day on which the full backup is created depends on the **Weekly backup** option (click the gear icon, then **Backup options > Weekly backup**).
- **Custom**
Specify schedules for full, differential, and incremental backups.
Differential backup is not available when backing up SQL data, Exchange data, or system state.

Additional scheduling options

With any destination, you can do the following:

- Set a date range for when the schedule is effective. Select the **Run the plan within a date range** check box, and then specify the date range.
- Disable the schedule. While the schedule is disabled, the retention rules are not applied unless a backup is started manually.
- Introduce a delay from the scheduled time. The delay value for each machine is selected randomly and ranges from zero to the maximum value you specify. You may want to use this setting when backing up multiple machines to a network location, to avoid excessive network load.
Click the gear icon, then **Backup options > Scheduling**. Select **Distribute backup start times within a time window**, and then specify the maximum delay. The delay value for each machine is determined when the backup plan is applied to the machine and remains the same until you edit the backup plan and change the maximum delay value.

Note *In cloud deployments, this option is enabled by default, with the maximum delay set to 30 minutes. In on-premise deployments, by default all backups start exactly as scheduled.*

5.5 Retention rules

1. Click **How long to keep**.
2. In **Cleanup**, choose one of the following:
 - **By backup age** (default)
Specify how long to keep backups created by the backup plan. By default, the retention rules are specified for each backup set (p. 123) separately. If you want to use a single rule for all backups, click **Switch to single rule for all backup sets**.
 - **By number of backups**
Specify the maximum number of backups to keep.
 - **Keep backups indefinitely**

Note A backup stored in a local or network folder cannot be deleted if it has dependent backups that are not subject to deletion. Such backup chains are deleted only when the lifetime of all their backups expires. This requires extra space for storing backups whose deletion is postponed. Also, the backup age and number of backups may exceed the values you specify.

5.6 Replication

If you enable backup replication, each backup will be copied to a second location immediately after creation. If earlier backups were not replicated (for example, the network connection was lost), the software also replicates all of the backups that appeared after the last successful replication.

Replicated backups do not depend on the backups remaining in the original location and vice versa. You can recover data from any backup, without access to other locations.

Usage examples

- **Reliable disaster recovery**
Store your backups both on-site (for immediate recovery) and off-site (to secure the backups from local storage failure or a natural disaster).
- **Using the cloud storage to protect data from a natural disaster**
Replicate the backups to the cloud storage by transferring only the data changes.
- **Keeping only the latest recovery points**
Delete older backups from a fast storage according to retention rules, in order to not overuse expensive storage space.

Supported locations

You can replicate a backup *from* any of these locations:

- A local folder
- A network folder

You can replicate a backup *to* any of these locations:

- A local folder
- A network folder
- The cloud storage

To enable backup replication

1. On the backup plan panel, enable the **Replicate backups** switch.

2. In **Where to replicate**, specify the replication destination, as described in "Selecting a destination" (p. 41).
3. In **How long to keep**, specify the retention rules, as described in "Retention rules" (p. 45).

5.7 Encryption

We recommend that you encrypt all backups that are stored in the cloud storage, especially if your company is subject to regulatory compliance.

Important *There is no way to recover encrypted backups if you lose or forget the password.*

Encryption in a backup plan

To enable encryption, specify the encryption settings when creating a backup plan. After a backup plan is applied, the encryption settings cannot be modified. To use different encryption settings, create a new backup plan.

To specify the encryption settings in a backup plan

1. On the backup plan panel, enable the **Encryption** switch.
2. Specify and confirm the encryption password.
3. Select one of the following encryption algorithms:
 - **AES 128** – the backups will be encrypted by using the Advanced Encryption Standard (AES) algorithm with a 128-bit key.
 - **AES 192** – the backups will be encrypted by using the AES algorithm with a 192-bit key.
 - **AES 256** – the backups will be encrypted by using the AES algorithm with a 256-bit key.
4. Click **OK**.

Encryption as a machine property

This option is intended for administrators who handle backups of multiple machines. If you need a unique encryption password for each machine or if you need to enforce encryption of backups regardless of the backup plan encryption settings, save the encryption settings on each machine individually.

Saving the encryption settings on a machine does not affect the currently applied backup plans, but will override the encryption settings of all backup plans applied later. Any backup created by these backup plans will be encrypted, even if encryption is disabled. After the settings are saved, they cannot be modified, but you can reset them as described below.

This option is available for machines running Windows or Linux. It is not supported for OS X.

This option can be used on a machine running Agent for VMware. However, be careful if you have more than one Agent for VMware connected to the same vCenter Server. It is mandatory to use the same encryption settings for all of the agents, because there is a kind of load balancing among them.

To save the encryption settings on a machine

1. Log on as an administrator (in Windows) or the root user (in Linux).
2. Run the following script:
 - In Windows: `<installation_path>\PyShell\bin\acropsh.exe -m manage_creds --set-password <encryption_password>`

Here, <installation_path> is the backup agent installation path. By default, it is %ProgramFiles%\BackupClient in cloud deployments and %ProgramFiles%\Acronis in on-premise deployments.

- In Linux: `/usr/sbin/acropsh -m manage_creds --set-password <encryption_password>`

The backups will be encrypted using the AES algorithm with a 256-bit key.

To reset the encryption settings on a machine

1. Log on as an administrator (in Windows) or root user (in Linux).
2. Run the following script:
 - In Windows: `<installation_path>\PyShell\bin\acropsh.exe -m manage_creds --reset`
Here, <installation_path> is the backup agent installation path. By default, it is %ProgramFiles%\BackupClient in cloud deployments and %ProgramFiles%\Acronis in on-premise deployments.
 - In Linux: `/usr/sbin/acropsh -m manage_creds --reset`

Important After you reset the encryption settings on a machine, the backups of this machine will fail. To continue backing up the machine, create a new backup plan.

How the encryption works

The AES cryptographic algorithm operates in the Cipher-block chaining (CBC) mode and uses a randomly generated key with a user-defined size of 128, 192 or 256 bits. The larger the key size, the longer it will take for the program to encrypt the backups and the more secure your data will be.

The encryption key is then encrypted with AES-256 using an SHA-256 hash of the password as a key. The password itself is not stored anywhere on the disk or in the backups; the password hash is used for verification purposes. With this two-level security, the backup data is protected from any unauthorized access, but recovering a lost password is not possible.

5.8 Starting a backup manually

1. Select a machine that has at least one applied backup plan.
2. Click **Backup**.
3. If more than one backup plans are applied, select the backup plan.
4. Click **Run now** on the backup plan panel.

The backup progress is shown in the **Status** column for the machine.

5.9 Backup options

To modify the backup options, click the gear icon next to the backup plan name, and then click **Backup options**.

Availability of the backup options

The set of available backup options depends on:

- The environment the agent operates in (Windows, Linux, OS X).
- The type of the data being backed up (disks, files, virtual machines, application data).
- The backup destination (the cloud storage, local or network folder).

The following table summarizes the availability of the backup options.

	Disk-level backup			File-level backup			Virtual machines		SQL and Exchange
	Windows	Linux	OS X	Windows	Linux	OS X	ESXi	Hyper-V	Windows
Backup consolidation (p. 49)	+	+	+	+	+	+	+	+	-
Backup validation (p. 49)	+	+	+	+	+	+	+	+	+
Changed block tracking (CBT) (p. 50)	+	-	-	-	-	-	+	+	-
Compression level (p. 50)	+	+	+	+	+	+	+	+	+
Email notifications (p. 50)	+	+	+	+	+	+	+	+	+
Error handling (p. 51)									
Re-attempt, if an error occurs	+	+	+	+	+	+	+	+	+
Do not show messages and dialogs while processing (silent mode)	+	+	+	+	+	+	+	+	+
Ignore bad sectors	+	+	+	+	+	+	+	+	-
Re-attempt, if an error occurs during VM snapshot creation	-	-	-	-	-	-	+	+	-
Fast incremental/differential backup (p. 52)	+	+	+	-	-	-	-	-	-
File-level backup snapshot (p. 53)	-	-	-	+	+	+	-	-	-
File-level security (p. 54)	-	-	-	+	-	-	-	-	-
File filters (p. 52)	+	+	+	+	+	+	+	+	-
Log truncation (p. 54)	-	-	-	-	-	-	+	+	SQL only
LVM snapshotting (p. 54)	-	+	-	-	-	-	-	-	-
Mount points (p. 54)	-	-	-	+	-	-	-	-	-
Multi-volume snapshot (p. 55)	+	-	-	+	-	-	-	-	-
Performance (p. 56)	+	+	+	+	+	+	+	+	+
Pre/Post commands (p. 57)	+	+	+	+	+	+	+	+	+
Pre/Post data capture commands (p. 58)	+	+	+	+	+	+	-	-	+
Scheduling (p. 60)									
Distribute start times within a time window	+	+	+	+	+	+	+	+	+
Limit the number of simultaneously running backups	-	-	-	-	-	-	+	+	-

	Disk-level backup			File-level backup			Virtual machines		SQL and Exchange
	Windows	Linux	OS X	Windows	Linux	OS X	ESXi	Hyper-V	Windows
Sector-by-sector backup (p. 60)	+	+	-	-	-	-	+	+	-
Splitting (p. 61)	+	+	+	+	+	+	+	+	+
Task failure handling (p. 61)	+	+	+	+	+	+	+	+	+
Volume Shadow Copy Service (VSS) (p. 61)	+	-	-	+	-	-	-	+	+
Volume Shadow Copy Service (VSS) for virtual machines (p. 62)	-	-	-	-	-	-	+	+	-
Weekly backup (p. 62)	+	+	+	+	+	+	+	+	+
Windows event log (p. 62)	+	-	-	+	-	-	+	+	+

5.9.1 Backup consolidation

This option is effective for the **Always full; Weekly full, Daily incremental;** and **Custom** backup schemes.

The preset is: **Disabled**.

Consolidation is the process of combining two or more subsequent backups into a single backup.

If this option is enabled, a backup that should be deleted during cleanup is consolidated with the next dependent backup (incremental or differential).

Otherwise, the backup is retained until all dependent backups become subject to deletion. This helps avoid the potentially time-consuming consolidation, but requires extra space for storing backups whose deletion is postponed. The backups' age or number can exceed the values specified in the retention rules.

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

5.9.2 Backup validation

Validation is an operation that checks the possibility of data recovery from a backup. When this option is enabled, each backup created by the backup plan is validated immediately after creation.

The preset is: **Disabled**.

Validation calculates a checksum for every data block that can be recovered from the backup. The only exception is validation of file-level backups that are located in the cloud storage. These backups are validated by checking consistency of the metadata saved in the backup.

Validation is a time-consuming process, even for an incremental or differential backup, which are small in size. This is because the operation validates not only the data physically contained in the backup, but all of the data recoverable by selecting the backup. This requires access to previously created backups.

While the successful validation means a high probability of successful recovery, it does not check all factors that influence the recovery process. If you back up the operating system, we recommend performing a test recovery under the bootable media to a spare hard drive or running a virtual machine from the backup (p. 103) in the ESXi or Hyper-V environment.

5.9.3 Changed block tracking (CBT)

This option is effective for disk-level backups of virtual machines and of physical machines running Windows.

The preset is: **Enabled**.

This option determines whether to use Changed Block Tracking (CBT) when performing an incremental or differential backup.

The CBT technology accelerates the backup process. Changes to the disk content are continuously tracked at the block level. When a backup starts, the changes can be immediately saved to the backup.

5.9.4 Compression level

The option defines the level of compression applied to the data being backed up. The available levels are: **None, Normal, High**.

The preset is: **Normal**.

A higher compression level means that the backup process takes longer, but the resulting backup occupies less space.

The optimal data compression level depends on the type of data being backed up. For example, even maximum compression will not significantly reduce the backup size if the backup contains essentially compressed files, such as .jpg, .pdf or .mp3. However, formats such as .doc or .xls will be compressed well.

5.9.5 Email notifications

The option enables you to set up email notifications about errors, warnings, and successfully completed backups.

This option is available only in on-premise deployments. In cloud deployments, the settings are configured per account when an account is created (p. 118).

The preset is: **Use the default settings**.

You can either use the default settings, or override them with custom values that will be specific for this plan only. The default settings are configured as described in "Email notifications" (p. 116).

Important When the default settings are changed, all backup plans that use the default settings are affected.

Before enabling this option, ensure that the **Email server** (p. 115) settings are configured.

To customize email notifications for a backup plan

1. Select **Customize the settings for this backup plan**.
2. In the **Recipients' email addresses** field, type the destination email address. You can enter several addresses separated by semicolons.
3. Select the types of notifications that you want to be sent. The following types are available:
 - **Errors**
 - **Warnings**
 - **Successful backups**

The subject of the email messages is based on the following template: **[subject] [machine name] [backup plan name]**. The **[subject]** placeholder will be replaced by one of the following phrases: **Backup succeeded, Backup failed, Backup succeeded with warnings**.

5.9.6 Error handling

These options enable you to specify how to handle errors that might occur during backup.

Re-attempt, if an error occurs

The preset is: **Enabled. Number of attempts: 30. Interval between attempts: 30 seconds**.

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

For example, if the backup destination on the network becomes unavailable or not reachable, the program will attempt to reach the destination every 30 seconds, but no more than 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.

Note If the cloud storage is selected as the primary or the second destination, the option value is automatically set to **Enabled. Number of attempts: 300**.

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

The preset is: **Enabled**.

With the silent mode enabled, the program will automatically handle situations requiring user interaction (except for handling bad sectors, which is defined as a separate option). If an operation cannot continue without user interaction, it will fail. Details of the operation, including errors, if any, can be found in the operation log.

Ignore bad sectors

The preset is: **Disabled**.

When this option is disabled, each time the program comes across a bad sector, the backup activity will be assigned the **Interaction required** status. In order to back up the valid information on a rapidly dying disk, enable ignoring bad sectors. The rest of the data will be backed up and you will be able to mount the resulting disk backup and extract valid files to another disk.

Re-attempt, if an error occurs during VM snapshot creation

The preset is: **Enabled**. **Number of attempts: 3**. **Interval between attempts: 5 minutes**.

When taking a virtual machine snapshot fails, the program re-attempts to perform the unsuccessful operation. You can set the time interval and the number of attempts. The attempts will be stopped as soon as the operation succeeds OR the specified number of attempts are performed, depending on which comes first.

5.9.7 Fast incremental/differential backup

This option is effective for incremental and differential disk-level backup.

The preset is: **Enabled**.

Incremental or differential backup captures only data changes. To speed up the backup process, the program determines whether a file has changed or not by the file size and the date/time when the file was last modified. Disabling this feature will make the program compare the entire file contents to those stored in the backup.

5.9.8 File filters

File filters define which files and folders to skip during the backup process.

File filters are available for both disk-level and file-level backup.

To enable file filters

1. Select the data to back up.
2. Click the gear icon next to the backup plan name, and then click **Backup options**.
3. Select **File filters**.
4. Use any of the options described below.

Exclude files matching specific criteria

There are two options that function in an inverse manner.

- **Back up only files matching the following criteria**
Example: If you select to back up the entire machine and specify **C:\File.exe** in the filter criteria, only this file will be backed up.
- **Do not back up files matching the following criteria**
Example: If you select to back up the entire machine and specify **C:\File.exe** in the filter criteria, only this file will be skipped.

It is possible to use both options simultaneously. The latter option overrides the former, i.e. if you specify **C:\File.exe** in both fields, this file will be skipped during a backup.

Criteria

▪ **Full path**

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

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

- **Name**

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

The criteria are *not* case-sensitive. For example, by specifying **C:\Temp**, you will also select **C:\TEMP**, **C:\temp**, and so on.

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

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

The question mark (?) substitutes for exactly one character in a file name. For example, the criterion **Doc?.txt** matches files such as **Doc1.txt** and **Docs.txt**, but not the files **Doc.txt** or **Doc11.txt**

Exclude hidden files and folders

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

Exclude system files and folders

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

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

5.9.9 File-level backup snapshot

This option is effective only for file-level backup.

This option defines whether to back up files one by one or by taking an instant data snapshot.

Note Files that are stored on network shares are always backed up one by one.

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

You can select one of the following:

- **Create a snapshot if it is possible**

Back up files directly if taking a snapshot is not possible.

- **Always create a snapshot**

The snapshot enables backing up of all files including files opened for exclusive access. The files will be backed up at the same point in time. Choose this setting only if these factors are critical, that is, backing up files without a snapshot does not make sense. If a snapshot cannot be taken, the backup will fail.

- **Do not create a snapshot**

Always back up files directly. Trying to back up files that are opened for exclusive access will result in a read error. Files in the backup may be not time-consistent.

5.9.10 File-level security

This option is effective only for file-level backup in Windows.

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

The preset is: **Enabled**.

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

If this option is disabled, the recovered files and folders will inherit the permissions from the folder to which they are recovered or from the disk, if recovered to the root.

Alternatively, you can disable recovery (p. 77) of the security settings. The result will be the same - the files will inherit the permissions from the parent folder.

5.9.11 Log truncation

This option is effective for backup of Microsoft SQL Server databases and for disk-level backup with enabled Microsoft SQL Server application backup.

This option defines whether the SQL Server transaction logs are truncated after a successful backup.

The preset is: **Enabled**.

When this option is enabled, a database can be recovered only to a point in time of a backup created by this software. Disable this option if you back up transaction logs by using the native backup engine of Microsoft SQL Server. You will be able to apply the transaction logs after a recovery and thus recover a database to any point in time.

5.9.12 LVM snapshotting

This option is effective only for physical machines.

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

This option defines how a snapshot of a logical volume is taken. The backup software can do this on its own or rely on Linux Logical Volume Manager (LVM).

The preset is: **By the backup software**.

- **By the backup software.** The snapshot data is kept mostly in RAM. The backup is faster and unallocated space on the volume group is not required. Therefore, we recommend changing the preset only if you are experiencing problems with backing up logical volumes.
- **By LVM.** The snapshot is stored on unallocated space of the volume group. If the unallocated space is missing, the snapshot will be taken by the backup software.

5.9.13 Mount points

This option is effective only in Windows for a file-level backup of a data source that includes mounted volumes or cluster shared volumes.

This option is effective only when you select for backup a folder that is higher in the folder hierarchy than the mount point. (A mount point is a folder on which an additional volume is logically attached.)

- If such folder (a parent folder) is selected for backup, and the **Mount points** option is enabled, all files located on the mounted volume will be included in the backup. If the **Mount points** option is disabled, the mount point in the backup will be empty.

During recovery of a parent folder, the mount point content will or will not be recovered, depending on whether the **Mount points** option for recovery (p. 77) is enabled or disabled.

- If you select the mount point directly, or select any folder within the mounted volume, the selected folders will be considered as ordinary folders. They will be backed up regardless of the state of the **Mount points** option and recovered regardless of the state of the **Mount points** option for recovery (p. 77).

The preset is: **Disabled**.

Tip. You can back up Hyper-V virtual machines residing on a cluster shared volume by backing up the required files or the entire volume with file-level backup. Just power off the virtual machines to be sure that they are backed up in a consistent state.

Example

Let's assume that the **C:\Data1** folder is a mount point for the mounted volume. The volume contains folders **Folder1** and **Folder2**. You create a backup plan for file-level backup of your data.

If you select the check box for volume C and enable the **Mount points** option, the **C:\Data1** folder in your backup will contain **Folder1** and **Folder2**. When recovering the backed-up data, be aware of proper using the **Mount points** option for recovery (p. 77).

If you select the check box for volume C, and disable the **Mount points** option, the **C:\Data1** folder in your backup will be empty.

If you select the check box for the **Data1**, **Folder1** or **Folder2** folder, the checked folders will be included in the backup as ordinary folders, regardless of the state of the **Mount points** option.

5.9.14 Multi-volume snapshot

This option is effective only for Windows operating systems.

This option applies to disk-level backup. This option also applies to file-level backup when the file-level backup is performed by taking a snapshot. (The "File-level backup snapshot" (p. 53) option determines whether a snapshot is taken during file-level backup).

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

The preset is: **Enabled**.

When this option is enabled, snapshots of all volumes being backed up are created simultaneously. Use this option to create a time-consistent backup of data spanning multiple volumes; for instance, for an Oracle database.

When this option is disabled, the volumes' snapshots are taken one after the other. As a result, if the data spans several volumes, the resulting backup may be not consistent.

5.9.15 Performance

Process priority

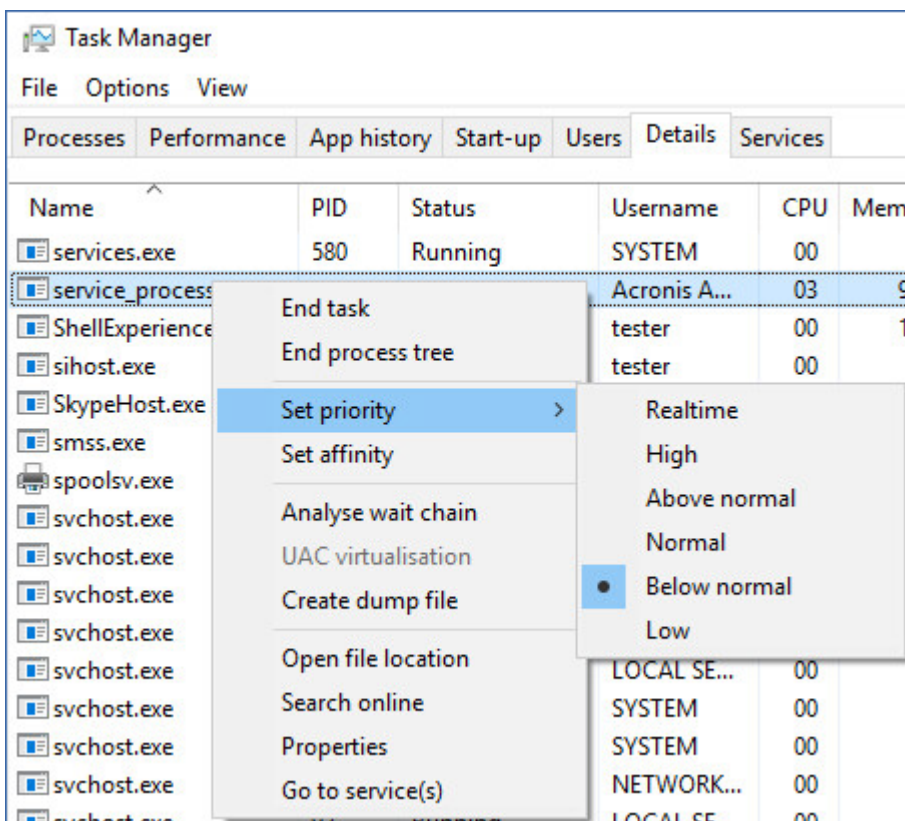
This option defines the priority of the backup process in the operating system.

The available settings are: **Low**, **Normal**, **High**.

The preset is: **Low** (in Windows, corresponds to **Below normal**).

The priority of a process running in a system determines the amount of CPU and system resources allocated to that process. Decreasing the backup priority will free more resources for other applications. Increasing the backup priority might speed up the backup process by requesting the operating system to allocate more resources like the CPU to the backup application. However, the resulting effect will depend on the overall CPU usage and other factors like disk in/out speed or network traffic.

This option sets the priority of the backup process (**service_process.exe**) in Windows and the niceness of the backup process (**service_process**) in Linux and OS X.



Output speed during backup

This option enables you to limit the hard drive writing speed (when backing up to a local folder) or the speed of transferring the backup data through the network (when backing up to a network share or to cloud storage).


The preset is: **Disabled**.

When this option is enabled, you can specify the maximum allowed output speed in KB/second.

5.9.16 Pre/Post commands

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

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



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

- Delete some temporary files from the disk before starting backup.
- Configure a third-party antivirus product to be started each time before the backup starts.
- Selectively copy backups to another location. This option may be useful because the replication configured in a backup plan copies *every* backup to subsequent locations.

The agent performs the replication *after* executing the post-backup command.

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

5.9.16.1 Pre-backup command

To specify a command/batch file to be executed before the backup process starts

1. Enable the **Execute a command before the backup** switch.
2. 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".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
	Selected	Cleared	Selected	Cleared
Fail the backup if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not back up until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the backup only after the command is successfully executed. Fail the backup if the command execution fails.	Perform the backup after the command is executed despite execution failure or success.	N/A	Perform the backup concurrently with the command execution and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

5.9.16.2 Post-backup command

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

1. Enable the **Execute a command after the backup** switch.
2. In the **Command...** field, type a command or browse to a batch file.
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field, specify the command execution arguments, if required.
5. Select the **Fail the backup if the command execution fails** check box if successful execution of the command is critical for you. The command is considered failed if its exit code is not equal to zero. If the command execution fails, the backup status will be set to **Error**.


When the check box is not selected, the command execution result does not affect the backup failure or success. You can track the command execution result by exploring the **Activities** tab.

6. Click **Done**.

5.9.17 Pre/Post data capture commands

The option enables you to define the commands to be automatically executed before and after data capture (that is, taking the data snapshot). Data capture is performed at the beginning of the backup procedure.

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



If the Volume Shadow Copy Service option (p. 61) is enabled, the commands' execution and the Microsoft VSS actions will be sequenced as follows:

"Before data capture" commands -> VSS Suspend -> Data capture -> VSS Resume -> "After data capture" commands.

By using the pre/post data capture commands, you can suspend and resume a database or application that is not compatible with VSS. Because the data capture takes seconds, the database or application idle time will be minimal.

5.9.17.1 Pre-data capture command

To specify a command/batch file to be executed before data capture

1. Enable the **Execute a command before the data capture** switch.
2. 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".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the backup if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not perform the data capture until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the data capture only after the command is successfully executed. Fail the backup if the command execution fails.	Perform the data capture after the command is executed despite execution failure or success.	N/A	Perform the data capture concurrently with the command and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

5.9.17.2 Post-data capture command

To specify a command/batch file to be executed after data capture

1. Enable the **Execute a command after the data capture** switch.
2. 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".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the backup if the command execution fails*	Selected	Cleared	Selected	Cleared
Do not back up until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Continue the backup only after the command is successfully executed.	Continue the backup after the command is executed despite command execution failure or success.	N/A	Continue the backup concurrently with the command execution and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

5.9.18 Scheduling

This option defines whether backups start as scheduled or with a delay, and how many virtual machines are backed up simultaneously.

The preset is:

- On-premise deployment: **Start all backups exactly as scheduled.**
- Cloud deployment: **Distribute backup start times within a time window. Maximum delay: 30 minutes.**

You can select one of the following:

- **Start all backups exactly as scheduled**
Backups of physical machines will start exactly as scheduled. Virtual machines will be backed up one by one.
- **Distribute start times within a time window**
Backups of physical machines will start with a delay from the scheduled time. The delay value for each machine is selected randomly and ranges from zero to the maximum value you specify. You may want to use this setting when backing up multiple machines to a network location, to avoid excessive network load. The delay value for each machine is determined when the backup plan is applied to the machine and remains the same until you edit the backup plan and change the maximum delay value.
Virtual machines will be backed up one by one.
- **Limit the number of simultaneously running backups by**
This option is available only when a backup plan is applied to multiple virtual machines. This option defines how many virtual machines an agent can back up simultaneously when executing the given backup plan.
If, according to the backup plan, an agent has to start backing up multiple machines at once, it will choose two machines. (To optimize the backup performance, the agent tries to match machines stored on different storages.) Once any of the two backups is completed, the agent chooses the third machine and so on.
You can change the number of virtual machines for an agent to simultaneously back up. The maximum value is 10.
Backups of physical machines will start exactly as scheduled.

5.9.19 Sector-by-sector backup

The option is effective only for disk-level backup.

This option defines whether an exact copy of a disk or volume on a physical level is created.

The preset is: **Disabled.**

If this option is enabled, all disk or volume's sectors will be backed up, including unallocated space and those sectors that are free of data. The resulting backup will be equal in size to the disk being backed up (if the "Compression level" (p. 50) option is set to **None**). The software automatically switches to the sector-by-sector mode when backing up drives with unrecognized or unsupported file systems.

5.9.20 Splitting

This option is effective for the **Always full; Weekly full, Daily incremental;** and **Custom** backup schemes.

This option enables you to select the method of splitting of large backups into smaller files.

The preset is: **Automatic**.

The following settings are available:

- **Automatic**
A backup will be split if it exceeds the maximum file size supported by the file system.
- **Fixed size**
Enter the desired file size or select it from the drop-down list.

5.9.21 Task failure handling

This option determines the program behavior when execution of a backup plan fails.

If this option is enabled, the program will try to execute the backup plan again. You can specify the number of attempts and the time interval between the attempts. The program stops trying as soon as an attempt completes successfully OR the specified number of attempts is performed, depending on which comes first.

The preset is: **Disabled**.

5.9.22 Volume Shadow Copy Service (VSS)

This option is effective only for Windows operating systems.

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

The preset is: **Enabled. Automatically select snapshot provider**.

You can select one of the following:

- **Automatically select snapshot provider**
Automatically select among the hardware snapshot provider, software snapshot providers, and Microsoft Software Shadow Copy provider.
- **Use Microsoft Software Shadow Copy provider**
We recommend choosing this option when backing up application servers (Microsoft Exchange Server, Microsoft SQL Server, Microsoft SharePoint, or Active Directory).

Disable this option if your database is incompatible with VSS. Snapshots are taken faster, but data consistency of the applications whose transactions are not completed at the time of taking a snapshot cannot be guaranteed. You may use Pre/Post data capture commands (p. 58) to ensure that the data is backed up in a consistent state. For instance, specify pre-data capture commands that will suspend the database and flush all caches to ensure that all transactions are completed; and specify post-data capture commands that will resume the database operations after the snapshot is taken.

Enable VSS full backup

If this option is enabled, logs of Microsoft Exchange Server and of other VSS-aware applications (except for Microsoft SQL Server) will be truncated after each successful full, incremental or differential disk-level backup.

The preset is: **Disabled**.

Leave this option disabled in the following cases:

- If you use Agent for Exchange or third-party software for backing up the Exchange Server data. This is because the log truncation will interfere with the consecutive transaction log backups.
- If you use third-party software for backing up the SQL Server data. The reason for this is that the third-party software will take the resulting disk-level backup for its "own" full backup. As a result, the next differential backup of the SQL Server data will fail. The backups will continue failing until the third-party software creates the next "own" full backup.
- If other VSS-aware applications are running on the machine and you need to keep their logs for any reason.

Enabling this option does not result in the truncation of Microsoft SQL Server logs. To truncate the SQL Server log after a backup, enable the Log truncation (p. 54) backup option.

5.9.23 Volume Shadow Copy Service (VSS) for virtual machines

This option defines whether quiesced snapshots of virtual machines are taken. To take a quiesced snapshot, the backup software applies VSS inside a virtual machine by using VMware Tools or Hyper-V Integration Services.

The preset is: **Enabled**.

If this option is enabled, transactions of all VSS-aware applications running in a virtual machine are completed before taking snapshot. If a quiesced snapshot fails after the number of re-attempts specified in the "Error handling" (p. 51) option, and application backup is disabled, a non-quiesced snapshot is taken. If application backup is enabled, the backup fails.

If this option is disabled, a non-quiesced snapshot is taken. The virtual machine will be backed up in a crash-consistent state.

5.9.24 Weekly backup

This option determines which backups are considered "weekly" in retention rules and backup schemes. A "weekly" backup is the first backup created after a week starts.

The preset is: **Monday**.

5.9.25 Windows event log

This option is effective only in Windows operating systems.

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

The preset is: **Disabled**.

6 Recovery

6.1 Recovery cheat sheet

The following table summarizes the available recovery methods. Use the table to choose a recovery method that best fits your need.

What to recover	Recovery method
Physical machine (Windows or Linux)	Using the web interface (p. 64) Using bootable media (p. 68)
Physical machine (Mac)	Using bootable media (p. 68)
Virtual machine (VMware or Hyper-V)	Using the web interface (p. 66) Using bootable media (p. 68)
ESXi configuration	Using bootable media (p. 74)
Files/Folders	Using the web interface (p. 71) Downloading files from the cloud storage (p. 72) Using bootable media (p. 72) Extracting files from local backups (p. 73)
System state	Using the web interface (p. 74)
SQL databases	Using the web interface (p. 96)
Exchange databases	Using the web interface (p. 99)
Exchange mailboxes	Using the web interface (p. 100)

Note for Mac users

Starting with 10.11 El Capitan, certain system files, folders, and processes are flagged for protection with an extended file attribute `com.apple.rootless`. This feature is called System Integrity Protection (SIP). The protected files include preinstalled applications and most of the folders in `/system`, `/bin`, `/sbin`, `/usr`.

The protected files and folders cannot be overwritten during a recovery under the operating system. If you need to overwrite the protected files, perform the recovery under bootable media.

6.2 Creating bootable media

Bootable media is a CD, DVD, USB flash drive, or other removable media that enables you to run the agent without the help of an operating system. The main purpose of bootable media is to recover an operating system that cannot start.

We highly recommend that you create and test a bootable media as soon as you start using disk-level backup. Also, it is a good practice to re-create the media after each major update of the backup agent.

You can recover either Windows or Linux by using the same media. To recover OS X, create a separate media on a machine running OS X.

To create bootable media in Windows or Linux

1. Download the bootable media ISO file. To download the file, select a machine, and then click **Recover > More ways to recover... > Download ISO image**.
2. Do any of the following:
 - Burn a CD/DVD using the ISO file.
 - Create a bootable USB flash drive by using the ISO file and one of the free tools available online.

Use ISO to USB or RUFUS if you need to boot an UEFI machine, Win32DiskImager for a BIOS machine. In Linux, using the dd utility makes sense.

- Connect the ISO file as a CD/DVD drive to the virtual machine that you want to recover.

Alternatively, you can create bootable media by using Bootable Media Builder (p. 83).

To create bootable media in OS X

1. On a machine where Agent for Mac is installed, click **Applications > Rescue Media Builder**.
2. The software displays the connected removable media. Select the one that you want to make bootable.

Warning All data on the disk will be erased.

3. Click **Create**.
4. Wait while the software creates the bootable media.

6.3 Recovering a machine

6.3.1 Physical machine

This section describes recovery of physical machines by using the web interface.

Use bootable media instead of the web interface if you need to recover:

- OS X
- Any operating system to bare metal or to an offline machine

Recovery of an operating system requires a reboot. You can choose whether to restart the machine automatically or assign it the **Interaction required** status. The recovered operating system goes online automatically.

To recover a physical machine

1. Select the backed-up machine.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Do any of the following:

- If the backup is located in the cloud storage or on a network share, click **Select machine**, select a target machine that is online, and then select a recovery point.
- Select a recovery point on the Backups tab (p. 80).
- Recover the machine as described in "Recovering disks by using bootable media" (p. 68).

4. Click **Recover > Entire machine**.

The software automatically maps the disks from the backup to the disks of the target machine.

- To recover to another physical machine, click **Target machine**, and then select a target machine that is online.

- If the disk mapping fails, recover the machine as described in "Recovering disks by using bootable media" (p. 68). The media enables you to choose disks for recovery and to map the disks manually.

RECOVER TO
Physical machine ▾

TARGET MACHINE
ABR11MMS

DISK MAPPING
Disk 1 → Disk 1

START RECOVERY

RECOVERY OPTIONS

5. Click **Start recovery**.
6. Confirm that you want to overwrite the disks with their backed-up versions. Choose whether to restart the machine automatically.

The recovery progress is shown on the **Activities** tab.

6.3.2 Physical machine to virtual


This section describes recovery of a physical machine as a virtual machine by using the web interface. This operation can be performed if at least one Agent for VMware or Agent for Hyper-V is installed and registered.

For more information about P2V migration, refer to "Machine migration" (p. 112).

To recover a physical machine as a virtual machine

1. Select the backed-up machine.
2. Click **Recovery**.
3. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Do any of the following:
 - If the backup is located in the cloud storage or on a network share, click **Select machine**, select a machine that is online, and then select a recovery point.
 - Select a recovery point on the Backups tab (p. 80).
 - Recover the machine as described in "Recovering disks by using bootable media" (p. 68).
4. Click **Recover > Entire machine**.
5. In **Recover to**, select **Virtual machine**.
6. Click **Target machine**.
 - a. Select the hypervisor (**VMware ESXi** or **Hyper-V**).
At least one Agent for VMware or Agent for Hyper-V must be installed.

- b. Select whether to recover to a new or existing machine. The new machine option is preferable as it does not require the disk configuration of the target machine to exactly match the disk configuration in the backup.
 - c. Select the host and specify the new machine name, or select an existing target machine.
 - d. Click **OK**.
7. [Optional] When recovering to a new machine, you can also do the following:
- Click **Datastore** for ESXi or **Path** for Hyper-V, and then select the datastore (storage) for the virtual machine.
 - Click **VM settings** to change the memory size, the number of processors, and the network connections of the virtual machine.



- 8. Click **Start recovery**.
 - 9. When recovering to an existing virtual machine, confirm that you want to overwrite the disks.
- The recovery progress is shown on the **Activities** tab.

6.3.3 Virtual machine

A virtual machine must be stopped during the recovery to this machine. The software stops the machine without a prompt. When the recovery is completed, you have to start the machine manually.

This behavior can be changed by using the VM power management recovery option (click **Recovery options > VM power management**).

To recover a virtual machine

- 1. Do one of the following:
 - Select a backed-up machine, click **Recovery**, and then select a recovery point.

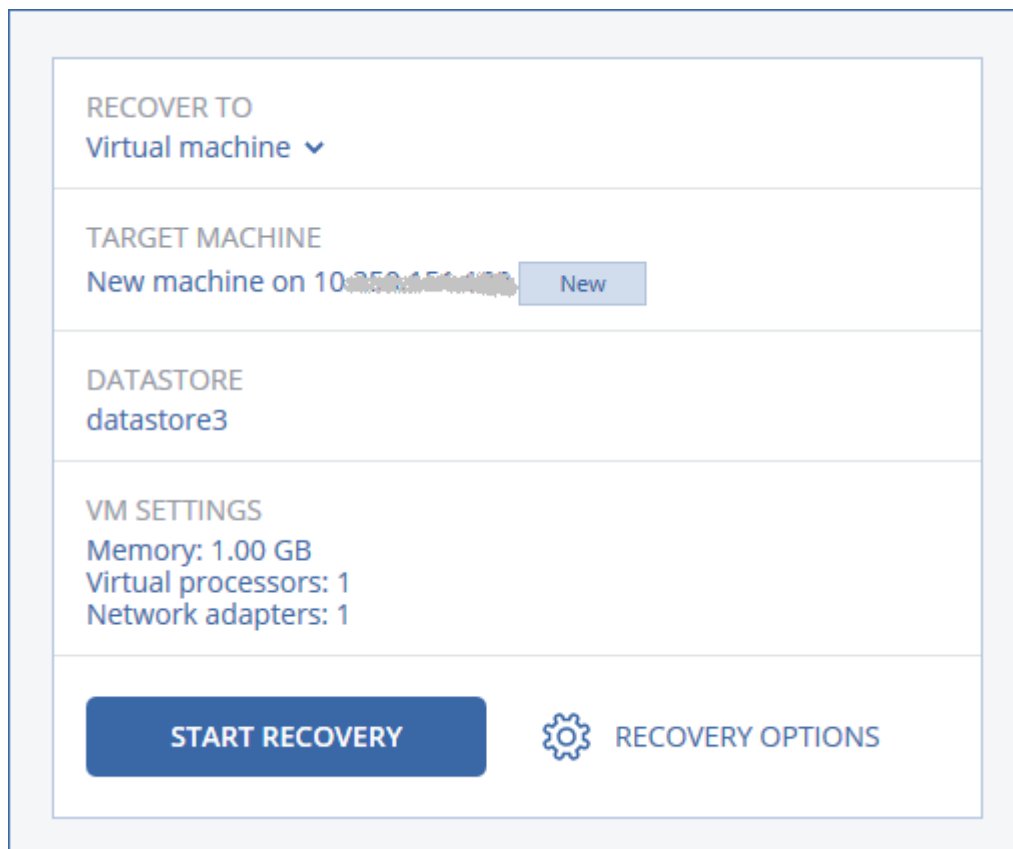
- Select a recovery point on the Backups tab (p. 80).
2. Click **Recover > Entire machine**.
 3. If you want to recover to a physical machine, select **Physical machine** in **Recover to**. Otherwise, skip this step.

Recovery to a physical machine is possible only if the disk configuration of the target machine exactly matches the disk configuration in the backup.

If this is the case, continue to step 4 in "Physical machine" (p. 64). Otherwise, we recommend that you perform the V2P migration by using bootable media (p. 68).
 4. The software automatically selects the original machine as the target machine.

To recover to another virtual machine, click **Target machine**, and then do the following:

 - a. Select the hypervisor (**VMware ESXi** or **Hyper-V**).
 - b. Select whether to recover to a new or existing machine.
 - c. Select the host and specify the new machine name, or select an existing target machine.
 - d. Click **OK**.
 5. [Optional] When recovering to a new machine, you can also do the following:
 - Click **Datastore** for ESXi or **Path** for Hyper-V, and then select the datastore (storage) for the virtual machine.
 - Click **VM settings** to change the memory size, the number of processors, and the network connections of the virtual machine.



6. Click **Start recovery**.
 7. When recovering to an existing virtual machine, confirm that you want to overwrite the disks.
- The recovery progress is shown on the **Activities** tab.

6.3.4 Recovering disks by using bootable media

For information about how to create bootable media, refer to "Creating bootable media" (p. 63).

To recover disks by using bootable media

1. Boot the target machine by using bootable media.
2. Click **Manage this machine locally** or click **Rescue Bootable Media** twice, depending on the media type you are using.
3. If a proxy server is enabled in your network, click **Tools > Proxy server**, and then specify the proxy server host name/IP address and port. Otherwise, skip this step.
4. On the welcome screen, click **Recover**.
5. Click **Select data**, and then click **Browse**.
6. Specify the backup location:
 - To recover from cloud storage, select **Cloud storage**. Enter the credentials of the account to which the backed up machine is assigned.
 - To recover from a local or a network folder, browse to the folder under **Local folders** or **Network folders**.Click **OK** to confirm your selection.
7. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
8. In **Backup contents**, select the disks that you want to recover. Click **OK** to confirm your selection.
9. Under **Where to recover**, the software automatically maps the selected disks to the target disks. If the mapping is not successful or if you are unsatisfied with the mapping result, you can re-map disks manually.

Changing disk layout may affect the operating system bootability. Please use the original machine's disk layout unless you feel fully confident of success.

10. [When recovering Linux] If the backed-up machine had logical volumes (LVM) and you want to reproduce the original LVM structure:
 - a. Ensure that the number of the target machine disks and each disk capacity are equal to or exceed those of the original machine, and then click **Apply RAID/LVM**.
 - b. Review the volume structure, and then click **Apply RAID/LVM** to create it.
11. [Optional] Click **Recovery options** to specify additional settings.
12. Click **OK** to start the recovery.

6.3.5 Using Universal Restore

The most recent operating systems remain bootable when recovered to dissimilar hardware, including the VMware or Hyper-V platforms. If a recovered operating system does not boot, use the Universal Restore tool to update the drivers and modules that are critical for the operating system startup.

Universal Restore is applicable to Windows and Linux.

To apply Universal Restore

1. Boot the machine from the bootable media.
2. Click **Apply Universal Restore**.
3. If there are multiple operating systems on the machine, choose the one to apply Universal Restore to.

4. [For Windows only] Configure the additional settings (p. 69).
5. Click **OK**.

6.3.5.1 Universal Restore in Windows

Preparation

Prepare drivers

Before applying Universal Restore to a Windows operating system, make sure that you have the drivers for the new HDD controller and the chipset. These drivers are critical to start the operating system. Use the CD or DVD supplied by the hardware vendor or download the drivers from the vendor's website. The driver files should have the *.inf extension. If you download the drivers in the *.exe, *.cab or *.zip format, extract them using a third-party application.

The best practice is to store drivers for all the hardware used in your organization in a single repository sorted by device type or by the hardware configurations. You can keep a copy of the repository on a DVD or a flash drive; pick some drivers and add them to the bootable media; create the custom bootable media with the necessary drivers (and the necessary network configuration) for each of your servers. Or, you can simply specify the path to the repository every time Universal Restore is used.

Check access to the drivers in bootable environment

Make sure you have access to the device with drivers when working under bootable media. Use WinPE-based media if the device is available in Windows but Linux-based media does not detect it.

Universal Restore settings

Automatic driver search

Specify where the program will search for the Hardware Abstraction Layer (HAL), HDD controller driver and network adapter driver(s):

- If the drivers are on a vendor's disc or other removable media, turn on the **Search removable media**.
- If the drivers are located in a networked folder or on the bootable media, specify the path to the folder by clicking **Add folder**.

In addition, Universal Restore will search the Windows default driver storage folder. Its location is determined in the registry value **DevicePath**, which can be found in the registry key **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion**. This storage folder is usually **WINDOWS/inf**.

Universal Restore will perform the recursive search in all the sub-folders of the specified folder, find the most suitable HAL and HDD controller drivers of all those available, and install them into the system. Universal Restore also searches for the network adapter driver; the path to the found driver is then transmitted by Universal Restore to the operating system. If the hardware has multiple network interface cards, Universal Restore will try to configure all the cards' drivers.

Mass storage drivers to install anyway

You need this setting if:

- The hardware has a specific mass storage controller such as RAID (especially NVIDIA RAID) or a fibre channel adapter.

- You migrated a system to a virtual machine that uses a SCSI hard drive controller. Use SCSI drivers bundled with your virtualization software or download the latest drivers versions from the software manufacturer website.
- If the automatic drivers search does not help to boot the system.

Specify the appropriate drivers by clicking **Add driver**. The drivers defined here will be installed, with appropriate warnings, even if the program finds a better driver.

Universal Restore process

After you have specified the required settings, click **OK**.

If Universal Restore cannot find a compatible driver in the specified locations, it will display a prompt about the problem device. Do one of the following:

- Add the driver to any of the previously specified locations and click **Retry**.
- If you do not remember the location, click **Ignore** to continue the process. If the result is not satisfactory, reapply Universal Restore. When configuring the operation, specify the necessary driver.

Once Windows boots, it will initialize the standard procedure for installing new hardware. The network adapter driver will be installed silently if the driver has the Microsoft Windows signature. Otherwise, Windows will ask for confirmation on whether to install the unsigned driver.

After that, you will be able to configure the network connection and specify drivers for the video adapter, USB and other devices.

6.3.5.2 Universal Restore in Linux

Universal Restore can be applied to Linux operating systems with a kernel version of 2.6.8 or later.

When Universal Restore is applied to a Linux operating system, it updates a temporary file system known as the initial RAM disk (initrd). This ensures that the operating system can boot on the new hardware.

Universal Restore adds modules for the new hardware (including device drivers) to the initial RAM disk. As a rule, it finds the necessary modules in the **/lib/modules** directory. If Universal Restore cannot find a module it needs, it records the module's file name into the log.

Universal Restore may modify the configuration of the GRUB boot loader. This may be required, for example, to ensure the system bootability when the new machine has a different volume layout than the original machine.

Universal Restore never modifies the Linux kernel.

Reverting to the original initial RAM disk

You can revert to the original initial RAM disk if necessary.

The initial RAM disk is stored on the machine in a file. Before updating the initial RAM disk for the first time, Universal Restore saves a copy of it to the same directory. The name of the copy is the name of the file, followed by the **_acronis_backup.img** suffix. This copy will not be overwritten if you run Universal Restore more than once (for example, after you have added missing drivers).

To revert to the original initial RAM disk, do any of the following:

- Rename the copy accordingly. For example, run a command similar to the following:

```
mv initrd-2.6.16.60-0.21-default_acronis_backup.img
initrd-2.6.16.60-0.21-default
```

- Specify the copy in the **initrd** line of the GRUB boot loader configuration.

6.4 Recovering files

6.4.1 Recovering files by using the web interface

1. Select the machine that originally contained the data that you want to recover.
2. Click **Recovery**.
3. Select the recovery point. Note that recovery points are filtered by location.
If the selected machine is physical and it is offline, recovery points are not displayed. Select a recovery point on the Backups tab (p. 80) or use other ways to recover:
 - Download the files from the cloud storage (p. 72)
 - Use bootable media (p. 72)
4. Click **Recover > Files/folders**.
5. Browse to the required folder or use search to obtain the list of the required files and folders.
You can use one or more wildcard characters (* and ?). For more details about using wildcards, refer to "File filters" (p. 52).
6. Select the files that you want to recover.
7. If you want to save the files as a .zip file, click **Download**, select the location to save the data to, and click **Save**. Otherwise, skip this step.
8. Click **Recover**.
In **Recover to**, you see one of the following:
 - The machine that originally contained the files that you want to recover (if it is a machine with an agent or an ESXi virtual machine).
 - The machine where Agent for Hyper-V is installed (if the files originate from a Hyper-V virtual machine). Files from a Hyper-V virtual machine cannot be recovered to the original machine.
This is the target machine for the recovery. You can select another machine, if necessary.
9. [Only when recovering to an ESXi virtual machine] Provide the credentials of a guest system user. The user must be a member of the **Administrators** group in Windows or a root user in Linux.
10. In **Path**, select the recovery destination. You can select one of the following:
 - The original location (when recovering to the original machine)
 - A local folder on the target machine
 - A network folder that is accessible from the target machine.
11. Click **Start recovery**.
12. Select one of the file overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**

The recovery progress is shown on the **Activities** tab.


6.4.2 Downloading files from the cloud storage

You can browse the cloud storage, view the contents of the backups, and download files that you need.

Limitation: Backups of system state, SQL databases, and Exchange databases cannot be browsed.

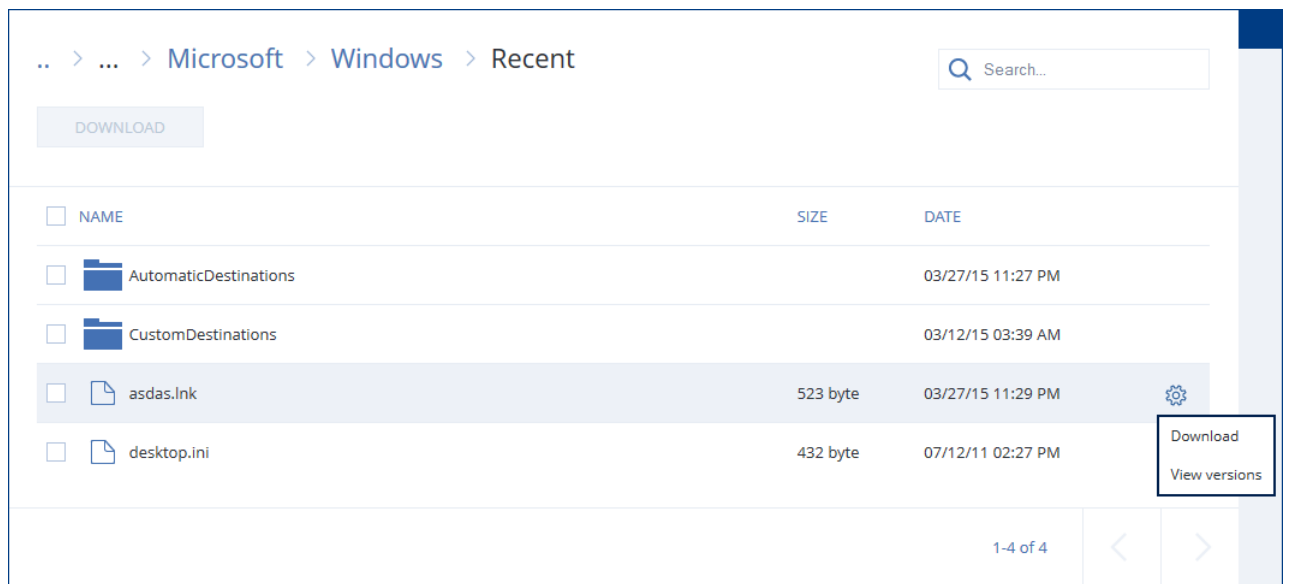
To download files from the cloud storage

1. Select a machine that was backed up.
2. Click **Recover > More ways to recover... > Download files**.
3. Enter the credentials of the account to which the backed up machine is assigned.
4. [When browsing disk-level backups] Under **Versions**, click the backup from which you want to recover the files.



[When browsing file-level backups] You can select the backup date and time in the next step, under the gear icon located to the right of the selected file. By default, files are recovered from the latest backup.

5. Browse to the required folder or use search to obtain the list of the required files and folders.



6. Select the check boxes for the items you need to recover, and then click **Download**.
If you select a single file, it will be downloaded as is. Otherwise, the selected data will be archived into a .zip file.
7. Select the location to save the data to, and then click **Save**.

6.4.3 Recovering files by using bootable media

For information about how to create bootable media, refer to "Creating bootable media" (p. 63).

To recover files by using bootable media

1. Boot the target machine by using the bootable media.
2. Click **Manage this machine locally** or click **Rescue Bootable Media** twice, depending on the media type you are using.
3. If a proxy server is enabled in your network, click **Tools > Proxy server**, and then specify the proxy server host name/IP address and port. Otherwise, skip this step.
4. On the welcome screen, click **Recover**.
5. Click **Select data**, and then click **Browse**.
6. Specify the backup location:
 - To recover from cloud storage, select **Cloud storage**. Enter the credentials of the account to which the backed up machine is assigned.
 - To recover from a local or a network folder, browse to the folder under **Local folders** or **Network folders**.Click **OK** to confirm your selection.
7. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
8. In **Backup contents**, select **Folders/files**.
9. Select the data that you want to recover. Click **OK** to confirm your selection.
10. Under **Where to recover**, specify a folder. Optionally, you can prohibit overwriting of newer versions of files or exclude some files from recovery.
11. [Optional] Click **Recovery options** to specify additional settings.
12. Click **OK** to start the recovery.

6.4.4 Extracting files from local backups

You can browse the contents of backups and extract files that you need.

Requirements

- This functionality is available only in Windows by using File Explorer.
- A backup agent must be installed on the machine from which you browse a backup.
- The backed-up file system must be one of the following: FAT16, FAT23, NTFS, ReFS, Ext2, Ext3, Ext4, XFS, or HFS+.
- The backup must be stored in a local folder, on a network share (SMB/CIFS), or in the Secure Zone.

To extract files from a backup

1. Browse to the backup location by using File Explorer.
2. Double-click the backup file. The file names are based on the following template:
<machine name> - <backup plan GUID>
3. If the backup is encrypted, enter the encryption password. Otherwise, skip this step.
File Explorer displays the recovery points.
4. Double-click the recovery point.
File Explorer displays the backed-up data.
5. Browse to the required folder.
6. Copy the required files to any folder on the file system.

6.5 Recovering system state

1. Select the machine for which you want to recover the system state.
2. Click **Recovery**.
3. Select a system state recovery point. Note that recovery points are filtered by location.
4. Click **Recover system state**.
5. Confirm that you want to overwrite the system state with its backed-up version.

The recovery progress is shown on the **Activities** tab.

6.6 Recovering ESXi configuration

To recover an ESXi configuration, you need Linux-based bootable media. For information about how to create bootable media, refer to "Creating bootable media" (p. 63).

If you are recovering an ESXi configuration to a non-original host and the original ESXi host is still connected to the vCenter Server, disconnect and remove this host from the vCenter Server to avoid unexpected issues during the recovery. If you want to keep the original host along with the recovered one, you can add it again after the recovery is complete.

The virtual machines running on the host are not included in an ESXi configuration backup. They can be backed up and recovered separately.

To recover an ESXi configuration

1. Boot the target machine by using the bootable media.
2. Click **Manage this machine locally**.
3. If the backup is located in cloud storage that is accessed via a proxy server, click **Tools > Proxy server**, and then specify the proxy server host name/IP address and port. Otherwise, skip this step.
4. On the welcome screen, click **Recover**.
5. Click **Select data**, and then click **Browse**.
6. Specify the backup location:
 - To recover from cloud storage, select **Cloud storage**. Enter the credentials of the account to which the backed-up host is assigned.
 - To recover from a local or a network folder, browse to the folder under **Local folders** or **Network folders**.Click **OK** to confirm your selection.
7. In **Show**, select **ESXi configurations**.
8. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
9. Click **OK**.
10. In **Disks to be used for new datastores**, do the following:
 - Under **Recover ESXi to**, select the disk where the host configuration will be recovered. If you are recovering the configuration to the original host, the original disk is selected by default.
 - [Optional] Under **Use for new datastore**, select the disks where new datastores will be created. Be careful because all data on the selected disks will be lost. If you want to preserve the virtual machines in the existing datastores, do not select any disks.

11. If any disks for new datastores are selected, select the datastore creation method in **How to create new datastores: Create one datastore per disk** or **Create one datastore on all selected HDDs**.
12. [Optional] In **Network mapping**, change the result of automatic mapping of the virtual switches present in the backup to the physical network adapters.
13. [Optional] Click **Recovery options** to specify additional settings.
14. Click **OK** to start the recovery.

6.7 Recovery options

To modify the recovery options, click **Recovery options** when configuring recovery.

Availability of the recovery options

The set of available recovery options depends on:

- The environment the agent that performs recovery operates in (Windows, Linux, OS X, or bootable media).
- The type of data being recovered (disks, files, virtual machines, application data).

The following table summarizes the availability of the recovery options.

	Disks			Files				Virtual machines	SQL and Exchange
	Windows	Linux	Bootable media	Windows	Linux	OS X	Bootable media	ESXi and Hyper-V	Windows
Backup validation (p. 76)	+	+	+	+	+	+	+	+	+
Date and time for files (p. 76)	-	-	-	+	+	+	+	-	-
Error handling (p. 76)	+	+	+	+	+	+	+	+	+
File exclusions (p. 77)	-	-	-	+	+	+	+	-	-
File-level security (p. 77)	-	-	-	+	+	+	+	-	-
Flashback (p. 77)	-	-	-	-	-	-	-	+	-
Full path recovery (p. 77)	-	-	-	+	+	+	+	-	-
Mount points (p. 77)	-	-	-	+	-	-	-	-	-
Performance (p. 78)	+	+	-	+	+	+	-	+	+
Pre/post commands (p. 78)	+	+	-	+	+	+	-	+	+
Restart during recovery	+	+	+	+	+	-	+	-	-

	Disks			Files				Virtual machines	SQL and Exchange
	Windows	Linux	Bootable media	Windows	Linux	OS X	Bootable media	ESXi and Hyper-V	Windows
SID changing (p. 79)	+	-	-	-	-	-	-	-	-
VM power management (p. 80)	-	-	-	-	-	-	-	+	-
Windows event log (p. 80)	+	-	-	+	-	-	-	Hyper-V only	+

6.7.1 Backup validation

This option defines whether to validate a backup to ensure that the backup is not corrupted, before data is recovered from it.

The preset is: **Disabled**.

Validation calculates a checksum for every data block saved in the backup. The only exception is validation of file-level backups that are located in the cloud storage. These backups are validated by checking consistency of the meta information saved in the backup.

Validation is a time-consuming process, even for an incremental or differential backup, which are small in size. This is because the operation validates not only the data physically contained in the backup, but all of the data recoverable by selecting the backup. This requires access to previously created backups.

6.7.2 Date and time for files

This option is effective only when recovering files.

This option defines whether to recover the files' date and time from the backup or assign the files the current date and time.

If this option is enabled, the files will be assigned the current date and time.

The preset is: **Enabled**.

6.7.3 Error handling

These options enable you to specify how to handle errors that might occur during recovery.

Re-attempt, if an error occurs

The preset is: **Enabled. Number of attempts: 30. Interval between attempts: 30 seconds.**

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

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

The preset is: **Disabled**.

With the silent mode enabled, the program will automatically handle situations requiring user interaction where possible. If an operation cannot continue without user interaction, it will fail. Details of the operation, including errors, if any, can be found in the operation log.

6.7.4 File exclusions

This option is effective only when recovering files.

The option defines which files and folders to skip during the recovery process and thus exclude from the list of recovered items.

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

6.7.5 File-level security

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

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

The preset is: **Enabled**.

If the file NTFS permissions were preserved during backup (p. 54), you can choose whether to recover the permissions or let the files inherit their NTFS permissions from the folder to which they are recovered.

6.7.6 Flashback

This option is effective when recovery to a virtual machine is performed by Agent for VMware or Agent for Hyper-V.

The Flashback technology accelerates recovery of virtual machines. If the option is enabled, only the differences between backup and the target are recovered. Data is compared at the block level.

The preset is: **Enabled**.

6.7.7 Full path recovery

This option is effective only when recovering data from a file-level backup.

If this option is enabled, the full path to the file will be re-created in the target location.

The preset is: **Disabled**.

6.7.8 Mount points

This option is effective only in Windows for recovering data from a file-level backup.

Enable this option to recover files and folders that were stored on the mounted volumes and were backed up with the enabled Mount points (p. 54) option.

The preset is: **Disabled**.

This option is effective only when you select for recovery a folder that is higher in the folder hierarchy than the mount point. If you select for recovery folders within the mount point or the mount point itself, the selected items will be recovered regardless of the **Mount points** option value.

Note Please be aware that if the volume is not mounted at the moment of recovery, the data will be recovered directly to the folder that has been the mount point at the time of backing up.

6.7.9 Performance

This option defines the priority of the recovery process in the operating system.

The available settings are: **Low, Normal, High**.

The preset is: **Normal**.

The priority of a process running in a system determines the amount of CPU and system resources allocated to that process. Decreasing the recovery priority will free more resources for other applications. Increasing the recovery priority might speed up the recovery process by requesting the operating system to allocate more resources to the application that will perform the recovery. However, the resulting effect will depend on the overall CPU usage and other factors like disk I/O speed or network traffic.

6.7.10 Pre/Post commands

The option enables you to define the commands to be automatically executed before and after the data recovery.

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

- Launch the **Checkdisk** command in order to find and fix logical file system errors, physical errors or bad sectors to be started before the recovery starts or after the recovery ends.

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

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

6.7.10.1 Pre-recovery command

To specify a command/batch file to be executed before the recovery process starts

1. Enable the **Execute a command before the recovery** switch.
2. 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".)
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field specify the command's execution arguments, if required.
5. Depending on the result you want to obtain, select the appropriate options as described in the table below.
6. Click **Done**.

Check box	Selection			
Fail the recovery if the	Selected	Cleared	Selected	Cleared

command execution fails*				
Do not recover until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the recovery only after the command is successfully executed. Fail the recovery if the command execution failed.	Perform the recovery after the command is executed despite execution failure or success.	N/A	Perform the recovery concurrently with the command execution and irrespective of the command execution result.

* A command is considered failed if its exit code is not equal to zero.

6.7.10.2 Post-recovery command

To specify a command/executable file to be executed after the recovery is completed

1. Enable the **Execute a command after the recovery** switch.
2. In the **Command...** field, type a command or browse to a batch file.
3. In the **Working directory** field, specify a path to a directory where the command/batch file will be executed.
4. In the **Arguments** field, specify the command execution arguments, if required.
5. Select the **Fail the recovery if the command execution fails** check box if successful execution of the command is critical for you. The command is considered failed if its exit code is not equal to zero. If the command execution fails, the recovery status will be set to **Error**.

When the check box is not selected, the command execution result does not affect the recovery failure or success. You can track the command execution result by exploring the **Activities** tab.

6. Click **Done**.

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

6.7.11 SID changing

This option is effective when recovering Windows 8.1/Windows Server 2012 R2 or earlier.

This option is not effective when recovery to a virtual machine is performed by Agent for VMware or Agent for Hyper-V.

The preset is: **Disabled**.

The software can generate a unique security identifier (Computer SID) for the recovered operating system. You only need this option to ensure operability of third-party software that depends on Computer SID.

Microsoft does not officially support changing SID on a deployed or recovered system. So use this option at your own risk.

6.7.12 VM power management

These options are effective when recovery to a virtual machine is performed by Agent for VMware or Agent for Hyper-V.

Power off target virtual machines when starting recovery

The preset is: **Enabled**.

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 starts. Users will be disconnected from the machine and any unsaved data will be lost.

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

Power on the target virtual machine when recovery is complete

The preset is: **Disabled**.

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

6.7.13 Windows event log

This option is effective only in Windows operating systems.

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

The preset is: **Disabled**.

7 Operations with backups

7.1 The Backups tab

The **Backups** tab shows backups of all machines ever registered on the management server. This includes offline machines and machines that are no longer registered.

In a cloud deployment, users have access only to their own backups. An administrator can view backups on behalf of any account that belongs to the same group and its child groups. This account is indirectly chosen in **Machine to browse from**. The **Backups** tab shows backups of all machines ever registered under the same account as this machine is registered.

Backup locations that are used in backup plans are automatically added to the **Backups** tab. To add a custom folder (for example, a detachable USB device) to the list of backup locations, click **Browse** and specify the folder path.

To select a recovery point by using the Backups tab

1. On the **Backups** tab, select the location where the backups are stored.

The software displays all backups that your account is allowed to view in the selected location. The backups are combined in groups. The group names are based on the following template:

<machine name> - <backup plan name>

2. Select a group from which you want to recover the data.
3. [Optional] Click **Change** next to **Machine to browse from**, and then select another machine. Some backups can only be browsed by specific agents. For example, you must select a machine running Agent for SQL to browse the backups of Microsoft SQL Server databases.

Important Please be aware that the **Machine to browse from** is a default destination for recovery from a physical machine backup. After you select a recovery point and click **Recover**, double check the **Target machine** setting to ensure that you want to recover to this specific machine. To change the recovery destination, specify another machine in **Machine to browse from**.

4. Click **Show backups**.
5. Select the recovery point.

7.2 Mounting volumes from a backup

Mounting volumes from a disk-level backup lets you access the volumes as though they were physical disks. Volumes are mounted in the read-only mode.

Requirements

- This functionality is available only in Windows by using File Explorer.
- Agent for Windows must be installed on the machine that performs the mount operation.
- The backed-up file system must be supported by the Windows version that the machine is running.
- The backup must be stored in a local folder, on a network share (SMB/CIFS), or in the Secure Zone.

To mount a volume from a backup

1. Browse to the backup location by using File Explorer.
2. Double-click the backup file. The file names are based on the following template:
<machine name> - <backup plan GUID>
3. If the backup is encrypted, enter the encryption password. Otherwise, skip this step. File Explorer displays the recovery points.
4. Double-click the recovery point. File Explorer displays the backed-up volumes.

Tip Double-click a volume to browse its content. You can copy files and folders from the backup to any folder on the file system.

5. Right-click a volume to mount, and then click **Mount in read-only mode**.
6. If the backup is stored on a network share, provide access credentials. Otherwise, skip this step. The software mounts the selected volume. The first unused letter is assigned to the volume.

To unmount a volume

1. Browse to **Computer (This PC)** in Windows 8.1 and later) by using File Explorer.
2. Right-click the mounted volume.
3. Click **Unmount**. The software unmounts the selected volume.

7.3 Deleting backups

To delete backups of a machine that is online and present in the backup console

1. On the **All devices** tab, select a machine whose backups you want to delete.
2. Click **Recovery**.
3. Select the location to delete the backups from.
4. Do one of the following:
 - To delete a single backup, select the backup to delete, and then click the recycle bin icon.
 - To delete all backups in the selected location, click **Delete all**.
5. Confirm your decision.

To delete backups of any machine

1. On the **Backups** tab, select the location from which you want to delete the backups.
The software displays all backups that your account is allowed to view in the selected location. The backups are combined in groups. The group names are based on the following template:
<machine name> - <backup plan name>
2. Select a group.
3. Do one of the following:
 - To delete a single backup, click **Show backups**, select the backup to delete, and then click the recycle bin icon.
 - To delete the selected group, click **Delete**.
4. Confirm your decision.

8 Operations with backup plans

To edit a backup plan

1. If you want to edit the backup plan for all machines to which it is applied, select one of these machines. Otherwise, select the machines for which you want to edit the backup plan.
2. Click **Backup**.
3. Select the backup plan that you want to edit.
4. Click the gear icon next to the backup plan name, and then click **Edit**.
5. To modify the plan parameters, click the corresponding section of the backup plan panel.
6. Click **Save changes**.
7. To change the backup plan for all machines to which it is applied, click **Apply the changes to this backup plan**. Otherwise, click **Create a new backup plan only for the selected devices**.

To revoke a backup plan from machines

1. Select the machines that you want to revoke the backup plan from.
2. Click **Backup**.
3. If several backup plans are applied to the machines, select the backup plan that you want to revoke.
4. Click the gear icon next to the backup plan name, and then click **Revoke**.

To delete a backup plan

1. Select any machine to which the backup plan that you want to delete is applied.
2. Click **Backup**.

3. If several backup plans are applied to the machine, select the backup plan that you want to delete.
4. Click the gear icon next to the backup plan name, and then click **Delete**.
As a result, the backup plan is revoked from all of the machines and completely removed from the web interface.

9 Bootable Media Builder

Bootable Media Builder is a dedicated tool for creating bootable media. It is available in on-premise deployments only.

Bootable Media Builder is installed by default when you install the management server. You can install the media builder separately on any machine running Windows or Linux. The supported operating systems are the same as for the corresponding agents.

Why use the media builder?

The bootable media that is available for downloading in the backup console can be used only for recovery. This media is based on a Linux kernel. Unlike Windows PE, it does not allow injecting custom drivers on the fly.

- The media builder enables you to create a customized Linux-based or WinPE-based bootable media with the backup functionality.
- Apart from creating physical media or its ISO, you can upload the media to Windows Deployment Services (WDS) and use network boot.
- Finally, you can write the media directly to a flash drive, without using third-party tools.

32- or 64-bit?

Bootable Media Builder can be installed from both 32-bit and 64-bit setup programs. The bitness of the media corresponds to the bitness of the setup program. However, you can create a 32-bit WinPE-based media by using the 64-bit media builder, if you download the 32-bit plugin.

Please remember that in most cases you need a 64-bit media to boot a machine that uses Unified Extensible Firmware Interface (UEFI).

9.1 Linux-based bootable media

To create a Linux-based bootable media

1. Start the Bootable Media Builder.
2. Specify the license key. The license will not get assigned or reassigned. It determines which functionality to enable for the created media. Without the license keys, you can create media only for recovery.
3. Select **Bootable media type: Default (Linux-based media)**.
Select the way volumes and network resources will be handled—called the media style:
 - A media with Linux-style volume handling displays the volumes as, for example, hda1 and sdb2. It tries to reconstruct MD devices and logical (LVM) volumes before starting a recovery.
 - A media with Windows-style volume handling displays the volumes as, for example, C: and D:. It provides access to dynamic (LDM) volumes.
4. [Optional] Specify the parameters of the Linux kernel. Separate multiple parameters with spaces.

For example, to be able to select a display mode for the bootable agent each time the media starts, type: **vga=ask**

For a list of parameters, see Kernel parameters (p. 84).

5. Select the components to be placed on the media: the bootable agent and/or Universal Restore. Using a media with the bootable agent, you can perform backup, recovery, and disk management operations on any PC-compatible hardware, including bare metal. Universal Restore enables you to boot an operating system recovered to dissimilar hardware or to a virtual machine if the system bootability issues occur. The tool finds and installs drivers for devices that are critical for the operating system start, such as storage controllers, motherboard, or chipset.
6. [Optional] Specify the timeout interval for the boot menu plus the component that will automatically start on timeout. If not configured, the loader waits for someone to select whether to boot the operating system (if present) or the component. If you set, say, **10 sec.** for the bootable agent, the agent will launch 10 seconds after the menu is displayed. This enables unattended onsite operation when booting from WDS/RIS.
7. [Optional] Specify the remote logon settings: the user name and password to be specified in a command string if the **acrocmbd** utility is running on a different machine. If you leave these boxes empty, the command does not need to contain credentials.
8. [Optional] Specify network settings (p. 86): TCP/IP settings to be assigned to the machine network adapters.
9. [Optional] Specify a network port (p. 87): The TCP port that the bootable agent listens for incoming connection.
10. [Optional] If a proxy server is enabled in your network, specify its host name/IP address and port.
11. Select the type of media to create. You can:
 - Create CD, DVD, or other bootable media such as removable USB flash drives if the hardware BIOS allows for boot from such media.
 - Build an ISO image to burn it later on a blank disc or to connect it to a virtual machine.
 - Upload the selected components to a WDS/RIS.
12. [Optional] Add Windows system drivers to be used by Universal Restore (p. 87). This window appears if Universal Restore is added to media and media other than WDS/RIS is selected.
13. If prompted, specify the host name/IP address and credentials for WDS/RIS, or a path to the media ISO file.
14. Check your settings in the summary screen and click **Proceed**.

9.1.1 Kernel parameters

This window lets you specify one or more parameters of the Linux kernel. They will be automatically applied when the bootable media starts.

These parameters are typically used when experiencing problems while working with the bootable media. Normally, you can leave this field empty.

You can also specify any of these parameters by pressing F11 while in the boot menu.

Parameters

When specifying multiple parameters, separate them with spaces.

acpi=off

Disables Advanced Configuration and Power Interface (ACPI). You may want to use this parameter when experiencing problems with a particular hardware configuration.

noapic

Disables Advanced Programmable Interrupt Controller (APIC). You may want to use this parameter when experiencing problems with a particular hardware configuration.

vga=ask

Prompts for the video mode to be used by the bootable media's graphical user interface. Without the **vga** parameter, the video mode is detected automatically.

vga=mode_number

Specifies the video mode to be used by the bootable media's graphical user interface. The mode number is given by *mode_number* in the hexadecimal format—for example: **vga=0x318**. Screen resolution and the number of colors corresponding to a mode number may be different on different machines. We recommend using the **vga=ask** parameter first to choose a value for *mode_number*.

quiet

Disables displaying of startup messages when the Linux kernel is loading, and starts the management console after the kernel is loaded.

This parameter is implicitly specified when creating the bootable media, but you can remove this parameter while in the boot menu.

Without this parameter, all startup messages will be displayed, followed by a command prompt. To start the management console from the command prompt, run the command: **/bin/product**

nousb

Disables loading of the USB (Universal Serial Bus) subsystem.

nousb2

Disables USB 2.0 support. USB 1.1 devices still work with this parameter. This parameter allows you to use some USB drives in the USB 1.1 mode if they do not work in the USB 2.0 mode.

nodma

Disables direct memory access (DMA) for all IDE hard disk drives. Prevents the kernel from freezing on some hardware.

nofw

Disables the FireWire (IEEE1394) interface support.

nopcmcia

Disables detection of PCMCIA hardware.

nomouse

Disables mouse support.

module_name=off

Disables the module whose name is given by *module_name*. For example, to disable the use of the SATA module, specify: **sata_sis=off**

pci=bios

Forces the use of PCI BIOS instead of accessing the hardware device directly. You may want to use this parameter if the machine has a non-standard PCI host bridge.

pci=nobios

Disables the use of PCI BIOS; only direct hardware access methods will be allowed. You may want to use this parameter when the bootable media fails to start, which may be caused by the BIOS.

pci=biosirq

Uses PCI BIOS calls to get the interrupt routing table. You may want to use this parameter if the kernel is unable to allocate interrupt requests (IRQs) or discover secondary PCI buses on the motherboard.

These calls might not work properly on some machines. But this may be the only way to get the interrupt routing table.

9.1.2 Network settings

While creating bootable media, you have an option to pre-configure network connections that will be used by the bootable agent. The following parameters can be pre-configured:

- IP address
- Subnet mask
- Gateway
- DNS server
- WINS server.

Once the bootable agent starts on a machine, the configuration is applied to the machine's network interface card (NIC). If the settings have not been pre-configured, the agent uses DHCP auto configuration. You also have the ability to configure the network settings manually when the bootable agent is running on the machine.

Pre-configuring multiple network connections

You can pre-configure TCP/IP settings for up to ten network interface cards. To ensure that each NIC will be assigned the appropriate settings, create the media on the server for which the media is customized. When you select an existing NIC in the wizard window, its settings are selected for saving on the media. The MAC address of each existing NIC is also saved on the media.

You can change the settings, except for the MAC address; or configure the settings for a non-existent NIC, if need be.

Once the bootable agent starts on the server, it retrieves the list of available NICs. This list is sorted by the slots the NICs occupy: the closest to the processor on top.

The bootable agent assigns each known NIC the appropriate settings, identifying the NICs by their MAC addresses. After the NICs with known MAC addresses are configured, the remaining NICs are assigned the settings that you have made for non-existent NICs, starting from the upper non-assigned NIC.

You can customize bootable media for any machine, and not only for the machine where the media is created. To do so, configure the NICs according to their slot order on that machine: NIC1 occupies the slot closest to the processor, NIC2 is in the next slot and so on. When the bootable agent starts on that machine, it will find no NICs with known MAC addresses and will configure the NICs in the same order as you did.

Example

The bootable agent could use one of the network adapters for communication with the management console through the production network. Automatic configuration could be done for this connection. Sizeable data for recovery could be transferred through the second NIC, included in the dedicated backup network by means of static TCP/IP settings.

9.1.3 Network port

While creating bootable media, you have an option to pre-configure the network port that the bootable agent listens to for an incoming connection from the **acrocmbd** utility. The choice is available among:

- the default port
- the currently used port
- the new port (enter the port number)

If the port has not been pre-configured, the agent uses port 9876.

9.1.4 Drivers for Universal Restore

While creating bootable media, you have an option to add Windows drivers to the media. The drivers will be used by Universal Restore to boot up Windows that was migrated to dissimilar hardware.

You will be able to configure Universal Restore:

- to search the media for the drivers that best fit the target hardware
- to get the mass-storage drivers that you explicitly specify from the media. This is necessary when the target hardware has a specific mass storage controller (such as a SCSI, RAID, or Fiber Channel adapter) for the hard disk.

The drivers will be placed in the visible Drivers folder on the bootable media. The drivers are not loaded into the targetmachine RAM, therefore, the media must stay inserted or connected throughout the Universal Restore operation.

Adding drivers to bootable media is available when you are creating a removable media or its ISO or detachable media, such as a flash drive. Drivers cannot be uploaded on WDS/RIS.

The drivers can be added to the list only in groups, by adding the INF files or folders containing such files. Selecting individual drivers from the INF files is not possible, but the media builder shows the file content for your information.

To add drivers:

1. Click **Add** and browse to the INF file or a folder that contains INF files.
2. Select the INF file or the folder.
3. Click **OK**.

The drivers can be removed from the list only in groups, by removing INF files.

To remove drivers:

1. Select the INF file.
2. Click **Remove**.

9.2 WinPE-based bootable media

Bootable Media Builder provides three methods of integrating Acronis Backup with WinPE:

- Creating the PE ISO with the plug-in from scratch.
- Adding the Acronis Plug-in to a WIM file for any future purpose (manual ISO building, adding other tools to the image and so on).

Bootable Media Builder supports WinPE distributions that are based on any the following kernels:

- Windows Vista (PE 2.0)
- Windows Vista SP1 and Windows Server 2008 (PE 2.1)
- Windows 7 (PE 3.0) with or without the supplement for Windows 7 SP1 (PE 3.1)
- Windows 8 (PE 4.0)
- Windows 8.1 (PE 5.0)
- Windows 10 (PE for Windows 10)

Bootable Media Builder supports both 32-bit and 64-bit WinPE distributions. The 32-bit WinPE distributions can also work on 64-bit hardware. However, you need a 64-bit distribution to boot a machine that uses Unified Extensible Firmware Interface (UEFI).

PE images based on WinPE 4 and later require approximately 1 GB of RAM to work.

9.2.1 Preparation: WinPE 2.x and 3.x

To be able to create or modify PE 2.x or 3.x images, install Bootable Media Builder on a machine where Windows Automated Installation Kit (AIK) is installed. If you do not have a machine with AIK, prepare it as follows.

To prepare a machine with AIK

1. Download and install Windows Automated Installation Kit.
Automated Installation Kit (AIK) for Windows Vista (PE 2.0):
<http://www.microsoft.com/Downloads/details.aspx?familyid=C7D4BC6D-15F3-4284-9123-679830D629F2&displaylang=en>
Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008 (PE 2.1):
<http://www.microsoft.com/downloads/details.aspx?FamilyID=94bb6e34-d890-4932-81a5-5b50c657de08&DisplayLang=en>
Automated Installation Kit (AIK) for Windows 7 (PE 3.0):
<http://www.microsoft.com/downloads/details.aspx?familyid=696DD665-9F76-4177-A811-39C26D3B3B34&displaylang=en>
Automated Installation Kit (AIK) Supplement for Windows 7 SP1 (PE 3.1):
<http://www.microsoft.com/download/en/details.aspx?id=5188>
You can find system requirements for installation by following the above links.
2. [Optional] Burn the WAIK to DVD or copy to a flash drive.
3. Install the Microsoft .NET Framework from this kit (NETFXx86 or NETFXx64, depending on your hardware).
4. Install Microsoft Core XML (MSXML) 5.0 or 6.0 Parser from this kit.
5. Install Windows AIK from this kit.
6. Install Bootable Media Builder on the same machine.

It is recommended that you familiarize yourself with the help documentation supplied with Windows AIK. To access the documentation, select **Microsoft Windows AIK -> Documentation** from the start menu.

9.2.2 Preparation: WinPE 4.0 and later

To be able to create or modify PE 4 or later images, install Bootable Media Builder on a machine where Windows Assessment and Deployment Kit (ADK) is installed. If you do not have a machine with ADK, prepare it as follows.

To prepare a machine with ADK

1. Download the setup program of Assessment and Deployment Kit.
Assessment and Deployment Kit (ADK) for Windows 8 (PE 4.0):
<http://www.microsoft.com/en-us/download/details.aspx?id=30652>.
Assessment and Deployment Kit (ADK) for Windows 8.1 (PE 5.0):
<http://www.microsoft.com/en-US/download/details.aspx?id=39982>.
Assessment and Deployment Kit (ADK) for Windows 10 (PE for Windows 10):
<https://msdn.microsoft.com/en-us/windows/hardware/dn913721%28v=vs.8.5%29.aspx>.
You can find system requirements for installation by following the above links.
2. Install Assessment and Deployment Kit on the machine.
3. Install Bootable Media Builder on the same machine.

9.2.3 Adding Acronis Plug-in to WinPE

To add Acronis Plug-in to WinPE:

1. Start the Bootable Media Builder.
2. Specify the license keys. The license keys will not get assigned or reassigned. They determine which functionality to enable for the created media. Without the license keys, you can create media only for recovery.
3. Select **Bootable media type: Windows PE** or **Bootable media type: Windows PE (64-bit)**. A 64-bit media is required to boot a machine that uses Unified Extensible Firmware Interface (UEFI).
If you have selected **Bootable media type: Windows PE**, do the following first:
 - Click **Download the Plug-in for WinPE (32-bit)**.
 - Save the plug-in to **%PROGRAM_FILES%\Acronis\BootableComponents\WinPE32**.If you plan to recover an operating system to dissimilar hardware or to a virtual machine and want to ensure the system bootability, select the **Include the Universal Restore tool...** check box.
4. Select **Create WinPE automatically**.
The software runs the appropriate script and proceeds to the next window.
5. Select whether to enable or disable the remote connection to a machine booted from the media. If enabled, enter a user name and password to be specified in a command line if the **acrocnd** utility is running on a different machine. If you leave these boxes empty, the remote connection will be disabled.
6. Specify network settings (p. 86) for the machine network adapters or choose DHCP auto configuration.
7. [Optional] Specify the Windows drivers to be added to Windows PE.

Once you boot a machine into Windows PE, the drivers can help you access the device where the backup is located. Add 32-bit drivers if you use a 32-bit WinPE distribution or 64-bit drivers if you use a 64-bit WinPE distribution.

Also, you will be able to point to the added drivers when configuring Universal Restore for Windows. For using Universal Restore, add 32-bit or 64-bit drivers depending on whether you are planning to recover a 32-bit or a 64-bit Windows operating system.

To add the drivers:

- Click **Add** and specify the path to the necessary *.inf file for a corresponding SCSI, RAID, SATA controller, network adapter, tape drive or other device.
 - Repeat this procedure for each driver you want to be included in the resulting WinPE media.
8. Choose whether you want to create ISO or WIM image or upload the media on a server (WDS or RIS).
 9. Specify the full path to the resulting image file including the file name, or specify the server and provide the user name and password to access it.
 10. Check your settings in the summary screen and click **Proceed**.
 11. Burn the .ISO to CD or DVD using a third-party tool or copy to a flash drive.

Once a machine boots into WinPE, the agent 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
```

Do not copy and paste this example. Type the command manually, otherwise it will fail.

For more information on customizing Windows PE 2.x and 3.x, see the Windows Preinstallation Environment User's Guide (Winpe.chm). The information on customizing Windows PE 4.0 and later is available in the Microsoft TechNet Library.

10 Protecting mobile devices

To back up and recover the data on your mobile devices, use the backup app.

Supported mobile devices

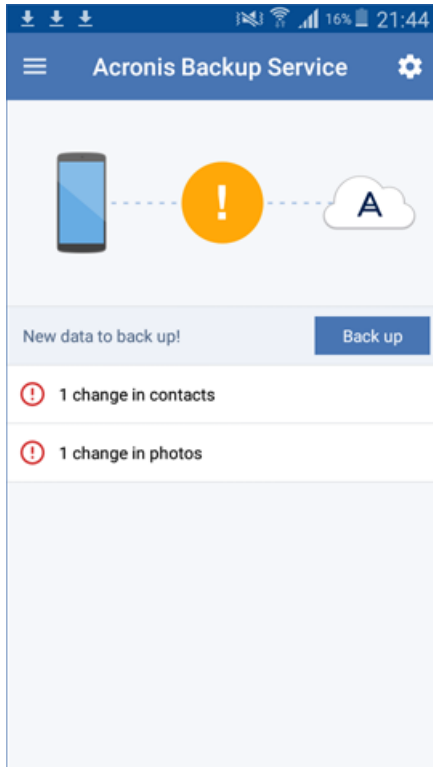
- Smartphones and tablets running Android 4.1 or later.
- iPhones, iPads, and iPods running iOS 8 or later.

What you can back up

- Contacts
- Photos
- Videos
- Calendars
- Text messages (only on Android devices)
- Reminders (only on iOS devices)

What you need to know

- You can back up the data only to the cloud storage.
- If an SD card is present, the data stored on this card is also backed up.
- Any time you open the app, you see the summary of data changes and can start a backup manually.



- The **Continuous backup** functionality is enabled by default. In this mode, the backup app checks for the data changes every six hours and runs a backup automatically if some data has changed. You can turn off continuous backup or change it to **Only when charging** in the app settings.
- You can access the backed-up data from any mobile device registered under your account. This helps you transfer the data from an old mobile device to a new one. Contacts and photos from an Android device can be recovered to an iOS device and vice versa. You can also download a photo, video, or contact to a computer by using the backup console.
- In the backup app, you can recover the data only from the latest backup. If you need to recover from older backups, use the backup console on either a tablet or a computer.
- Retention rules are not applied to backups of mobile devices. If the storage quota is exceeded, you can delete backups manually by using the backup console.

Step-by-step instructions

To get the backup app

1. In the backup console, click **All devices > Add**.
2. Select the device type.

You will get the direct link to the backup app. If you specify your email address, you will receive a message with this link.

To start backing up

1. After the backup app is installed, open it.
2. Sign in with your account.

3. Tap **Create backup**.
4. Select the data categories that you want to back up. By default, all categories are selected.
5. Tap **Start Backup**.
6. [Only on iOS devices] Allow the app access to your personal data.

To access data via the backup console

1. On a computer, open a browser and type the backup console URL.
2. Sign in with your account.
3. In **All devices**, select your mobile device name, and then click **Recovery**.
4. Select the recovery point.
5. Select the data category, and then select the data item. You can preview text messages, photos, and contacts.
6. [Optional] Click **Download** to download a photo, a video, or a contact.

For more information, refer to <https://docs.acronis.com/mobile-backup>. This help is also available in the backup app (tap **Settings** > **Help** on the app menu).

11 Protecting applications

Protecting Microsoft SQL Server and Microsoft Exchange Server

There are two methods of protecting these applications:

- **Database backup**
This is a file-level backup of the databases and the metadata associated with them. The databases can be recovered to a live application or as files.
- **Application-aware backup**
This is a disk-level backup that also collects the applications' metadata. This metadata enables browsing and recovery of the application data without recovering the entire disk or volume. The disk or volume can also be recovered as a whole. This means that a single solution and a single backup plan can be used for both disaster recovery and data protection purposes.

Protecting Microsoft SharePoint

A Microsoft SharePoint farm consists of front-end servers that run SharePoint services, database servers that run Microsoft SQL Server, and (optionally) application servers that offload some SharePoint services from the front-end servers. Some front-end and application servers may be identical to each other.

To protect an entire SharePoint farm:

- Back up all of the database servers with application-aware backup.
- Back up all of the unique front-end servers and application servers with usual disk-level backup.

The backups of all servers should be done on the same schedule.

To protect only the content, you can back up the content databases separately.

Protecting a domain controller

A machine running Active Directory Domain Services can be protected by application-aware backup. If a domain contains more than one domain controller, and you recover one of them, a nonauthoritative restore is performed and a USN rollback will not occur after the recovery.

Recovering applications

The following table summarizes the available application recovery methods.

	From a database backup	From an application-aware backup	From a disk backup
Microsoft SQL Server	Databases to a live SQL Server instance (p. 96) Databases as files (p. 96)	Entire machine (p. 64) Databases to a live SQL Server instance (p. 96) Databases as files (p. 96)	Entire machine (p. 64)
Microsoft Exchange Server	Databases to a live Exchange (p. 99) Databases as files (p. 99) Granular recovery to a live Exchange (p. 100)	Entire machine (p. 64) Databases to a live Exchange (p. 99) Databases as files (p. 99) Granular recovery to a live Exchange (p. 100)	Entire machine (p. 64)
Microsoft SharePoint database servers	Databases to a live SQL Server instance (p. 96) Databases as files (p. 96) Granular recovery by using SharePoint Explorer	Entire machine (p. 64) Databases to a live SQL Server instance (p. 96) Databases as files (p. 96) Granular recovery by using SharePoint Explorer	Entire machine (p. 64)
Microsoft SharePoint front-end web servers	-	-	Entire machine (p. 64)
Active Directory Domain Services	-	Entire machine (p. 64)	-

11.1 Prerequisites

Before configuring the application backup, ensure that the requirements listed below are met.

To check the VSS writers state, use the **vssadmin list writers** command.

Common requirements

For Microsoft SQL Server, ensure that:

- At least one Microsoft SQL Server instance is started.
- SQL Server Browser Service and TCP/IP protocol are enabled. For instructions on how to start SQL Server Browser Service, refer to: <http://msdn.microsoft.com/en-us/library/ms189093.aspx>. You can enable the TCP/IP protocol by using a similar procedure.
- The SQL writer for VSS is turned on.

For Microsoft Exchange Server, ensure that:

- The Microsoft Exchange Information Store service is started.
- Windows PowerShell is installed. For Exchange 2010 or later, the Windows PowerShell version must be at least 2.0.
- Microsoft .NET Framework is installed.
For Exchange 2007, the Microsoft .NET Framework version must be at least 2.0.
For Exchange 2010 or later, the Microsoft .NET Framework version must be at least 3.5.
- The Exchange writer for VSS is turned on.

On a domain controller, ensure that:

- The Active Directory writer for VSS is turned on.

When creating a backup plan, ensure that:

- For physical machines, the Volume Shadow Copy Service (VSS) (p. 61) backup option is enabled.
- For virtual machines, the Volume Shadow Copy Service (VSS) for virtual machines (p. 62) backup option is enabled.

Additional requirements for application-aware backups

When creating a backup plan, ensure that **Entire machine** is selected for backup.

If the applications run on virtual machines that are backed up by Agent for VMware, ensure that:

- The virtual machines being backed up meet the requirements for application-consistent quiescing listed in the following VMware knowledge base article:
<https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vddk.pg.doc%2FvddkBackupVadp.9.6.html>
- VMware Tools is installed and up-to-date on the machines.
- User Account Control (UAC) is disabled on the machines. If you do not want to disable UAC, you must provide the credentials of a built-in domain administrator (DOMAIN\Administrator) when enabling application backup.

11.2 Database backup

Before backing up databases, ensure that the requirements listed in "Prerequisites" (p. 93) are met.

Select the databases as described below, and then specify other settings of the backup plan as appropriate (p. 37).

11.2.1 Selecting SQL databases

A backup of an SQL database contains the database files (.mdf, .ndf), log files (.ldf), and other associated files. The files are backed with the help of the SQL Writer service. The service must be running at the time that the Volume Shadow Copy Service (VSS) requests a backup or recovery.

The SQL transaction logs are truncated after each successful backup. SQL log truncation can be disabled in the backup plan options (p. 54).

To select SQL databases

1. Click **Microsoft SQL**.
Machines with Agent for SQL installed are shown.
2. Browse to the data that you want to back up.
Double-click a machine to view the SQL Server instances it contains. Double-click an instance to view the databases it contains.
3. Select the data that you want to back up. You can select entire instances or individual databases.
 - If you select entire SQL Server instances, all current databases and all databases that are added to the selected instances in the future will be backed up.
 - If you select databases directly, only the selected databases will be backed up.

4. Click **Backup**. If prompted, provide credentials to access the SQL Server data. The account must be a member of the **Backup Operators** or **Administrators** group on the machine and a member of the **sysadmin** role on each of the instances that you are going to back up.

11.2.2 Selecting Exchange Server data

The following table summarizes the Microsoft Exchange Server data that you can select for backup and the minimal user rights required to back up the data.

Exchange version	Data items	User rights
2007	Storage groups	Membership in the Exchange Organization Administrators role group
2010/2013/2016	Databases	Membership in the Organization Management role group.

A full backup contains all of the selected Exchange Server data.

An incremental backup contains the changed blocks of the database files, the checkpoint files, and a small number of the log files that are more recent than the corresponding database checkpoint. Because changes to the database files are included in the backup, there is no need to back up all the transaction log records since the previous backup. Only the log that is more recent than the checkpoint needs to be replayed after a recovery. This makes for faster recovery and ensures successful database backup, even with circular logging enabled.

The transaction log files are truncated after each successful backup.

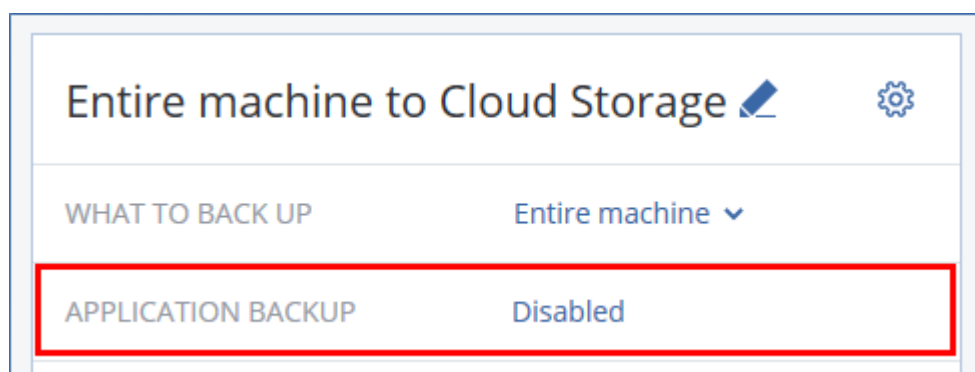
To select Exchange Server data

1. Click **Microsoft Exchange**.
Machines with Agent for Exchange installed are shown.
2. Browse to the data that you want to back up.
Double-click a machine to view the databases (storage groups) it contains.
3. Select the data that you want to back up. If prompted, provide the credentials to access the data.
4. Click **Backup**.

11.3 Application-aware backup

Application-aware disk-level backup is available for physical machines and for ESXi virtual machines.

When you back up a machine running Microsoft SQL Server, Microsoft Exchange Server, or Active Directory Domain Services, enable **Application backup** for additional protection of these applications' data.



Why use application-aware backup?

By using application-aware backup, you ensure that:

1. The applications are backed up in a consistent state and thus will be available immediately after the machine is recovered.
2. You can recover the SQL and Exchange databases, mailboxes, and mailbox items without recovering the entire machine.
3. The SQL transaction logs are truncated after each successful backup. SQL log truncation can be disabled in the backup plan options (p. 54). The Exchange transaction logs are truncated on virtual machines only. You can enable the VSS full backup option (p. 61) if you want to truncate Exchange transaction logs on a physical machine.
4. If a domain contains more than one domain controller, and you recover one of them, a nonauthoritative restore is performed and a USN rollback will not occur after the recovery.

What do I need to use application-aware backup?

On a physical machine, Agent for SQL and/or Agent for Exchange must be installed, in addition to Agent for Windows. On a virtual machine, no agent installation is required; it is presumed that the machine is backed up by Agent for VMware (Windows).

Other requirements are listed in the "Prerequisites" (p. 93) and "Required user rights" (p. 96) sections.

11.3.1 Required user rights

An application-aware backup contains metadata of VSS-aware applications that are present on the disk. To access this metadata, the agent needs an account with the appropriate rights, which are listed below. You are prompted to specify this account when enabling application backup.

- For SQL Server:
The account must be a member of the **Backup Operators** or **Administrators** group on the machine, and a member of the **sysadmin** role on each of the instances that you are going to back up.
- For Exchange Server:
Exchange 2007: The account must be a member of the **Exchange Organization Administrators** role group.
Exchange 2010 and later: The account must be a member of the **Organization Management** role group.
- For Active Directory:
The account must be a domain administrator.

11.4 Recovering SQL databases

This section describes recovery from both database backups and application-aware backups.

You can recover SQL databases to a SQL Server instance, if Agent for SQL is installed on the machine running the instance. You will need to provide credentials for an account that is a member of the **Backup Operators** or **Administrators** group on the machine and a member of the **sysadmin** role on the target instance.

Alternatively, you can recover the databases as files. This can be useful if you need to recover databases to a machine where Agent for SQL is not installed, or you need to extract data for data

mining, audit, or further processing by third-party tools. You can attach the SQL database files to a SQL Server instance, as described in "Attaching SQL Server databases" (p. 98).

If you use only Agent for VMware, recovering databases as files is the only available recovery method.

System databases are basically recovered in the same way as user databases. The peculiarities of system database recovery are described in "Recovering system databases" (p. 98).

To recover SQL databases

1. When recovering from a database backup, click **Microsoft SQL**. Otherwise, skip this step.
2. Select the machine that originally contained the data that you want to recover.
3. Click **Recovery**.
4. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Do one of the following:

- If the backup is located in the cloud storage or on a network share, click **Select machine**, select an online machine that has Agent for SQL, and then select a recovery point.
- Select a recovery point on the Backups tab (p. 80).

The machine chosen for browsing in either of the above actions becomes a target machine for the SQL databases recovery.

5. Do one of the following:
 - When recovering from a database backup, click **Recover SQL databases**.
 - When recovering from an application-aware backup, click **Recover > SQL databases**.
6. Select the data that you want to recover. Double-click an instance to view the databases it contains.
7. If you want to recover the databases as files, click **Recover as files**, select a local or a network folder to save the files to, and then click **Recover**. Otherwise, skip this step.
8. Click **Recover**.
9. By default, the databases are recovered to the original ones. If the original database does not exist, it will be recreated. You can select another machine or another SQL Server instance to recover the databases to.

To recover a database as a different one to the same instance:

- a. Click the database name.
 - b. In **Recover to**, select **New database**.
 - c. Specify the new database name.
 - d. Specify the new database path and log path. The folder you specify must not contain the original database and log files.
10. [Optional] To change the database state after recovery, click the database name, and then choose one of the following states:
 - **Ready to use (RESTORE WITH RECOVERY)** (default)
After the recovery completes, the database will be ready for use. Users will have full access to it. The software will roll back all uncommitted transactions of the recovered database that are stored in the transaction logs. You will not be able to recover additional transaction logs from the native Microsoft SQL backups.
 - **Non-operational (RESTORE WITH NORECOVERY)**
After the recovery completes, the database will be non-operational. Users will have no access to it. The software will keep all uncommitted transactions of the recovered database.

You will be able to recover additional transaction logs from the native Microsoft SQL backups and thus reach the necessary recovery point.

- **Read-only (RESTORE WITH STANDBY)**

After the recovery completes, users will have read-only access to the database. The software will undo any uncommitted transactions. However, it will save the undo actions in a temporary standby file so that the recovery effects can be reverted.

This value is primarily used to detect the point in time when a SQL Server error occurred.

11. Click **Start recovery**.

The recovery progress is shown on the **Activities** tab.

11.4.1 Recovering system databases

All system databases of an instance are recovered at once. When recovering system databases, the software automatically restarts the destination instance in the single-user mode. After the recovery completes, the software restarts the instance and recovers other databases (if any).

Other things to consider when recovering system databases:

- System databases can only be recovered to an instance of the same version as the original instance.
- System databases are always recovered in the "ready to use" state.

Recovering the master database

System databases include the **master** database. The **master** database records information about all databases of the instance. Hence, the **master** database in a backup contains information about databases which existed in the instance at the time of the backup. After recovering the **master** database, you may need to do the following:

- Databases that have appeared in the instance after the backup was done are not visible by the instance. To bring these databases back to production, attach them to the instance manually by using SQL Server Management Studio.
- Databases that have been deleted after the backup was done are displayed as offline in the instance. Delete these databases by using SQL Server Management Studio.

11.4.2 Attaching SQL Server databases

This section describes how to attach a database in SQL Server by using SQL Server Management Studio. Only one database can be attached at a time.

Attaching a database requires any of the following permissions: **CREATE DATABASE**, **CREATE ANY DATABASE**, or **ALTER ANY DATABASE**. Normally, these permissions are granted to the **sysadmin** role of the instance.

To attach a database

1. Run Microsoft SQL Server Management Studio.
2. Connect to the required SQL Server instance, and then expand the instance.
3. Right-click **Databases** and click **Attach**.
4. Click **Add**.
5. In the **Locate Database Files** dialog box, find and select the .mdf file of the database.

6. In the **Database Details** section, make sure that the rest of database files (.ndf and .ldf files) are found.

Details. SQL Server database files may not be found automatically, if:

- They are not in the default location, or they are not in the same folder as the primary database file (.mdf). Solution: Specify the path to the required files manually in the **Current File Path** column.
- You have recovered an incomplete set of files that make up the database. Solution: Recover the missing SQL Server database files from the backup.

7. When all of the files are found, click **OK**.

11.5 Recovering Exchange databases

This section describes recovery from both database backups and application-aware backups.

You can recover Exchange Server data to a live Exchange Server. This may be the original Exchange Server or an Exchange Server of the same version running on the machine with the same fully qualified domain name (FQDN). Agent for Exchange must be installed on the target machine.

The following table summarizes the Exchange Server data that you can select for recovery and the minimal user rights required to recover the data.

Exchange version	Data items	User rights
2007	Storage groups	Membership in the Exchange Organization Administrators role group.
2010/2013/2016	Databases	Membership in the Organization Management role group.

Alternatively, you can recover the databases (storage groups) as files. The database files, along with transaction log files, will be extracted from the backup to a folder that you specify. This can be useful if you need to extract data for an audit or further processing by third-party tools, or when the recovery fails for some reason and you are looking for a workaround to mount the databases manually (p. 100).

If you use only Agent for VMware, recovering databases as files is the only available recovery method.

To recover Exchange data

We will refer to both databases and storage groups as "databases" throughout this procedure.

1. When recovering from a database backup, click **Microsoft Exchange**. Otherwise, skip this step.
2. Select the machine that originally contained the data that you want to recover.
3. Click **Recovery**.
4. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Use other ways to recover:

- If the backup is located in the cloud storage or on a network share, click **Select machine**, select an online machine that has Agent for Exchange, and then select a recovery point.
- Select a recovery point on the Backups tab (p. 80).

The machine chosen for browsing in either of the above actions becomes a target machine for the Exchange data recovery.

5. Click **Recover > Exchange databases**.

6. Select the data that you want to recover.
7. If you want to recover the databases as files, click **Recover as files**, select a local or a network folder to save the files to, and then click **Recover**. Otherwise, skip this step.
8. Click **Recover**. If prompted, provide credentials to access the Exchange Server.
9. By default, the databases are recovered to the original ones. If the original database does not exist, it will be recreated.
To recover a database as a different one:
 - a. Click the database name.
 - b. In **Recover to**, select **New database**.
 - c. Specify the new database name.
 - d. Specify the new database path and log path. The folder you specify must not contain the original database and log files.
10. Click **Start recovery**.

The recovery progress is shown on the **Activities** tab.

11.5.1 Mounting Exchange Server databases

After recovering the database files, you can bring the databases online by mounting them. Mounting is performed by using Exchange Management Console, Exchange System Manager, or Exchange Management Shell.

The recovered databases will be in a Dirty Shutdown state. A database that is in a Dirty Shutdown state can be mounted by the system if it is recovered to its original location (that is, information about the original database is present in Active Directory). When recovering a database to an alternate location (such as a new database or as the recovery database), the database cannot be mounted until you bring it to a Clean Shutdown state by using the **Eseutil /r <Enn>** command. **<Enn>** specifies the log file prefix for the database (or storage group that contains the database) into which you need to apply the transaction log files.

The account you use to attach a database must be delegated an Exchange Server Administrator role and a local Administrators group for the target server.

For details about how to mount databases, see the following articles:

- Exchange 2016: <http://technet.microsoft.com/en-us/library/aa998871.aspx>
- Exchange 2013: [http://technet.microsoft.com/en-us/library/aa998871\(v=EXCHG.150\).aspx](http://technet.microsoft.com/en-us/library/aa998871(v=EXCHG.150).aspx)
- Exchange 2010: [http://technet.microsoft.com/en-us/library/aa998871\(v=EXCHG.141\).aspx](http://technet.microsoft.com/en-us/library/aa998871(v=EXCHG.141).aspx)
- Exchange 2007: [http://technet.microsoft.com/en-us/library/aa998871\(v=EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa998871(v=EXCHG.80).aspx)

11.6 Recovering Exchange mailboxes and mailbox items

This section describes how to recover Exchange mailboxes and mailbox items from database backups and from application-aware backups.

Overview

Granular recovery can be performed to Microsoft Exchange Server 2010 and later. The source backup may contain databases of any supported Exchange version.

Granular recovery can be performed by Agent for Exchange or Agent for VMware (Windows). The target Exchange Server and the machine running the agent must belong to the same Active Directory forest.

The following items can be recovered:

- Mailboxes
- Public folders
- Public folder items
- Email folders
- Email messages
- Calendar events
- Tasks
- Contacts
- Journal entries
- Notes

You can use search to locate the items.

When a mailbox is recovered to an existing mailbox, the existing items with matching IDs are overwritten.

Recovery of mailbox items does not overwrite anything. The mailbox items are always recovered to the **Recovered items** folder of the target mailbox.

Requirements on user accounts

A mailbox being recovered from a backup must have an associated user account in Active Directory.

User mailboxes and their contents can be recovered only if their associated user accounts are *enabled*. Shared, room, and equipment mailboxes can be recovered only if their associated user accounts are *disabled*.

A mailbox that does not meet the above conditions is skipped during recovery.

If some mailboxes are skipped, the recovery will succeed with warnings. If all mailboxes are skipped, the recovery will fail.

11.6.1 Recovering mailboxes

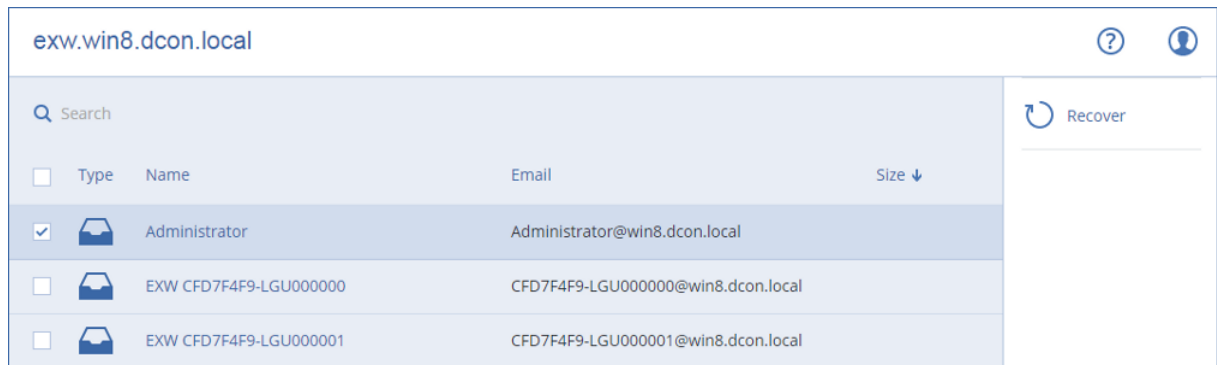
1. When recovering from a database backup, click **Microsoft Exchange**. Otherwise, skip this step.
2. Select the machine that originally contained the data that you want to recover.
3. Click **Recovery**.
4. Select a recovery point. Note that recovery points are filtered by location.

If the machine is offline, the recovery points are not displayed. Use other ways to recover:

- If the backup is located in the cloud storage or on a network share, click **Select machine**, select an online machine that has Agent for Exchange or Agent for VMware, and then select a recovery point.
- Select a recovery point on the Backups tab (p. 80).

The machine chosen for browsing in either of the above actions will perform the recovery instead of the original machine that is offline.

5. Click **Recover > Exchange mailboxes**.
6. Select the mailboxes that you want to recover.
You can search mailboxes by name. Wildcards are not supported.



7. Click **Recover**.
8. Click **Target machine with Microsoft Exchange Server** to select or change the target machine. This step allows recovery to a machine that is not running Agent for Exchange.
Specify the fully qualified domain name (FQDN) of the machine where the **Client Access** role of Microsoft Exchange Server is enabled. The machine must belong to the same Active Directory forest as the machine that performs the recovery.
If prompted, provide the credentials of an account that is a member of the **Organization Management** role group.
9. [Optional] Click **Database to re-create any missing mailboxes** to change the automatically selected database.
10. Click **Start recovery**.
11. Confirm your decision.

The recovery progress is shown on the **Activities** tab.

11.6.2 Recovering mailbox items

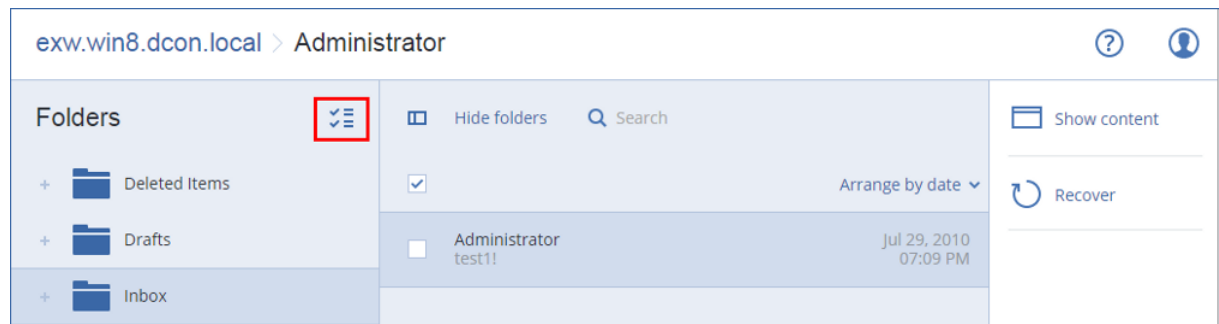
1. When recovering from a database backup, click **Microsoft Exchange**. Otherwise, skip this step.
2. Select the machine that originally contained the data that you want to recover.
3. Click **Recovery**.
4. Select a recovery point. Note that recovery points are filtered by location.
If the machine is offline, the recovery points are not displayed. Use other ways to recover:
 - If the backup is located in the cloud storage or on a network share, click **Select machine**, select an online machine that has Agent for Exchange or Agent for VMware, and then select a recovery point.
 - Select a recovery point on the Backups tab (p. 80).
The machine chosen for browsing in either of the above actions will perform the recovery instead of the original machine that is offline.
5. Click **Recover > Exchange mailboxes**.
6. Click the mailbox that originally contained the items that you want to recover.
7. Select the items that you want to recover.
The following search options are available. Wildcards are not supported.
 - For email messages: search by subject, sender, recipient, and date.
 - For events: search by title and date.

- For tasks: search by subject and date.
- For contacts: search by name, email address, and phone number.

When an email message is selected, you can click **Show content** to view its contents, including attachments.

Tip Click the name of an attached file to download it.

To be able to select folders, click the recover folders icon.



8. Click **Recover**.
9. Click **Target machine with Microsoft Exchange Server** to select or change the target machine. This step allows recovery to a machine that is not running Agent for Exchange. Specify the fully qualified domain name (FQDN) of the machine where the **Client Access** role of Microsoft Exchange Server is enabled. The machine must belong to the same Active Directory forest as the machine that performs the recovery. If prompted, provide the credentials of an account that is a member of the **Organization Management** role group.
10. In **Target mailbox**, view, change, or specify the target mailbox. By default, the original mailbox is selected. If this mailbox does not exist or a non-original target machine is selected, you must specify the target mailbox.
11. Click **Start recovery**.
12. Confirm your decision.

The recovery progress is shown on the **Activities** tab.

12 Advanced operations with virtual machines

12.1 Running a virtual machine from a backup (Instant Restore)

You can run a virtual machine from a disk-level backup that contains an operating system. This operation, also known as instant recovery, enables you to spin up a virtual server in seconds. The virtual disks are emulated directly from the backup and thus do not consume space on the datastore (storage). The storage space is required only to keep changes to the virtual disks.

We recommend running this temporary virtual machine for up to three days. Then, you can completely remove it or convert it to a regular virtual machine (finalize) without downtime.

As long as the temporary virtual machine exists, retention rules cannot be applied to the backup being used by that machine. Backups of the original machine can continue to run.

Usage examples

- **Disaster recovery**
Instantly bring a copy of a failed machine online.
- **Testing a backup**
Run the machine from the backup and ensure that the guest OS and applications are functioning properly.
- **Accessing application data**
While the machine is running, use application's native management tools to access and extract the required data.

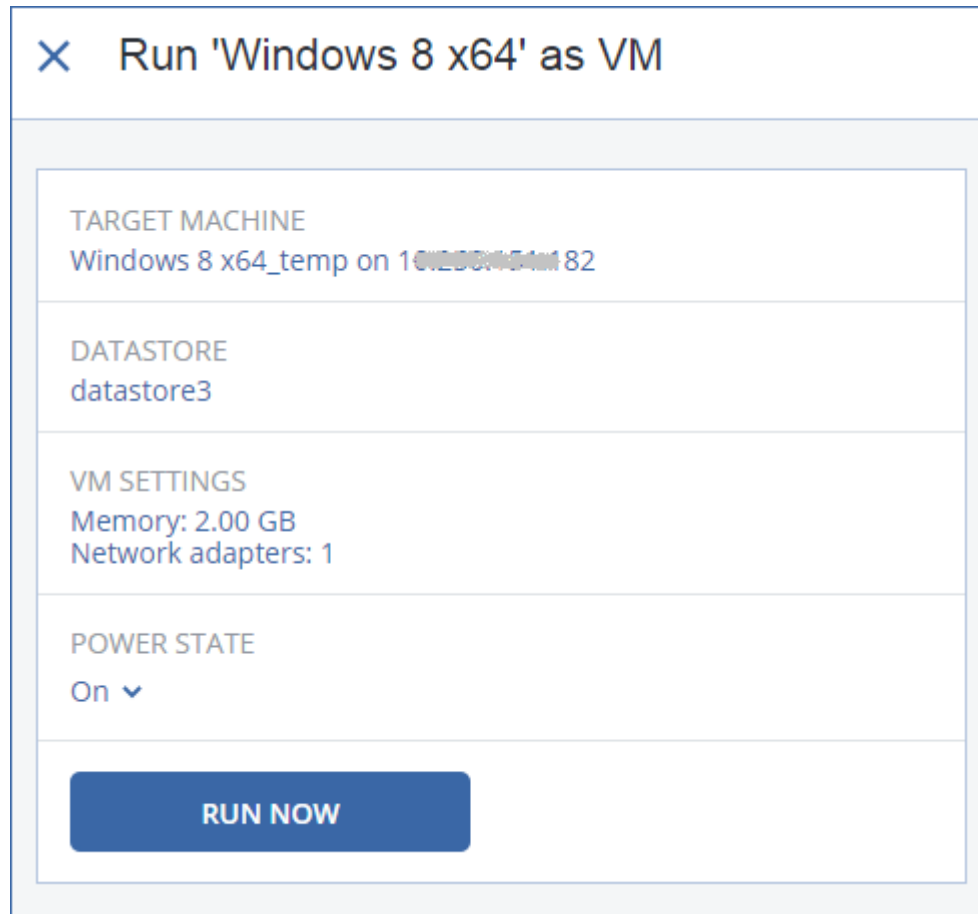
Prerequisites

- At least one Agent for VMware or Agent for Hyper-V must be registered in the backup service.
- The backup can be stored in a network folder or in a local folder of the machine where Agent for VMware or Agent for Hyper-V is installed. If you select a network folder, it must be accessible from that machine. A virtual machine can also be run from a backup stored in the cloud storage, but it works slower because this operation requires intense random-access reading from the backup.
- The backup must contain an entire machine or all of the volumes that are required for the operating system to start.
- Backups of both physical and virtual machines can be used. Backups of *Virtuozzo containers* cannot be used.

12.1.1 Running the machine

1. Do one of the following:
 - Select a backed-up machine, click **Recovery**, and then select a recovery point.
 - Select a recovery point on the Backups tab (p. 80).
2. Click **Run as VM**.

The software automatically selects the host and other required parameters.



3. [Optional] Click **Target machine**, and then change the virtual machine type (ESXi or Hyper-V), the host, or the virtual machine name.
4. [Optional] Click **Datastore** for ESXi or **Path** for Hyper-V, and then select the datastore for the virtual machine.
Changes to the virtual disks accumulate while the machine is running. Ensure that the selected datastore has enough free space.
5. [Optional] Click **VM settings** to change the memory size and network connections of the virtual machine.
6. [Optional] Select the VM power state (**On/Off**).
7. Click **Run now**.

As a result, the machine appears in the web interface with one of the following icons:



. Such virtual machines cannot be selected for backup.



or

12.1.2 Deleting the machine

We do not recommend to delete a temporary virtual machine directly in vSphere/Hyper-V. This may lead to artifacts in the web interface. Also, the backup from which the machine was running may remain locked for a while (it cannot be deleted by retention rules).

To delete a virtual machine that is running from a backup

1. On the **All devices** tab, select a machine that is running from a backup.
2. Click **Delete**.

The machine is removed from the web interface. It is also removed from the vSphere or Hyper-V inventory and datastore (storage). All changes that occurred to the data while the machine was running are lost.

12.1.3 Finalizing the machine

While a virtual machine is running from a backup, the virtual disks' content is taken directly from that backup. Therefore, the machine will become inaccessible or even corrupted if the connection is lost to the backup location or to the backup agent.

For an ESXi machine, you have the option to make this machine permanent, i.e. recover all of its virtual disks, along with the changes that occurred while the machine was running, to the datastore that stores these changes. This process is named finalization.

Finalization is performed without downtime. The virtual machine will *not* be powered off during finalization.

To finalize a machine that is running from a backup

1. On the **All devices** tab, select a machine that is running from a backup.
2. Click **Finalize**.
3. [Optional] Specify a new name for the machine.
4. [Optional] Change the disk provisioning mode. The default setting is **Thin**.
5. Click **Finalize**.

The machine name changes immediately. The recovery progress is shown on the **Activities** tab. Once the recovery is completed, the machine icon changes to that of a regular virtual machine.

12.2 Replication of virtual machines

Replication is available only for VMware ESXi virtual machines.

Replication is the process of creating an exact copy (replica) of a virtual machine, and then maintaining the replica in sync with the original machine. By replicating a critical virtual machine, you will always have a copy of this machine in a ready-to-start state.

The replication can be started manually or on the schedule you specify. The first replication is full (copies the entire machine). All subsequent replications are incremental and are performed with Changed Block Tracking (p. 110), unless this option is disabled.

Replication vs. backing up

Unlike scheduled backups, a replica keeps only the latest state of the virtual machine. A replica consumes datastore space, while backups can be kept on a cheaper storage.

However, powering on a replica is much faster than a recovery and faster than running a virtual machine from a backup. When powered on, a replica works faster than a VM running from a backup and does not load the Agent for VMware.

Usage examples

- **Replicate virtual machines to a remote site.**

Replication enables you to withstand partial or complete datacenter failures, by cloning the virtual machines from a primary site to a secondary site. The secondary site is usually located in a remote facility that is unlikely to be affected by environmental, infrastructure, or other factors that might cause the primary site failure.

- **Replicate virtual machines within a single site (from one host/datastore to another).**

Onsite replication can be used for high availability and disaster recovery scenarios.

What you can do with a replica

- **Test a replica** (p. 108)

The replica will be powered on for testing. Use vSphere Client or other tools to check if the replica works correctly. Replication is suspended while testing is in progress.

- **Failover to a replica** (p. 108)

Failover is a transition of the workload from the original virtual machine to its replica. Replication is suspended while a failover is in progress.

- **Back up the replica**

Both backup and replication require access to virtual disks, and thus impact the performance of the host where the virtual machine is running. If you want to have both a replica and backups of a virtual machine, but don't want to put additional load on the production host, replicate the machine to a different host, and set up backups of the replica.

Restrictions

The following types of virtual machines cannot be replicated:

- Fault-tolerant machines running on ESXi 5.5 and lower.
- Machines running from backups.
- Replicas of virtual machines.

12.2.1 Creating a replication plan

A replication plan must be created for each machine individually. It is not possible to apply an existing plan to other machines.

To create a replication plan

1. Select a virtual machine to replicate.
2. Click **Replication**.
The software displays a new replication plan template.
3. [Optional] To modify the replication plan name, click the default name.
4. Click **Target machine**, and then do the following:
 - a. Select whether to create a new replica or use an existing replica of the original machine.
 - b. Select the ESXi host and specify the new replica name, or select an existing replica.
The default name of a new replica is **[Original Machine Name]_replica**.
 - c. Click **OK**.
5. [Only when replicating to a new machine] Click **Datastore**, and then select the datastore for the virtual machine.
6. [Optional] Click **Schedule** to change the replication schedule.

By default, replication is performed on a daily basis, Monday to Friday. You can select the time to run the replication.

If you want to change the replication frequency, move the slider, and then specify the schedule.

You can also do the following:

- Set a date range for when the schedule is effective. Select the **Run the plan within a date range** check box, and then specify the date range.
 - Disable the schedule. In this case, replication can be started manually.
7. [Optional] Click the gear icon to modify the replication options (p. 110).
 8. Click **Apply**.
 9. [Optional] To run the plan manually, click **Run now** on the plan panel.

As a result of running a replication plan, the virtual machine replica appears in the **All devices** list



with the following icon:

12.2.2 Testing a replica

To prepare a replica for testing

1. Select a replica to test.
2. Click **Test replica**.
3. Click **Start testing**.
4. Select whether to connect the powered-on replica to a network. By default, the replica will not be connected to a network.
5. [Optional] If you chose to connect the replica to the network, select the **Stop original virtual machine** check box to stop the original machine before powering on the replica.
6. Click **Start**.

To stop testing a replica

1. Select a replica for which testing is in progress.
2. Click **Test replica**.
3. Click **Stop testing**.
4. Confirm your decision.

12.2.3 Failing over to a replica

To failover a machine to a replica

1. Select a replica to failover to.
2. Click **Replica actions**.
3. Click **Failover**.
4. Select whether to connect the powered-on replica to a network. By default, the replica will be connected to the same network as the original machine.
5. [Optional] If you chose to connect the replica to the network, clear the **Stop original virtual machine** check box to keep the original machine online.
6. Click **Start**.

While the replica is in a failover state, you can choose one of the following actions:

- **Stop failover** (p. 109)

Stop failover if the original machine was fixed. The replica will be powered off. Replication will be resumed.

- **Perform permanent failover to the replica** (p. 109)

This instant operation removes the 'replica' flag from the virtual machine, so that replication to it is no longer possible. If you want to resume replication, edit the replication plan to select this machine as a source.

- **Failback** (p. 109)

Perform failback if you failed over to the site that is not intended for continuous operations. The replica will be recovered to the original or a new virtual machine. Once the recovery to the original machine is complete, it is powered on and replication is resumed. If you choose to recover to a new machine, edit the replication plan to select this machine as a source.

12.2.3.1 Stopping failover

To stop a failover

1. Select a replica that is in the failover state.
2. Click **Replica actions**.
3. Click **Stop failover**.
4. Confirm your decision.

12.2.3.2 Performing a permanent failover

To perform a permanent failover

1. Select a replica that is in the failover state.
2. Click **Replica actions**.
3. Click **Permanent failover**.
4. [Optional] Change the name of the virtual machine.
5. [Optional] Select the **Stop original virtual machine** check box.
6. Click **Start**.

12.2.3.3 Failing back

To failback from a replica

1. Select a replica that is in the failover state.
2. Click **Replica actions**.
3. Click **Failback from replica**.

The software automatically selects the original machine as the target machine.
4. [Optional] Click **Target machine**, and then do the following:
 - a. Select whether to failback to a new or existing machine.
 - b. Select the ESXi host and specify the new machine name, or select an existing machine.
 - c. Click **OK**.
5. [Optional] When failing back to a new machine, you can also do the following:
 - Click **Datastore** to select the datastore for the virtual machine.
 - Click **VM settings** to change the memory size, the number of processors, and the network connections of the virtual machine.
6. [Optional] Click **Recovery options** to modify the failback options (p. 110).
7. Click **Start recovery**.

8. Confirm your decision.

12.2.4 Replication options

To modify the replication options, click the gear icon next to the replication plan name, and then click **Replication options**.

Changed Block Tracking (CBT)

This option is similar to the backup option "Changed Block Tracking (CBT)" (p. 50).

Disk provisioning

This option defines the disk provisioning settings for the replica.

The preset is: **Thin provisioning**.

The following values are available: **Thin provisioning, Thick provisioning, Keep the original setting**.

Error handling

This option is similar to the backup option "Error handling" (p. 51).

Pre/Post commands

This option is similar to the backup option "Pre/Post commands" (p. 57).

Volume Shadow Copy Service VSS for virtual machines

This option is similar to the backup option "Volume Shadow Copy Service VSS for virtual machines" (p. 62).

12.2.5 Failback options

To modify the failback options, click **Recovery options** when configuring failback.

Error handling

This option is similar to the recovery option "Error handling" (p. 76).

Performance

This option is similar to the recovery option "Performance" (p. 78).

Pre/Post commands

This option is similar to the recovery option "Pre/Post commands" (p. 78).

VM power management

This option is similar to the recovery option "VM power management" (p. 80).

12.2.6 Seeding an initial replica

To speed up replication to a remote location and save network bandwidth, you can perform replica seeding.

Important To perform replica seeding, Agent for VMware (Virtual Appliance) must be running on the target ESXi.

To seed an initial replica

1. Do one of the following:
 - If the original virtual machine can be powered off, power it off, and then skip to step 4.
 - If the original virtual machine cannot be powered off, continue to the next step.
2. Create a replication plan (p. 107).

When creating the plan, in **Target machine**, select **New replica** and the ESXi that hosts the original machine.
3. Run the plan once.

A replica is created on the original ESXi.
4. Export the virtual machine (or the replica) files to an external hard drive.
 - a. Connect the external hard drive to the machine where vSphere Client is running.
 - b. Connect vSphere Client to the original vCenter\ESXi.
 - c. Select the newly created replica in the inventory.
 - d. Click **File > Export > Export OVF template**.
 - e. In **Directory**, specify the folder on the external hard drive.
 - f. Click **OK**.
5. Transfer the hard drive to the remote location.
6. Import the replica to the target ESXi.
 - a. Connect the external hard drive to the machine where vSphere Client is running.
 - b. Connect vSphere Client to the target vCenter\ESXi.
 - c. Click **File > Deploy OVF template**.
 - d. In **Deploy from a file or URL**, specify the template that you exported in step 4.
 - e. Complete the import procedure.
7. Edit the replication plan that you created in step 2. In **Target machine**, select **Existing replica**, and then select the imported replica.

As a result, the software will continue updating the replica. All replications will be incremental.

12.3 Managing virtualization environments

You can view the vSphere and Hyper-V environments in their native presentation. Once the corresponding agent is installed and registered, the **VMware** or **Hyper-V** tab appears under **Devices**.

The **VMware** tab enables you to change access credentials for the vCenter Server or stand-alone ESXi host without re-installing the agent.

To change the vCenter Server or ESXi host access credentials

1. Under **Devices**, click **VMware**.
2. Click **Hosts and Clusters**.
3. In the **Hosts and Clusters** list (to the right of the **Hosts and Clusters** tree), select the vCenter Server or stand-alone ESXi host that was specified during the Agent for VMware installation.
4. Click **Overview**.
5. Under **Credentials**, click the user name.
6. Specify the new access credentials, and then click **OK**.

12.4 Machine migration

You can perform machine migration by recovering its backup to a non-original machine.

The following table summarizes the available migration options.

Backed-up machine type	Available recovery destinations		
	Physical machine	ESXi virtual machine	Hyper-V virtual machine
Physical machine	+	+	+
VMware ESXi virtual machine	+	+	+
Hyper-V virtual machine	+	+	+

For instructions on how to perform migration, refer to the following sections:

- Physical-to-virtual (P2V) - "Physical machine to virtual" (p. 65)
- Virtual-to-virtual (V2V) - "Virtual machine" (p. 66)
- Virtual-to-physical (V2P) - "Virtual machine" (p. 66) or "Recovering disks by using bootable media" (p. 68)

Although it is possible to perform V2P migration in the web interface, we recommend using bootable media in specific cases. Sometimes, you may want to use the media for migration to ESXi or Hyper-V.

The media enables you to do the following:

- Choose individual disks or volumes for recovery.
- Manually map the disks from the backup to the target machine disks.
- Recreate logical volumes (LVM) or Linux Software RAID on the target machine.
- Provide drivers for specific hardware that is critical for the system bootability.

12.5 Agent for VMware - necessary privileges

This section describes the privileges required for operations with ESXi virtual machines and, additionally, for virtual appliance deployment. Agent for VMware (Virtual Appliance) is available in on-premise deployment only.

To perform operations on all hosts and clusters managed by a vCenter Server, Agent for VMware needs the privileges on the vCenter Server. If you want the agent to operate on a specific ESXi host only, provide the agent with the same privileges on the host.

Specify the account with the necessary privileges during Agent for VMware installation or configuration. If you need to change the account at a later time, refer to the "Managing virtualization environments" (p. 111) section.

Object	Privilege	Operation				
		Back up a VM	Recover to a new VM	Recover to an existing VM	Run VM from backup	VA deployment
Datastore	Allocate space		+	+	+	+
	Browse datastore				+	+
	Configure datastore	+	+	+	+	+

		Operation				
Object	Privilege	Back up a VM	Recover to a new VM	Recover to an existing VM	Run VM from backup	VA deployment
	Low level file operations				+	+
Global	Licenses	+	+	+	+	
	Disable methods	+				
	Enable methods	+				
Host > Configuration	VM autostart configuration					+
	Storage partition configuration				+	
Host > Inventory	Modify cluster					+
Host > Local operations	Create VM				+	+
	Delete VM				+	+
	Reconfigure VM				+	+
Network	Assign network		+	+	+	+
Resource	Assign VM to resource pool		+	+	+	+
vApp	Import					+
Virtual machine > Configuration	Add existing disk	+	+		+	
	Add new disk		+	+	+	+
	Add or remove device		+		+	+
	Advanced		+			+
	Change CPU count		+			
	Disk lease	+				
	Memory		+			
	Remove disk	+	+	+	+	
	Rename		+			
	Settings			+		
Virtual machine > Interaction	Acquire guest control ticket				+	
	Configure CD media		+			
	Console interaction					+
	Guest operating system management by VIX API					+

		Operation				
Object	Privilege	Back up a VM	Recover to a new VM	Recover to an existing VM	Run VM from backup	VA deployment
	Power off			+	+	+
	Power on		+	+	+	+
Virtual machine > Inventory	Create from existing		+	+	+	
	Create new		+	+	+	+
	Move					+
	Register				+	
	Remove		+	+	+	+
	Unregister				+	
Virtual machine > Provisioning	Allow disk access		+	+	+	
	Allow virtual machine download	+	+	+	+	
Virtual machine > State	Create snapshot	+		+	+	+
	Remove snapshot	+		+	+	+

12.6 Windows Azure and Amazon EC2 virtual machines

To back up a Windows Azure or Amazon EC2 virtual machine, install a backup agent on the machine. The backup and recovery operations are the same as with a physical machine. Nevertheless, the machine is counted as virtual when you set quotas for the number of machines in a cloud deployment.

The difference from a physical machine is that Windows Azure and Amazon EC2 virtual machines cannot be booted from bootable media. If you need to recover to a new Windows Azure or Amazon EC2 virtual machine, follow the procedure below.

To recover a machine as a Windows Azure or Amazon EC2 virtual machine

1. Create a new virtual machine from an image/template in Windows Azure or Amazon EC2. The new machine must have the same disk configuration as the machine that you want to recover.
2. Install Agent for Windows or Agent for Linux on the new machine.
3. Recover the backed-up machine as described in "Physical machine" (p. 64). When configuring the recovery, select the new machine as the target machine.

Network requirements

The agents installed on the backed-up machines must be able to communicate with the management server over the network.

On-premise deployment

- If both the agents and the management server are installed in the Azure/EC2 cloud, all machines are already located in the same network. No additional actions are required.

- If the management server is located outside the Azure/EC2 cloud, the machines in the cloud will not have network access to the local network where the management server is installed. To enable the agents installed on such machines to communicate with the management server, a virtual private network (VPN) connection between the local (on-premises) and the cloud (Azure/EC2) network must be created. For instructions about how to create the VPN connection, refer to the following articles:

Amazon EC2:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html#vpn-create-cgw

Windows Azure:

<https://azure.microsoft.com/en-us/documentation/articles/vpn-gateway-site-to-site-create>

Cloud deployment

In a cloud deployment, the management server is located in one of the Acronis data centers and is thus reachable by the agents. No additional actions are required.

13 Management server settings

These settings are only available in on-premise deployments.

To access these settings, click **Settings > System settings**.

For information about managing Acronis Backup licenses, refer to "Managing licenses" (p. 26).

13.1 Email server

You can specify an email server that will be used to send email notifications from the management server.

To specify the email server

1. Click **Settings > System settings > Email server**.
2. In **Email service**, select one of the following:
 - **Custom**
 - **Gmail**

The **Less secure apps** setting must be turned on in your Gmail account. For more information, refer to <https://support.google.com/accounts/answer/6010255>.
 - **Yahoo Mail**
 - **Outlook.com**
3. [Only for a custom email service] Specify the following settings:
 - In **SMTP server**, enter the name of the outgoing mail server (SMTP).
 - In **SMTP port**, set the port of the outgoing mail server. By default, the port is set to 25.
 - Select whether to use SSL or TLS encryption. Select **None** to disable encryption.
 - If the SMTP server requires authentication, select the **SMTP server requires authentication** check box, and then specify the credentials of an account that will be used to send messages. If you are not sure whether the SMTP server requires authentication, contact your network administrator or your email service provider for assistance.
 - Some Internet service providers require authentication on the incoming mail server before being allowed to send something. If this is the case for you, select the **Configure an incoming mail server (POP)** check box to enable a POP server and to set up its settings:

- In **POP server**, enter the name of the POP server.
 - In **POP port**, set the port of the incoming mail server. By default, the port is set to 110.
 - Specify the access credentials for the incoming mail server.
4. [Only for Gmail, Yahoo Mail, and Outlook.com] Specify the credentials of an account that will be used to send messages.
 5. In **Sender**, type the name of the sender. This name will be shown in the **From** field of the email notifications. If you leave this field empty, the messages will contain the account specified in step 3 or 4.
 6. [Optional] Click **Send test message** to check whether the email notifications work correctly with the specified settings. Enter an email address to send the test message to.

13.2 Email notifications

You can configure default settings that are common for all email notifications sent from the management server.

When creating a backup plan, you can either use the default settings, or override them with custom values that will be specific for this plan only.

Important When the default settings are changed, all backup plans that use the default settings are affected.

Before configuring these settings, ensure that the **Email server** (p. 115) settings are configured.

To configure default email notification settings

1. Click **Settings > System settings > Email notifications**.
2. In the **Recipients' email addresses** field, type the destination email address. You can enter several addresses separated by semicolons.
3. Select the types of notifications that you want to be sent. The following types are available:
 - **Errors**
 - **Warnings**
 - **Successful backups**

The subject of the email messages is based on the following template: **[subject] [machine name] [backup plan name]**. The **[subject]** placeholder will be replaced by one of the following phrases: **Backup succeeded, Backup failed, Backup succeeded with warnings**.

14 Managing groups and accounts

The functionality described in this section is available only in a cloud deployment for accounts that have administrative privileges.

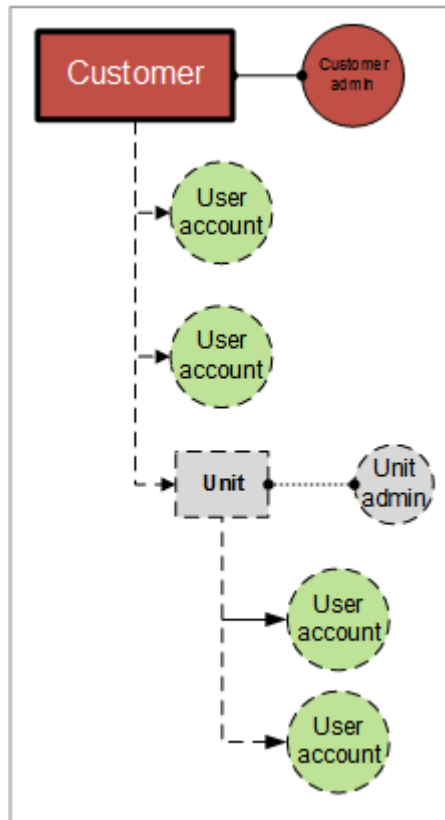
14.1 Accounts and groups

There are two account types: **administrator accounts** and **user accounts**. Both the user and the administrators can manage backups of the user's data.

Each account belongs to a group. The **Customer** group is automatically created for your organization. Optionally, you can create **Unit** groups, which typically correspond to units or departments of the organization.

An administrator can create and manage groups, administrator accounts, and user accounts on or below their level in the hierarchy.

The following diagram illustrates two hierarchy levels—the customer and unit groups. Optional groups and accounts are shown by a dotted line.



The following table summarizes operations that can be performed by the administrators and users.

Operation	Users	Administrators
Create groups	No	Yes
Create accounts	No	Yes
Download and install the backup software	Yes	Yes
Manage backup	Yes	Yes
Manage recovery	Yes	Yes
Create reports about the service usage	No	Yes

14.2 Creating a group

You may want to create a new **Unit** group within your company when expanding the backup service to a new organizational unit.

To create a group

1. Log in to the backup console.
2. Click **Manage accounts**.

3. Select a group in which you want to create the new group.
4. On the bottom of the **Groups** pane, click "+".
5. In **Name**, specify a name for the new group.
6. [Optional] In **Identifier**, type a string that will act as the identifier for the group. This identifier will appear in monthly reports, together with the group's usage data. You can use this identifier to refer to this group in other software, such as in your billing or monitoring systems.
The identifier can consist of up to 256 Unicode characters (for example, numbers and Latin letters). It does not need to be unique across groups.
7. In **Default language**, select the default language of notifications, reports, and backup software that will be used within this group.
8. [Optional] In **Backup locations**, select the backup locations for this group and its child groups. The following values are available:
 - **Local and cloud**
 - **Cloud only**
9. [Optional] Disable the **Agent auto update** switch. If you do this, the agents that are registered under the accounts within this group and its child groups will not be updated automatically when a new version is released.
10. [Optional] In **Contact information**, specify the contact information for the group.
11. Click **Create**.

The newly created group appears in the **Groups** tree.

If you want to specify the billing information for a group, select the group in the **Groups** list, click **Properties**, and then complete the billing information section.

14.3 Creating an account

At least one account (either an administrator or a user) must exist within a unit.

To create an account

1. Log in to the backup console.
2. Click **Manage accounts**.
3. Select a group in which you want to create the account.
4. Click the **Accounts** tab.
5. Click "+".
6. Specify the following contact information for the account.
 - **Login**

Important Each account must have a unique login. You can create multiple logins using the same email address.

 - **Email address**
 - [Optional] **First name**
 - [Optional] **Last name**
7. If you want this account to be an administrator account, enable the **Administrator privileges** switch.
8. [Optional] Disable the **Agent auto update** switch. If you do this, the agents that are registered under this account not be updated automatically when a new version is released.

9. [Optional] Specify the storage quota and the maximum number of machines the user is allowed to back up.

- **Physical workstations**
- **Physical servers**
- **Windows Server Essentials**
- **Virtual hosts**
- **Mobile devices**
- **Storage quota**

These quotas are "soft". If any of these values are exceeded, a notification will be sent to the email address specified in step 6. Restrictions on using the backup service are not applied.

10. [Optional] Specify the quota overages. An overage allows the user to exceed the quota by the specified value. When the overage is exceeded, backups fail.

Important *If you set both a quota and its overage to zero, the corresponding functionality will be hidden from the user.*

11. [Optional] In **Backup locations**, select the backup locations for this account. The following values are available:

- **Local and cloud**
- **Cloud only**

12. [Optional] Change the **Backup notifications** level. If you disable backup notifications, the notifications about backup failures, warnings, and successful backups will not be sent to the specified email address.

13. [Optional] Disable **Business notifications**. If you do this, notifications about exceeded quotas will not be sent to the specified email address.

14. Click **Add**.

As a result:

- A new account appears in the **Accounts** tab.
- An email message containing the activation link is sent to the email address you specified.

14.4 Creating a report about the service usage

Usage reports provide historical data about using the backup service.

Only administrators can create these reports.

Reporting parameters

The report includes the following data about a unit and its accounts:

- Size of backups by group, by account, by machine type.
- Amount of protected machines by group, by account, by machine type.
- Price value by group, by account, by machine type.
- The total size of backups.
- The total amount of protected machines.
- Total price value.

Report scope

You can select the scope of the report from the following values:

- **Direct customers and partners**

The report will include the values of the reporting parameters only for the immediate child groups of your group.

- **All customers and partners**

The report will include the values of the reporting parameters for all child groups of your group.

- **All customers and partners (including account details)**

The report will include the values of the reporting parameters for all child groups of your group and for all user accounts within the groups.

Enabling or disabling scheduled usage reports

A scheduled report covers system usage data for the last full calendar month. The reports are generated at 23:59:59 by UTC time on the first day of a month and sent on the second day of that month to all administrators of your group.

1. In the account management console, click **Reports**.
2. Select the **Scheduled** tab.
3. Enable or disable the scheduled usage reports by clicking the on/off switch.
4. In **Level of detail**, select the report scope as described above.

Generating a custom usage report

This type of report can be generated on demand and cannot be scheduled. The report will be sent to your email address.

1. In the account management console, click **Reports**.
2. Select the **Custom** tab.
3. In **Period**, select the reporting period:
 - **Current calendar month**
 - **Previous calendar month**
 - **Custom**
4. If you want to specify a custom reporting period, select the start and the end dates. Otherwise, skip this step.
5. In **Type**, select the report type:
 - **Summary report**: The report will include the total values of the reporting parameters for the specified period, including the total price value.
 - **Daily statistics**: The report will include the values of the reporting parameters for each day of the specified period, excluding the price values.
6. In **Level of detail**, select the report scope as described above.
7. To generate the report, click **Generate and send**.

14.5 Limiting access to the web interface

You can limit access to the web interface by specifying a list of IP addresses from which the members of a group are allowed to log in.

This restriction is *not* applied to the members of the child groups.

To limit access to the web interface

1. Log in to the account management console.

2. Select a group for which you want to limit the access.
3. Click **Security**.
4. Enable the **Logon control** switch.
5. In **Allowed IP addresses**, specify the allowed IP addresses.

You can enter any of the following parameters, separated by a semicolon:

- IP addresses, for example 192.0.2.0
- IP ranges, for example 192.0.2.0-192.0.2.255
- Subnets, for example 192.0.2.0/24

15 Troubleshooting

This section describes how to save an agent log to a .zip file. If a backup fails for an unclear reason, this file will help the technical support personnel to identify the problem.

To collect logs

1. Select the machine that you want to collect the logs from.
2. Click **Activities**.
3. Click **Collect system information**.
4. If prompted by your web browser, specify where to save the file.

Copyright Statement

Copyright © Acronis International GmbH, 2002-2016. All rights reserved.

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

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

Linux is a registered trademark of Linus Torvalds.

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

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

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

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

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

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

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

Acronis patented technologies

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

16 Glossary

B

Backup set

A group of backups to which an individual retention rule can be applied.

For the **Custom** backup scheme, the backup sets correspond to the backup methods (**Full**, **Differential**, and **Incremental**).

In all other cases, the backup sets are **Monthly**, **Daily**, **Weekly**, and **Hourly**.

- A monthly backup is the first backup created after a month starts.
- A weekly backup is the first backup created on the day of the week selected in the **Weekly backup** option (click the gear icon, then **Backup options** > **Weekly backup**).
- A daily backup is the first backup created after a day starts.
- An hourly backup is the first backup created after an hour starts.

D

Differential backup

A differential backup stores changes to the data against the latest full backup (p. 123). You need access to the corresponding full backup to recover the data from a differential backup.

F

Full backup

A self-sufficient backup containing all data chosen for backup. You do not need access to any other backup to recover the data from a full backup.

I

Incremental backup

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

S

Single-file backup format

A new backup format, in which the initial full and subsequent incremental backups are saved to a single .tib file, instead of a chain of files. This format leverages the speed of the incremental backup method, while avoiding its main disadvantage—difficult deletion of outdated backups. The software marks the blocks used by outdated backups as "free" and writes new backups to these blocks. This results in extremely fast cleanup, with minimal resource consumption.