

Cyber Protection

Version 20.09

Table of contents

1	Welcome to Cyber Protection.....	8
2	Cyber Protection editions.....	10
3	Comparison of editions	11
4	Software requirements	13
4.1	Supported Cyber Protect features by operating system	13
4.2	Supported web browsers.....	16
4.3	Supported operating systems and environments	17
4.4	Supported Microsoft SQL Server versions	19
4.5	Supported Microsoft Exchange Server versions	20
4.6	Supported Microsoft SharePoint versions.....	20
4.7	Supported Oracle Database versions	20
4.8	Supported SAP HANA versions	20
4.9	Supported virtualization platforms.....	20
4.10	Compatibility with encryption software	24
5	Supported file systems.....	25
6	Activating the account	26
6.1	Two-factor authentication.....	27
7	Accessing the Cyber Protection service	28
8	Installing the software	28
8.1	Preparation	29
8.2	Linux packages	33
8.3	Proxy server settings.....	35
8.4	Installing agents	38
8.4.1	Changing the logon account on Windows machines	41
8.5	Unattended installation or uninstallation	42
8.5.1	Unattended installation or uninstallation in Windows	42
8.5.2	Unattended installation or uninstallation in Linux	47
8.5.3	Unattended installation and uninstallation in macOS.....	51
8.6	Registering machines manually	53
8.7	Autodiscovery of machines.....	56
8.7.1	Autodiscovery and manual discovery.....	58
8.7.2	Managing discovered machines.....	62
8.7.3	Troubleshooting	63
8.8	Deploying Agent for VMware (Virtual Appliance) from an OVF template	64
8.8.1	Before you start.....	64
8.8.2	Deploying the OVF template	64
8.8.3	Configuring the virtual appliance	65
8.9	Deploying Agent for Virtuozzo Infrastructure Platform (Virtual Appliance) from a QCOW2 template.....	66

8.9.1	Before you start.....	66
8.9.2	Configuring networks in Virtuozzo Hybrid Infrastructure	67
8.9.3	Configuring user accounts in Virtuozzo Hybrid Infrastructure	71
8.9.4	Deploying the QCOW2 template	72
8.9.5	Configuring the virtual appliance	73
8.10	Deploying agents through Group Policy	76
8.11	Updating agents.....	77
8.12	Uninstalling agents	79
8.13	Security settings.....	80
8.14	Changing quota for a device	82
8.15	Cyber Protection services installed in your environment	82
9	Service console	83
10	Voice control for operations in the console.....	85
11	Device groups	87
11.1	Creating a static group.....	89
11.2	Adding devices to static groups	89
11.3	Creating a dynamic group.....	89
11.4	Applying a protection plan to a group.....	93
12	Multitenancy support	94
13	Protection plan and modules	95
13.1	Creating a protection plan	95
13.2	Default protection plans	97
13.3	Resolving plan conflicts.....	101
13.4	Operations with protection plans.....	101
14	#CyberFit Score for machines	103
14.1	Running a #CyberFit Score scan.....	107
15	Backup and recovery.....	109
15.1	Backup	109
15.2	Protection plan cheat sheet.....	111
15.3	Selecting data to back up.....	113
15.3.1	Selecting disks/volumes.....	113
15.3.2	Selecting files/folders.....	115
15.3.3	Selecting system state	117
15.3.4	Selecting ESXi configuration	117
15.4	Continuous data protection (CDP).....	118
15.5	Selecting a destination.....	123
15.5.1	About Secure Zone	125
15.6	Schedule.....	127
15.6.1	Schedule by events	129
15.6.2	Start conditions	131
15.7	Retention rules	136

15.8	Replication	137
15.9	Encryption	138
15.10	Notarization	140
15.11	Starting a backup manually	140
15.12	Default backup options.....	141
15.13	Backup options	141
15.13.1	Alerts	144
15.13.2	Backup consolidation	144
15.13.3	Backup file name	145
15.13.4	Backup format	148
15.13.5	Backup validation	149
15.13.6	Changed block tracking (CBT).....	149
15.13.7	Cluster backup mode	150
15.13.8	Compression level	151
15.13.9	Error handling	151
15.13.10	Fast incremental/differential backup	152
15.13.11	File filters	153
15.13.12	File-level backup snapshot	154
15.13.13	Forensic data	154
15.13.14	Log truncation.....	162
15.13.15	LVM snapshotting	162
15.13.16	Mount points.....	163
15.13.17	Multi-volume snapshot.....	163
15.13.18	Performance and backup window	164
15.13.19	Physical Data Shipping	166
15.13.20	Pre/Post commands.....	167
15.13.21	Pre/Post data capture commands	169
15.13.22	Scheduling.....	170
15.13.23	Sector-by-sector backup	171
15.13.24	Splitting	171
15.13.25	Task failure handling	172
15.13.26	Task start conditions	172
15.13.27	Volume Shadow Copy Service (VSS)	172
15.13.28	Volume Shadow Copy Service (VSS) for virtual machines	173
15.13.29	Weekly backup	174
15.13.30	Windows event log	174
15.14	Recovery	174
15.14.1	Recovery cheat sheet.....	174
15.14.2	Safe recovery	175
15.14.3	Creating bootable media	176
15.14.4	Startup Recovery Manager.....	177
15.14.5	Recovering a machine	178
15.14.6	Recovering files.....	187
15.14.7	Recovering system state	192
15.14.8	Recovering ESXi configuration.....	192
15.14.9	Recovery options	193
15.15	Operations with backups	199
15.15.1	The Backup storage tab	199
15.15.2	Mounting volumes from a backup.....	200
15.15.3	Deleting backups	201
15.16	Protecting Microsoft applications	202
15.16.1	Prerequisites	203
15.16.2	Database backup	205

15.16.3	Application-aware backup	210
15.16.4	Mailbox backup	211
15.16.5	Recovering SQL databases.....	212
15.16.6	Recovering Exchange databases	215
15.16.7	Recovering Exchange mailboxes and mailbox items	217
15.16.8	Changing the SQL Server or Exchange Server access credentials.....	223
15.17	Protecting mobile devices	223
15.18	Protecting Hosted Exchange data.....	227
15.18.1	Selecting mailboxes.....	227
15.18.2	Recovering mailboxes and mailbox items	228
15.19	Protecting Office 365 data	230
15.19.1	Using the locally installed Agent for Office 365	232
15.19.2	Using the cloud Agent for Office 365.....	234
15.20	Protecting G Suite data	252
15.20.1	Adding a G Suite organization	253
15.20.2	Protecting Gmail data	254
15.20.3	Protecting Google Drive files.....	257
15.20.4	Protecting Shared drive files.....	260
15.20.5	Notarization	263
15.21	Protecting Oracle Database	264
15.22	Protecting SAP HANA.....	264
15.23	Protecting websites and hosting servers.....	264
15.23.1	Protecting websites.....	264
15.23.2	Protecting web hosting servers.....	267
15.24	Special operations with virtual machines	268
15.24.1	Running a virtual machine from a backup (Instant Restore)	268
15.24.2	Working in VMware vSphere	271
15.24.3	Backing up clustered Hyper-V machines	285
15.24.4	Limiting the total number of simultaneously backed-up virtual machines	286
15.24.5	Machine migration	287
15.24.6	Windows Azure and Amazon EC2 virtual machines.....	288
16	Disaster recovery	289
17	About Cyber Disaster Recovery Cloud	290
18	Software requirements	291
19	Setting up the disaster recovery functionality	293
20	Setting up connectivity.....	294
20.1.1	Networking concepts.....	294
20.1.2	Initial connectivity configuration.....	301
20.1.3	Network management.....	304
21	Setting up recovery servers	312
21.1.1	How failover and failback work.....	312
21.1.2	Recovery server lifecycle	312
21.1.3	Creating a recovery server.....	314
21.1.4	Performing a test failover	316
21.1.5	Performing a failover	317
21.1.6	Performing a failback	319
21.1.7	Working with encrypted backups	320

22 Setting up primary servers	321
22.1.1 Creating a primary server	321
22.1.2 Operations with a primary server	322
23 Managing the cloud servers	323
23.1 Backing up the cloud servers	324
24 Orchestration (runbooks)	325
24.1.1 Creating a runbook	325
24.1.2 Operations with runbooks	326
25 Antimalware and web protection	328
25.1 Antivirus and Antimalware protection	328
25.1.1 Antivirus and Antimalware protection settings	329
25.2 Active Protection	335
25.3 Windows Defender Antivirus	336
25.4 Microsoft Security Essentials	338
25.5 URL filtering	338
25.6 Quarantine	345
25.7 Corporate allowlist	346
25.8 Antimalware scan of backups	347
26 Protection of collaboration and communication applications	348
27 Vulnerability assessment and patch management	349
27.1 Supported Microsoft and third-party products	349
27.2 Vulnerability assessment	350
27.2.1 Vulnerability assessment settings	351
27.2.2 Managing found vulnerabilities	352
27.2.3 Vulnerability assessment for Linux machines	353
27.3 Patch management	354
27.3.1 Patch management settings	356
27.3.2 Managing list of patches	358
27.3.3 Automatic patch approval	360
27.3.4 Manual patch approval	362
27.3.5 On-demand patch installation	362
27.3.6 Patch lifetime in the list	363
28 Remote desktop access	364
28.1 Remote access (RDP and HTML5 clients)	364
28.2 Sharing remote connection with end-users	368
29 Remote wipe	369
30 Smart protection	370
30.1 Threat feed	370
30.2 Data protection map	372
30.2.1 Data protection map settings	373
31 The Plans tab	374
31.1 Protection plan	375

31.2	Backup scanning plan	377
31.3	Backup plans for cloud applications	378
32	Monitoring	379
32.1	Cyber Protection	380
32.2	Protection status.....	380
32.3	#CyberFit Score by machine	381
32.4	Disk health forecast	382
32.5	Data protection map.....	386
32.6	Vulnerability assessment widgets	387
32.7	Patch installation widgets	388
32.8	Backup scanning details.....	390
32.9	Recently affected	391
32.10	Cloud applications.....	391
33	Reports.....	393
34	Troubleshooting	396
35	Glossary	397

1 Welcome to Cyber Protection

Cyber Protection represents an all-in-one cyber protection solution that integrates backup and recovery, disaster recovery, malware prevention, security controls, remote assistance, monitoring, and reporting.

It protects your entire business and businesses of your customers through layered protection approach, an innovative combination of proactive, active, and reactive data protection technologies:

- **Proactive** actions, such as vulnerability assessment/patch management, predictive analysis of hard drive health based on machine learning technologies, allow you to prevent any threats to your machines.
- **Active** actions, such as protection against malware, self-protection, allow you to detect threats.
- **Reactive** actions, such as backup and recovery (on-premises and cloud), disaster recovery (on-premises and cloud), allow you to respond to any failures.

Cyber Protection provides you with one protection agent, one service console that is easy to manage, and one protection plan that covers all security and data protection aspects.

Key functionality



Cyber Protection provides the following functionality:

- Backup and recovery allows you to back up and recover physical machines, virtual machines, and applications.

- Disaster recovery allows you to protect your local environment from disasters by launching the exact copies of machines in the cloud and switching the workload to the cloud servers.
- Antimalware and web protection provides you with the top multi-layered protection based on four different antimalware technologies inside. You will also be able to manage Microsoft Security Essentials and Windows Defender Antivirus from the service console. The URL filtering functionality allows you to prevent malicious file download and block access to suspicious web resources.
- Autodiscovery of machines provides you with an easy and automatic way to register a large number of machines and install a protection agent and additional components.
- Vulnerability assessment allows you to scan both Microsoft and third-party products on vulnerabilities.
- Patch management integrated with backup provides you with the following capabilities: automatic and manual patch approval, scheduled and on-demand patch installation, flexible reboot and maintenance window options, staged deployment.
- Hard drive health control allows you to track the status of hard disks and prevent their crash. Disk control uses a combination of machine learning and SMART reporting to predict disk failures.
- Remote management and assistance allows you to connect and manage machines remotely.
- #CyberFit Score provides you with a security assessment and scoring mechanism that evaluates the security posture of your machine.

Why Cyber Protection is special

Cyber Protection has the following unique features:

- Backup scanning in non-endpoint environments that ensures malware-free restores. It increases the potential of rootkits and bootkits detection and reduces loads on your machines.
- Safe recovery based on integrated antimalware scanning and malware deletion to prevent recurrence of an infection.
- Smart protection built on the basis of alarms received from Cyber Protection Operations Center (CPOC). This feature allows you to minimize business process downtime when facing issues like a malware attack, natural disaster, to reduce reaction time, and to avoid data loss.
- Protection from bad patches by creating pre-update backups.
- Continuous data protection that ensures that you will not lose your data changes made between scheduled backups. You can control what will be continuously backed up – Office documents, financial forms, graphic files, and so on. You get improved RPOs due to continuous backups.
- A data protection map that ensures tracking data distribution across the machines, monitoring the protection status of files, and using the collected data as the basis for compliance reports.
- Forensic backup that allows you to collect digital evidence data, include them in a disk-level backup, and use them for future investigations.
- Company-wide allowlist built on backups that allows you to prevent false detections. This feature eliminates time-consuming manual allowlisting of corporate applications, ensures greater productivity, and improves the detection rate by using improved heuristics.

2 Cyber Protection editions

To match the needs and budgets of various customers, the Cyber Protection service has different editions. A customer company can use one edition at a time.

Edition	Description
Cyber Backup – Standard	Provides: ▪ Backup and recovery functionality that covers small environment needs ▪ Vulnerability assessment, basic ransomware, and cryptomining protection functionality ▪ Remote installation functionality
Cyber Protect – Essentials	Provides: ▪ Basic file-level backup capabilities ▪ Vulnerability assessment ▪ Basic auto-discovery and remote agent installation ▪ Patch management ▪ Antivirus and antimalware protection ▪ URL filtering ▪ Remote desktop ▪ Remote device wipe ▪ Windows Defender Antivirus and Microsoft Security Essentials management
Cyber Protect – Standard	Provides: ▪ Backup and recovery functionality that covers small environment needs ▪ Remote installation functionality ▪ Vulnerability assessment and patch management functionality ▪ Advanced antimalware protection and web protection functionality ▪ Remote desktop functionality ▪ Security controls functionality such as Windows Defender management ▪ Alarms based on data from Acronis Cyber Protection Operations Center ▪ Data discovery functionality ▪ #CyberFit Score

Cyber Protect – Advanced	Provides: ▪ Backup and recovery functionality dedicated to protect advanced workloads such as Microsoft Exchange and Microsoft SQL clusters designed for big environments ▪ Group management and plan management ▪ Remote installation functionality ▪ Vulnerability assessment and patch management functionality ▪ Advanced antimalware protection and web protection functionality ▪ Remote desktop functionality ▪ Security controls functionality, such as Windows Defender management ▪ Alarms based on data from Acronis Cyber Protection Operations Center ▪ Data discovery functionality ▪ #CyberFit Score

Disaster Recovery add-on

The Disaster Recovery add-on provides disaster recovery functionality designed for companies that have high requirements for the Recovery Time Objective. The add-on can be used with the following editions:

- Cyber Backup Standard
- Cyber Protect Standard
- Cyber Protect Advanced

The following backup and recovery functionalities are available only in the Advanced edition:

- A new section in the UI shows all protection plans and VM replication plans (p. 374)
- Support for Microsoft SQL Server Always On Availability Groups (AAG) (p. 206)
- Support for Microsoft Exchange Server Database Availability Groups (DAG) (p. 208)
- The capability to create static and dynamic groups of devices (p. 87)
- Storing each machine's backups in a folder defined by a script (for machines running Windows) (p. 123)
- Protecting Oracle Database by application-aware backup and database backup that employs RMAN (p. 264)
- Protecting SAP HANA by an entire machine backup that employs the internal SAP HANA snapshot (p. 264)
- File notarization provides proof that a file is authentic and unchanged since it was backed up (p. 140)
- ASign allows a backed-up file to be signed by several people (p. 189)

The disaster recovery functionality is described in the "Disaster Recovery" (p. 289) chapter.

3 Comparison of editions

The differences between the different editions of the Cyber Protection service are shown below.

Features	Cyber Protect			Cyber Backup
	Essentials	Standard	Advanced	Standard
Protection features				
Vulnerability assessment (p. 350)	✓	✓	✓	✓
Active protection (p. 335)	✓	✓	✓	✓
Antimalware protection and web protection (p. 328)	✓	✓	✓	–
Security controls functionality such as Windows Defender Antivirus management (p. 336)	–	✓	✓	–
Alarms based on data from Cyber Protection Operations Center (p. 370)	✓	✓	✓	–
Data protection map (p. 372)	–	✓	✓	–
Patch management (p. 354)	✓	✓	✓	–
Remote desktop access (p. 364)	✓	✓	✓	–
Antimalware backup scanning at a centralized cloud location (p. 347)	–	–	✓	–
Disk health forecast (p. 382)	–	✓	✓	–
#CyberFit Score (p. 103)	✓	✓	✓	–
Backup features				
Backup and recovery (p. 109)	Small environment (up to 5)	Small environment (up to 5)	Big environment (more than 5)	Small environment (up to 5)
Backup to cloud storage	✓	✓	✓	✓
Backup to local storage	–	✓	✓	✓
Network shares backup	–	✓	✓	✓
Continuous data protection (CDP) (p. 118)	–	✓	✓	–
Forensic backups (p. 154)	–	–	✓	–
Safe recovery (p. 175)	–	–	✓	–
Standard applications backup (Microsoft 365, G Suite, MS Exchange, MS SQL, MS Sharepoint)	–	✓	✓	✓
Advanced workloads protection (for example, MS Exchange, MS SQL cluster) (p. 202)	–	–	✓	–

Management features				
Remote installation (p. 56)	✓	✓	✓	✓
Group management (p. 87) and plan management (p. 374)	✓	✓	✓	–
Disaster recovery features				
Disaster recovery (p. 289)	–	Add-on	Add-on	Add-on

4 Software requirements

4.1 Supported Cyber Protect features by operating system

The Cyber Protect features are supported on the following operating systems:

- Windows: Windows 7 and later, Windows 2008 R2 and later.
Windows Defender Antivirus management is supported on Windows 8.1 and later.
- Linux: CentOS 7.x, CentOS 8.0, Virtuozzo 7.x, Acronis Cyber Infrastructure 3.x.
Other Linux distributions and versions might be supported, but have not been tested.
- macOS: 10.13.x and later (only Antivirus and Antimalware protection is supported).

Important: The Cyber Protect features are only supported for machines on which a protection agent is installed. For virtual machines protected in agentless mode, for example by Agent for Hyper-V, Agent for VMware, or Agent for Virtuozzo Hybrid Infrastructure, only backup is supported.

Cyber Protect features	Windows	Linux	macOS
Default protection plans			
Remote Workers	Yes	No	Yes
Office Workers (third-party antivirus)	Yes	No	Yes
Office Workers (Cyber Protectantivirus)	Yes	No	Yes
Cyber Protect Essentials (only for Cyber Protect Essentials edition)	Yes	No	No
Forensic backup			
Collecting memory dump	Yes	No	No
Snapshot of running processes	Yes	No	No
Forensic backup for machines with one drive without reboot	Yes	No	No
Notarization of local image forensic backup	Yes	No	No
Notarization of cloud image forensic backup	Yes	No	No
Continuous data protection (CDP)			
CDP for files and folders	Yes	No	No
CDP for changed files via application tracking	Yes	No	No
Autodiscovery and remote installation			
Network-based discovery	Yes	No	No

Active Directory-based discovery	Yes	No	No
Template-based discovery (importing machines from a file)	Yes	No	No
Manual adding of devices	Yes	No	No
Active Protection			
Ransomware detection	Yes	No	No
Cryptomining processes detection	Yes	No	No
Process Injects detection	Yes	No	No
Automatic recovery of affected files from the local cache	Yes	No	No
Self-defense for Acronis backup files	Yes	No	No
Self-defense for Acronis software	Yes	No	No
Trusted/blocked process management	Yes	No	No
Processes/folders exclusions	Yes	No	No
Ransomware detection based on a process behavior (AI-based)	Yes	No	No
Cryptomining detection based on a process behavior	Yes	No	No
External drives protection (HDD, flash drives, SD cards)	Yes	No	No
Network folder protection	Yes	No	No
Server-side protection	Yes	No	No
Zoom, Cisco Webex, Citrix Workspace, and Microsoft Teams protection	Yes	No	No
Antivirus and Antimalware protection			
Fully-integrated Active Protection functionality	Yes	No	No
Real-time antimalware protection	Yes	No	Yes
Static analysis for portable executable files	Yes	No	Yes*
On-demand antimalware scanning	Yes	No	Yes
Network folder protection	Yes	No	No
Server-side protection	Yes	No	No
Scan of archive files	Yes	No	Yes
Scan of removable drives	Yes	No	Yes
Scan of only new and changed files	Yes	No	Yes
File/folder exclusions	Yes	No	Yes**
Processes exclusions	Yes	No	No
Behavioral analysis engine	Yes	No	No
Exploit prevention	Yes	No	No
Quarantine	Yes	No	Yes
Quarantine auto clean-up	Yes	No	Yes
URL filtering (http/https)	Yes	No	No

Corporate-wide allowlist	Yes	No	Yes
Windows Defender Antivirus management	Yes	No	No
Microsoft Security Essentials management	Yes	No	No
Registering and managing Antivirus and Antimalware protection via Windows Security Center	Yes	No	No
Exploit prevention in antivirus and antimalware protection	Yes	No	No
Vulnerability and configuration assessment			
Vulnerability assessment for Windows	Yes	No	No
Vulnerability assessments of Cyber Infrastructure (Linux)***	No	Yes	No
Vulnerability assessment for 3rd-party Windows applications	Yes	No	No
Patch management			
Patch auto-approval	Yes	No	No
Patch auto-installation	Yes	No	No
Patch testing	Yes	No	No
Manual patch installation	Yes	No	No
Patch scheduling	Yes	No	No
Fail-safe patching: backup of machine before installing patches as part of protection plan	Yes	No	No
Cancellation of a machine reboot if a backup is running	Yes	No	No
Data protection map			
Adjustable definition of important files	Yes	No	No
Scanning machines to find unprotected files	Yes	No	No
Unprotected locations overview	Yes	No	No
Ability to start the protection action from the Data protection map widget (Protect all files action)	Yes	No	No
Disk health			
AI-based HDD and SSD health control	Yes	No	No
Smart protection plans based on Acronis Cyber Protection Operations Center (CPOC) alerts			
Threat feed	Yes	No	No
Remediation wizard	Yes	No	No
Backup scanning			
Antimalware scan of image backups as part of backup plan	Yes	No	No
Scanning of image backups for malware in cloud	Yes	No	No
Malware scan of encrypted backups	Yes	No	No
Safe recovery			
Antimalware scanning with Antivirus and Antimalware protection during the recovery process	Yes	No	No
Safe recovery for encrypted backups	Yes	No	No

Remote desktop connection			
Connection via HTML5-based client	Yes	No	No
Connection via native Windows RDP client	Yes	No	No
Remote assistance	Yes	No	No
#CyberFit Score			
#CyberFit Score status	Yes	No	No
#CyberFit Score standalone tool	Yes	No	No
#CyberFit Score recommendations	Yes	No	No
Management options			
Upsell scenarios to promote Cyber Protect editions	Yes	Yes	Yes
Web-based centralized and remote management console	Yes	Yes	Yes
Protection options			
Remote wipe (Windows 10 only)	Yes	No	No
Cyber Protect Monitor			
Cyber Protect Monitor app	Yes	No	Yes
Protection status for Zoom	Yes	No	No
Protection status for Cisco Webex	Yes	No	No
Protection status for Citrix Workspace	Yes	No	No
Protection status for Microsoft Teams	Yes	No	No

* Static analysis for portable executable files is supported only for scheduled scans on macOS.

** File/folder exclusions are only supported for the case when you specify files and folders that will not be scanned by real-time protection or scheduled scans on macOS.

*** The vulnerability assessment depends on the availability of official security advisories for specific distribution, for example <https://lists.centos.org/pipermail/centos-announce/>, <https://lists.centos.org/pipermail/centos-announce/>, <https://lists.centos.org/pipermail/centos-cr-announce/>, <https://lists.centos.org/pipermail/centos-cr-announce/>, and others.

4.2 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 11 or later
- Microsoft Edge 25 or later
- Safari 8 or later running in the macOS 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.

4.3 Supported operating systems and environments

Agent for Windows

Windows XP Professional SP1 (x64), SP2 (x64), SP3 (x86)
Windows Server 2003 SP1/2003 R2 and later – Standard and Enterprise editions (x86, x64)
Windows Small Business Server 2003/2003 R2
Windows Vista – all editions
Windows Server 2008 – Standard, Enterprise, Datacenter, Foundation, and Web editions (x86, x64)
Windows Small Business Server 2008
Windows 7 – all editions
Windows Server 2008 R2 – Standard, Enterprise, Datacenter, Foundation, and Web editions
Windows Home Server 2011
Windows MultiPoint Server 2010/2011/2012
Windows Small Business Server 2011 – all editions
Windows 8/8.1 – all editions (x86, x64), except for the Windows RT editions
Windows Server 2012/2012 R2 – all editions
Windows Storage Server 2003/2008/2008 R2/2012/2012 R2/2016
Windows 10 – Home, Pro, Education, Enterprise, IoT Enterprise and LTSC (formerly LTSB) editions
Windows Server 2016 – all installation options, except for Nano Server
Windows Server 2019 – all installation options, except for Nano Server

Agent for SQL, Agent for Active Directory, Agent for Exchange (for database backup and application-aware backup)

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 Exchange (for mailbox backup)

Windows Server 2008 – Standard, Enterprise, Datacenter, Foundation, and Web editions (x86, x64)
Windows Small Business Server 2008
Windows 7 – all editions
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 (x86, x64), except for the Windows RT editions
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 – all installation options, except for Nano Server
Windows Server 2019 – all installation options, except for Nano Server

Agent for Office 365

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

Windows Small Business Server 2008

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

Windows Home Server 2011

Windows Small Business Server 2011 – all editions

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

Windows Server 2012/2012 R2 – all editions

Windows Storage Server 2008/2008 R2/2012/2012 R2/2016 (x64 only)

Windows 10 – Home, Pro, Education, and Enterprise editions (x64 only)

Windows Server 2016 – all installation options (x64 only), except for Nano Server

Windows Server 2019 – all installation options (x64 only), except for Nano Server

Agent for Oracle

Windows Server 2008R2 – Standard, Enterprise, Datacenter, and Web editions (x86, x64)

Windows Server 2012R2 – Standard, Enterprise, Datacenter, and Web editions (x86, x64)

Linux – any kernel and distribution supported by Agent for Linux (listed below)

Agent for Linux

Linux with kernel from 2.6.9 to 5.1 and glibc 2.3.4 or later, including the following x86 and x86_64 distributions:

Red Hat Enterprise Linux 4.x, 5.x, 6.x, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 8.0*, 8.1*

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, 16.10, 17.04, 17.10, 18.04, 18.10, 19.04, 19.10, 20.04

Fedora 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31

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, 8.6, 8.7, 8.8, 8.11, 9.0, 9.1, 9.2, 9.3, 9.4, 9.5, 9.6, 9.7, 9.8, 10

CentOS 5.x, 6.x, 7, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 8.0, 8.1

Oracle Linux 5.x, 6.x, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 8.0, 8.1 – both Unbreakable Enterprise Kernel and Red Hat Compatible Kernel

CloudLinux 5.x, 6.x, 7, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7

ClearOS 5.x, 6.x, 7, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6

ALT Linux 7.0

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**

* Configurations with Stratis are not supported.

Agent for Mac

OS X Mavericks 10.9

OS X Yosemite 10.10

OS X El Capitan 10.11

macOS Sierra 10.12
macOS High Sierra 10.13
macOS Mojave 10.14
macOS Catalina 10.15

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, 6.0, 6.5, 6.7

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 only) with Hyper-V role, including Server Core installation mode
Windows Server 2008 R2 with Hyper-V role, including Server Core installation mode
Microsoft Hyper-V Server 2008/2008 R2
Windows Server 2012/2012 R2 with Hyper-V role, including Server Core installation mode
Microsoft Hyper-V Server 2012/2012 R2
Windows 8, 8.1 (x64 only) with Hyper-V
Windows 10 – Pro, Education, and Enterprise editions with Hyper-V
Windows Server 2016 with Hyper-V role – all installation options, except for Nano Server
Microsoft Hyper-V Server 2016
Windows Server 2019 with Hyper-V role – all installation options, except for Nano Server
Microsoft Hyper-V Server 2019

Agent for Virtuozzo

Virtuozzo 6.0.10, 6.0.11, 6.0.12, 7.0.13, 7.0.14

Agent for Virtuozzo Hybrid Infrastructure

Virtuozzo Hybrid Infrastructure 3.5

4.4 Supported Microsoft SQL Server versions

- Microsoft SQL Server 2019
- Microsoft SQL Server 2017
- 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

4.5 Supported Microsoft Exchange Server versions

- **Microsoft Exchange Server 2019** – all editions.
- **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. Mailbox backup and granular recovery from database backups are supported starting with Service Pack 1 (SP1).
- **Microsoft Exchange Server 2007** – all editions, all service packs. Mailbox backup and granular recovery from database backups are not supported.

4.6 Supported Microsoft SharePoint versions

Cyber Protection 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.

4.7 Supported Oracle Database versions

- Oracle Database version 11g, all editions
- Oracle Database version 12c, all editions.

Only single-instance configurations are supported.

4.8 Supported SAP HANA versions

HANA 2.0 SPS 03 installed in RHEL 7.6 running on a physical machine or VMware ESXi virtual machine.

Because SAP HANA does not support recovery of multitenant database containers by using storage snapshots, this solution supports SAP HANA containers with only one tenant database.

4.9 Supported virtualization platforms

The following table summarizes how various virtualization platforms are supported.

Platform	Backup at a hypervisor level (agentless backup)	Backup from inside a guest OS
VMware		

Platform	Backup at a hypervisor level (agentless backup)	Backup from inside a guest OS
VMware vSphere versions: 4.1, 5.0, 5.1, 5.5, 6.0, 6.5, 6.7 VMware vSphere editions: VMware vSphere Essentials* VMware vSphere Essentials Plus* VMware vSphere Standard* VMware vSphere Advanced VMware vSphere Enterprise VMware vSphere Enterprise Plus	+	+
VMware vSphere Hypervisor (Free ESXi)**		+
VMware Server (VMware Virtual server) VMware Workstation VMware ACE VMware Player		+
Microsoft		
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 with Hyper-V Windows Server 2016 with Hyper-V – all installation options, except for Nano Server Microsoft Hyper-V Server 2016 Windows Server 2019 with Hyper-V – all installation options, except for Nano Server Microsoft Hyper-V Server 2019	+	+
Microsoft Virtual PC 2004 and 2007 Windows Virtual PC		+
Microsoft Virtual Server 2005		+
Citrix		
Citrix XenServer 4.1.5, 5.5, 5.6, 6.0, 6.1, 6.2, 6.5, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5		Only fully virtualized (aka HVM) guests. Paravirtualized (aka PV) guests are not supported.
Red Hat and Linux		

Platform	Backup at a hypervisor level (agentless backup)	Backup from inside a guest OS
Red Hat Enterprise Virtualization (RHEV) 2.2, 3.0, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6 Red Hat Virtualization (RHV) 4.0, 4.1		+
Kernel-based Virtual Machines (KVM)		+
Parallels		
Parallels Workstation		+
Parallels Server 4 Bare Metal		+
Oracle		
Oracle VM Server 3.0, 3.3, 3.4		Only fully virtualized (aka HVM) guests. Paravirtualized (aka PV) guests are not supported.
Oracle VM VirtualBox 4.x		+
Nutanix		
Nutanix Acropolis Hypervisor (AHV) 20160925.x through 20180425.x		+
Virtuozzo		
Virtuozzo 6.0.10, 6.0.11, 6.0.12	+	Virtual machines only. Containers are not supported.
Virtuozzo 7.0.13, 7.0.14	Ploop containers only. Virtual machines are not supported.	Virtual machines only. Containers are not supported.
Virtuozzo Hybrid Infrastructure		
Virtuozzo Hybrid Infrastructure 3.5	+	+
Amazon		
Amazon EC2 instances		+
Microsoft Azure		
Azure virtual machines		+

* In these editions, the HotAdd transport for virtual disks is supported on vSphere 5.0 and later. On version 4.1, backups may run slower.

** Backup at a hypervisor level is not supported for vSphere Hypervisor because this product restricts access to Remote Command Line Interface (RCLI) to read-only mode. The agent works during the vSphere Hypervisor evaluation period while no serial key is entered. Once you enter a serial key, the agent stops functioning.

Limitations

▪ Fault tolerant machines

Agent for VMware backs up a fault tolerant machine only if fault tolerance was enabled in VMware vSphere 6.0 and later. If you upgraded from an earlier vSphere version, it is enough to disable and enable fault tolerance for each machine. If you are using an earlier vSphere version, install an agent in the guest operating system.

- **Independent disks and RDM**

Agent for VMware does not back up Raw Device Mapping (RDM) disks in physical compatibility mode or independent disks. The agent skips these disks and adds warnings to the log. You can avoid the warnings by excluding independent disks and RDMs in physical compatibility mode from the protection plan. If you want to back up these disks or data on these disks, install an agent in the guest operating system.

- **Pass-through disks**

Agent for Hyper-V does not back up pass-through disks. During backup, the agent skips these disks and adds warnings to the log. You can avoid the warnings by excluding pass-through disks from the protection plan. If you want to back up these disks or data on these disks, install an agent in the guest operating system.

- **Hyper-V guest clustering**

Agent for Hyper-V does not support backup of Hyper-V virtual machines that are nodes of a Windows Server Failover Cluster. A VSS snapshot at the host level can even temporarily disconnect the external quorum disk from the cluster. If you want to back up these machines, install agents in the guest operating systems.

- **In-guest iSCSI connection**

Agent for VMware and Agent for Hyper-V do not back up LUN volumes connected by an iSCSI initiator that works within the guest operating system. Because the ESXi and Hyper-V hypervisors are not aware of such volumes, the volumes are not included in hypervisor-level snapshots and are omitted from a backup without a warning. If you want to back up these volumes or data on these volumes, install an agent in the guest operating system.

- **Linux machines containing logical volumes (LVM)**

Agent for VMware and Agent for Hyper-V do not support the following operations for Linux machines with LVM:

- P2V migration, V2P migration, and V2V migration from Virtuozzo. Use Agent for Linux to create the backup and bootable media to recover.
- Running a virtual machine from a backup created by Agent for Linux.

- **Encrypted virtual machines** (introduced in VMware vSphere 6.5)

- Encrypted virtual machines are backed up in an unencrypted state. If encryption is critical to you, enable encryption of backups when creating a protection plan (p. 138).
- Recovered virtual machines are always unencrypted. You can manually enable encryption after the recovery is complete.
- If you back up encrypted virtual machines, we recommend that you also encrypt the virtual machine where Agent for VMware is running. Otherwise, operations with encrypted machines may be slower than expected. Apply the **VM Encryption Policy** to the agent's machine by using vSphere Web Client.
- Encrypted virtual machines will be backed up via LAN, even if you configure the SAN transport mode for the agent. The agent will fall back on the NBD transport because VMware does not support SAN transport for backing up encrypted virtual disks.

- **Secure Boot**

- VMware virtual machines: (introduced in VMware vSphere 6.5) **Secure Boot** is disabled after a virtual machine is recovered as a new virtual machine. You can manually enable this option after the recovery is complete. This limitation applies to VMware.
- Hyper-V virtual machines: For all GEN2 VMs, Secure Boot is disabled after the virtual machine is recovered to both new virtual machine or an existing virtual machine.

- **ESXi configuration backup** is not supported for VMware vSphere 6.7.

4.10 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 protection 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.

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>

5 Supported file systems

A protection 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 (bootable media supports only recovery). The limitations apply to both the agents and bootable media.

File system	Supported by			Limitations
	Agents	Bootable media for Windows and Linux	Bootable media for Mac	
FAT16/32	All agents	+	+	No limitations
NTFS		+	+	
ext2/ext3/ext4		+	-	
HFS+	Agent for Mac	-	+	- Supported starting with macOS High Sierra 10.13 - Disk configuration should be re-created manually when recovering to a non-original machine or bare metal.
APFS		-	+	
JFS	Agent for Linux	+	-	- Files cannot be excluded from a disk backup - Fast incremental/differential backup cannot be enabled
ReiserFS3		+	-	
ReiserFS4		+	-	- Files cannot be excluded from a disk backup - Fast incremental/differential backup cannot be enabled - Volumes cannot be resized during a recovery
ReFS	All agents	+	+	
XFS		+	+	
Linux swap	Agent for Linux	+	-	No limitations

File system	Supported by			Limitations
	Agents	Bootable media for Windows and Linux	Bootable media for Mac	
exFAT	All agents	+ Bootable media cannot be used for recovery if the backup <i>is stored on exFAT</i>	+	- Only disk/volume backup is supported - Files cannot be excluded from a backup - Individual files cannot be recovered from a backup

The software automatically switches to the sector-by-sector mode when backing up drives with unrecognized or unsupported file systems (for example, Btrfs). 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.

Data Deduplication

In Windows Server 2012 and later, you can enable the Data Deduplication feature for an NTFS volume. Data Deduplication reduces the used space on the volume by storing duplicate fragments of the volume's files only once.

You can back up and recover a data deduplication-enabled volume at a disk level, without limitations. File-level backup is supported, except when using Acronis VSS Provider. To recover files from a disk backup, either run a virtual machine (p. 268) from your backup, or mount the backup (p. 200) on a machine running Windows Server 2012 or later, and then copy the files from the mounted volume.

The Data Deduplication feature of Windows Server is unrelated to the Acronis Backup Deduplication feature.

6 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. Ensure that your password is at least eight characters long. Remember your login that is shown on the account activation page.
If your administrator has enabled two-factor authentication, you will be prompted to set up two-factor authentication for your account (p. 27).

- **A link to the service console login page.** Use this link to access the console in the future. The login and password are the same as in the previous step.

6.1 Two-factor authentication

Two-factor authentication provides extra protection from unauthorized access to your account. When two-factor authentication is set up, you are required to enter your password (the first factor) and a one-time code (the second factor) to log in to the service console. The one-time code is generated by a special application that must be installed on your mobile phone or another device that belongs to you. Even if someone finds out your login and password, they still will not be able to login without access to your second-factor device.

The one-time code is generated based on the device's current time and the secret provided by the Cyber Protection service as the QR code or alphanumeric code. During the first login, you need to enter this secret to the authentication application.

To set up two-factor authentication for your account

1. Choose the second-factor device.
Most commonly it is a mobile phone, but you can also use a tablet, laptop, or desktop.
2. Ensure that the device time settings are correct and reflect the actual current time. Ensure that the device locks itself after a period of inactivity.
3. Install the authentication application on the device. The recommended applications are Google Authenticator or Microsoft Authenticator.
4. Go to the service console login page and set your password.
The service console shows the QR code and the alphanumeric code.
5. Save the QR code and the alphanumeric code in any convenient way (such as, print out the screen, write down the code, or save the screenshot in cloud storage). If you lose the second-factor device, you will be able to reset the two-factor authentication by using these codes.
6. Open the authentication application, and then do one of the following:
 - Scan the QR code
 - Manually enter the alphanumeric code to the applicationThe authentication application generates a one-time code. A new code will be generated every 30 seconds.
7. Return to the service console login page and enter the generated code.
A one-time code is valid for 30 seconds. If you wait longer than 30 seconds, use the next generated code.

When logging in the next time, you can select the checkbox **Trust this browser....** If you do this, the one-time code will not be required when you log in by using this browser on this machine.

What if...

...I lost the second-factor device?

If you have a trusted browser, you will be able to log in by using this browser. Nevertheless, when you have a new device, repeat steps 1-3 and 6-7 of the above procedure by using the new device and the saved QR code or alphanumeric code.

If you have not saved the code, ask the administrator or service provider to reset the two-factor authentication for your account, and then repeat steps 1-3 and 6-7 of the above procedure by using the new device.

...I want to change the second-factor device?

When logging in, click the **Reset two-factor authentication settings** link, confirm the operation by entering the one-time code, and then repeat the above procedure by using the new device.

7 Accessing the Cyber Protection service

You can log in to the Cyber Protection service if you activated your account.

To log in to the Cyber Protection service

1. Go to the Cyber Protection service login page. The login page address was included in the activation email message.
2. Type the login, and then click **Next**.
3. Type the password, and then click **Next**.
4. If you have the administrator role in the Cyber Protection service, click **Cyber Protection**.
Users who do not have the administrator role log in directly to the service console.

The timeout period for the service console is 24 hours for active sessions and 1 hour for idle sessions.

To reset your password

1. Go to the Cyber Protection service login page.
2. Type your login, and then click **Next**.
3. Click **Forgot password?**
4. Confirm that you want further instructions by clicking **Send**.
5. Follow the instructions in the email that you have received.
6. Set up your new password. Ensure that your password is at least eight characters long.

You can change the language of the web interface by clicking the account icon in the top-right corner.

If **Cyber Protection** is not the only service you are subscribed to, you can switch between the services by using the  icon in the top-right corner. Administrators can also use this icon for switching to the management portal.

If you are subscribed to any of the Cyber Protection editions, you can send feedback about the product from the service console. In the left navigation menu, click **Send feedback**, fill in the fields, attach files (if any) and click **Send**.

8 Installing the software

Related topics

Selecting components for installation 61

8.1 Preparation

Step 1

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

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

In Linux, Agent for Oracle and Agent for Virtuozzo require that Agent for Linux (64-bit) is also installed. These three agents share one installer.

What are you going to back up?	Which agent to install?	Where to install it?
Physical machines		
Physical machines running Windows	Agent for Windows	On the machine that will be backed up.
Physical machines running Linux	Agent for Linux	
Physical machines running macOS	Agent for Mac	
Applications		
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.*
Microsoft Office 365 mailboxes	Agent for Office 365	On a Windows machine that is connected to the Internet. Depending on the desired functionality, you may or may not need to install Agent for Office 365. For more information, refer to "Protecting Office 365 data" (p. 230).
Microsoft Office 365 OneDrive files and SharePoint Online sites	—	This data can be backed up only by an agent that is installed in the cloud. For more information, refer to "Protecting Office 365 data" (p. 230).
G Suite Gmail mailboxes, Google Drive files, and Shared drive files	—	This data can be backed up only by an agent that is installed in the cloud. For more information, refer to "Protecting G Suite" (p. 252).
Machines running Active Directory Domain Services	Agent for Active Directory	On the domain controller.
Machines running Oracle Database	Agent for Oracle	On the machine running Oracle Database.
Virtual machines		
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.**

What are you going to back up?	Which agent to install?	Where to install it?
	Agent for VMware (Virtual Appliance)	On the ESXi host.
Hyper-V virtual machines	Agent for Hyper-V	On the Hyper-V host.
Virtuozzo virtual machines and containers***	Agent for Virtuozzo	On the Virtuozzo host.
Virtual machines hosted on Amazon EC2	The same as for physical machines****	On the machine that will be backed up.
Virtual machines hosted on Windows Azure		
Citrix XenServer virtual machines		
Red Hat Virtualization (RHV/RHEV)		
Kernel-based Virtual Machines (KVM)		
Oracle virtual machines		
Nutanix AHV virtual machines		
Mobile devices		
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	

*During the installation, Agent for Exchange checks for enough free space on the machine where it will run. Free space equal to 15 percent of the biggest Exchange database is temporarily needed during a granular recovery.

**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. For detailed instructions, refer to "Agent for VMware - LAN-free backup" (p. 276).

***For Virtuozzo 7, only ploop containers are supported. Virtual machines are not supported.

****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.

Step 2

Review the system requirements for the agents.

Agent	Disk space occupied by the agent(s)
Agent for Windows	550 MB
Agent for Linux	2 GB
Agent for Mac	450 MB

Agent for SQL	600 MB (50 MB + 550 MB Agent for Windows)
Agent for Exchange	750 MB (200 MB + 550 MB Agent for Windows)
Agent for Office 365	550 MB
Agent for Active Directory	600 MB (50 MB + 550 MB Agent for Windows)
Agent for VMware	700 MB (150 MB + 550 MB Agent for Windows)
Agent for Hyper-V	600 MB (50 MB + 550 MB Agent for Windows)
Agent for Virtuozzo	500 MB
Agent for Oracle	450 MB

The typical memory consumption is 300 MB above the operating system and running applications. The peak consumption may reach 2 GB, depending on the amount and type of data being processed by the agents.

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

Step 3

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.

To download Agent for Office 365 setup program, click the account icon in the top-right corner, and then click **Downloads > Agent for Office 365**.

Installation in Linux and macOS is performed from ordinary setup programs.

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

Step 4

Cyber Protect features require Microsoft Visual C++ 2017 Redistributable. Please ensure that it is already installed on your machine or install it before installing the agent. After the installation of Microsoft Visual C++, a restart may be required. You can find the Microsoft Visual C++ Redistributable package here

<https://support.microsoft.com/help/2999226/update-for-universal-c-runtime-in-windows>.

Step 5

Verify 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 service 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 and 55556** The agents use these ports for data transfer during backup and recovery.

If a proxy server is enabled in your network, refer to the "Proxy server settings" (p. 35) section to understand whether you need to configure these settings on each machine that runs a protection agent.

The minimum Internet connection speed required for managing an agent from the cloud is 1 Mbit/s (not to be confused with the data transfer rate acceptable for backing up to the cloud). Consider this if you use a low-bandwidth connection technology such as ADSL.

TCP ports required for backup and replication of VMware virtual machines

- **TCP 443** Agent for VMware (both Windows and Virtual Appliance) connects to this port on the ESXi host/vCenter server to perform VM management operations, such as create, update, and delete VMs on vSphere during backup, recovery, and VM replication operations.
- **TCP 902** Agent for VMware (both Windows and Virtual Appliance) connects to this port on the ESXi host to establish NFC connections to read/write data on VM disks during backup, recovery, and VM replication operations.
- **TCP 3333** If the Agent for VMware (Virtual Appliance) is running on the ESXi host/cluster that is the target for VM replication, VM replication traffic does not go directly to the ESXi host on port 902. Instead, the traffic goes from the source Agent for VMware to TCP port 3333 on the Agent for VMware (Virtual Appliance) located on the target ESXi host/cluster.

The source Agent for VMware that reads data from the original VM disks can be anywhere else and can be of any type: Virtual Appliance or Windows.

The service that is responsible for accepting VM replication data on the target Agent for VMware (Virtual Appliance) is called "Replica disk server." This service is responsible for the WAN optimization techniques, such as traffic compression and deduplication during VM replication, including replica seeding (see Seeding an initial replica (p. 275)). When no Agent for VMware (Virtual Appliance) is running on the target ESXi host, this service is not available, and therefore the replica seeding scenario is not supported.

Ports required by the Downloader component

The Downloader component is responsible for delivering updates to a computer and distributing them to other Downloader instances. It can run in agent mode which turns its computer into Downloader agent. The Downloader agent downloads updates from the internet and servers as the source of updates distribution to other computers. The Downloader requires the following ports to operate.

- **6888** Used by BitTorrent protocol for torrent peer to peer updates.
- **6771** Used as the local peer discovery port. Also takes part in peer to peer updates.
- **18018** Used for communication between updaters working in different modes: Updater and UpdaterAgent.
- **18019** Local port, used for communication between the Updater and the Acronis Cyber Protection agent.

Step 6

On the machine where you plan to install the Cyber Protection agent, verify that the following local ports are not in use by other processes.

- 127.0.0.1:9999
- 127.0.0.1:43234
- 127.0.0.1:9850

Note You do not have to open them in the Firewall.

Changing the ports used by the Cyber Protection agent

Some of the ports required by the Cyber Protection agent might be in use by other applications in your environment. To avoid conflicts, you can change the default ports used by the Cyber Protection agent by modifying the following files.

- In Linux: `/opt/Acronis/etc/aakore.yaml`
- In Windows: `\ProgramData\Acronis\Agent\etc\aakore.yaml`

8.2 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 **libelf-dev**, **libelf-devel**, or **elfutils-libelf-devel** libraries for building kernels starting with 4.15 and configured with `CONFIG_UNWINDER_ORC=y`. For some distributions, such as Fedora 28, they need to be installed separately from kernel headers.

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.

5. In Red Hat Enterprise Linux, CentOS, and Fedora, run the following command to check whether **elfutils-libelf-devel** is installed:

```
yum list installed | grep elfutils-libelf-devel
```

If you see the information about the library version, the library 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 elfutils-libelf-devel	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 elfutils-libelf-devel	The setup program will download and install the packages automatically.
	perl	Run the following command: <pre>yum install perl</pre>
Ubuntu Debian	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>
SUSE Linux OpenSUSE	kernel-source gcc make perl	<pre>sudo zypper install kernel-source sudo zypper install gcc sudo zypper install make sudo zypper 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.

8.3 Proxy server settings

The protection agents can transfer data through an HTTP/HTTPS proxy server. The server must work through an HTTP tunnel without scanning or interfering with the HTTP traffic. Man-in-the-middle proxies are not supported.

Because the agent registers itself in the cloud during the installation, the proxy server settings must be provided during the installation or in advance.

In Windows

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. Also, you can enter the proxy settings during the installation, or specify them in advance by using the procedure described below. To change the proxy settings after the installation, use the same procedure.

To specify the proxy settings 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
"Login"="proxy_login"
"Password"="proxy_password"
```

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. If your proxy server requires authentication, replace `proxy_login` and `proxy_password` with the proxy server credentials. Otherwise, delete these lines from the file.
5. Save the document as **proxy.reg**.
6. Run the file as an administrator.
7. Confirm that you want to edit the Windows registry.
8. If the protection agent is not installed yet, you can install it now.
9. Open file `%programdata%\Acronis\Agent\etc\akore.yaml` in a text editor.
10. Locate the **env** section or create it and add the following lines:

```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

11. Replace `proxy_login` and `proxy_password` with the proxy server credentials, and `proxy_address:port` with the address and port number of the proxy server.
12. In the **Start** menu, click **Run**, type: **cmd**, and click **OK**.
13. Restart the `akore` service by using the following commands:

```
net stop akore
net start akore
```

14. Restart the agent by using the following commands:

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

In Linux

Run the installation file with the parameters **--http-proxy-host=ADDRESS**

--http-proxy-port=PORT --http-proxy-login=LOGIN

--http-proxy-password=PASSWORD. To change the proxy settings after the installation, use the procedure described below.

To change the proxy settings in Linux

1. Open the file `/etc/Acronis/Global.config` in a text editor.
2. Do one of the following:
 - If the proxy settings were specified during the agent installation, find the following section:

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdwor">"1"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdwor">"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```


- Otherwise, copy the above lines and paste them into the file between the **<registry name="Global">...</registry>** tags.
3. Replace ADDRESS with the new proxy server host name/IP address, and PORT with the decimal value of the port number.
 4. If your proxy server requires authentication, replace LOGIN and PASSWORD with the proxy server credentials. Otherwise, delete these lines from the file.
 5. Save the file.
 6. Open file **/opt/acronis/etc/aakore.yaml** in a text editor.
 7. Locate the **env** section or create it and add the following lines:


```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```
 8. Replace proxy_login and proxy_password with the proxy server credentials, and proxy_address:port with the address and port number of the proxy server.
 9. Restart the aakore service by using the following command:


```
sudo service aakore restart
```
 10. Restart the agent by executing the following command in any directory:


```
sudo service acronis_mms restart
```

In macOS

You can enter the proxy settings during the installation, or specify them in advance by using the procedure described below. To change the proxy settings after the installation, use the same procedure.

To specify the proxy settings in macOS

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>
    <value name="Login" type="TString">"proxy_login"</value>
    <value name="Password" type="TString">"proxy_password"</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. If your proxy server requires authentication, replace proxy_login and proxy_password with the proxy server credentials. Otherwise, delete these lines from the file.
5. Save the file.
6. If the protection agent is not installed yet, you can install it now.
7. Open file **/Library/Application Support/Acronis/Agent/etc/aakore.yaml** in a text editor.
8. Locate the **env** section or create it and add the following lines:


```
env:
  http-proxy: proxy_login:proxy_password@proxy_address:port
  https-proxy: proxy_login:proxy_password@proxy_address:port
```

9. Replace `proxy_login` and `proxy_password` with the proxy server credentials, and `proxy_address:port` with the address and port number of the proxy server.
10. Go to **Applications > Utilities > Terminal**
11. Restart the aakore service by using the following commands:

```
sudo launchctl stop aakore
sudo launchctl start aakore
```

12. Restart the agent by using the following commands:

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

In bootable media

When working under bootable media, you may need to access the cloud storage via a proxy server. To specify the proxy server settings, click **Tools > Proxy server**, and then specify the proxy server host name/IP address, port, and credentials.

8.4 Installing agents

You can install agents on machines running any of the operating systems listed in "Supported operating systems and environments" (p. 17). The operating systems that support the Cyber Protect features are listed in "Supported Cyber Protect features by operating system" (p. 13).

In Windows

Cyber Protect features require Microsoft Visual C++ 2017 Redistributable. Please ensure that it is already installed on your machine or install it before installing the agent. After the installation, a restart may be required. The Microsoft Visual C++ Redistributable package can be found here: <https://support.microsoft.com/help/2999226/update-for-universal-c-runtime-in-windows>.

1. Ensure that the machine is connected to the Internet.
2. Log on as an administrator and start the setup program.
3. [Optional] Click **Customize installation settings** and make the appropriate changes if you want:
 - To change the components to install (in particular, to disable installation of Cyber Protection Monitor and Command-Line Tool).
 - To change the method of registering the machine in the Cyber Protection service. You can switch from **Use service console** (default) to **Use credentials** or **Use registration token**.
 - To change the installation path.
 - To change the user account under which the agent service will run. For details, refer to "Changing the logon account on Windows machines" (p. 41).
 - To verify or change the proxy server host name/IP address, port, and credentials. If a proxy server is enabled in Windows, it is detected and used automatically.
4. Click **Install**.
5. [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, and then click **Done**. We recommend using an account that has the **Administrator** role assigned. Otherwise, provide an account with the necessary privileges (p. 282) on the vCenter Server or ESXi.
6. [Only when installing on a domain controller] Specify the user account under which the agent service will run, and then click **Done**. For security reasons, the setup program does not automatically create new accounts on a domain controller.

7. If you kept the default registration method **Use service console** in step 3, wait until the registration screen appears, and then proceed to the next step. Otherwise, no more actions are required.
8. Do one of the following:
 - Click **Register the machine**. In the opened browser window, sign in to the service console, review the registration details, and then click **Confirm registration**.
 - Click **Show registration info**. The setup program shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. In this case, you will need to enter the registration code in the registration form. The registration code is valid for one hour.

Alternatively, you can access the registration form by clicking **All devices > Add**, scrolling down to **Registration via code**, and then clicking **Register**.

Tip Do not quit the setup program until you confirm the registration. To initiate the registration again, you will have to restart the setup program, and then click **Register the machine**.

As a result, the machine will be assigned to the account that was used to log in to the service console.

- Register the machine manually by using the command line. For more information on how to do this, refer to "Registering machines manually" (p. 53).

In Linux

To install Agent for Linux, you need at least 2 GB of free disk space.

1. Ensure that the machine is connected to the Internet.
2. As the root user, run the installation file.

If a proxy server is enabled in your network, when running the file, specify the server host name/IP address and port in the following format: **--http-proxy-host=ADDRESS**
--http-proxy-port=PORT **--http-proxy-login=LOGIN**
--http-proxy-password=PASSWORD.

If you want to change the default method of registering the machine in the Cyber Protection service, run the installation file with one of the following parameters:

- **--register-with-credentials** – to ask for a user name and password during the installation
 - **--token=STRING** – to use a registration token
 - **--skip-registration** – to skip the registration
3. Select the check boxes for the agents that you want to install. The following agents are available:
 - **Agent for Linux**
 - **Agent for Virtuozzo**
 - **Agent for Oracle**

Agent for Virtuozzo and Agent for Oracle require that Agent for Linux (64-bit) is also installed.

4. If you kept the default registration method in step 2, proceed to the next step. Otherwise, enter the user name and password for the Cyber Protection service, or wait until the machine will be registered by using the token.
5. Do one of the following:
 - Click **Register the machine**. In the opened browser window, sign in to the service console, review the registration details, and then click **Confirm registration**.

- Click **Show registration info**. The setup program shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. In this case, you will need to enter the registration code in the registration form. The registration code is valid for one hour.

Alternatively, you can access the registration form by clicking **All devices > Add**, scrolling down to **Registration via code**, and then clicking **Register**.

Tip Do not quit the setup program until you confirm the registration. To initiate the registration again, you will have to restart the setup program and repeat the installation procedure.

As a result, the machine will be assigned to the account that was used to log in to the service console.

- Register the machine manually by using the command line. For more information on how to do this, refer to "Registering machines manually" (p. 53).
6. If UEFI Secure Boot is enabled on the machine, you are informed that you need to restart the system after the installation. Be sure to remember what password (the one of the root user or "acronis") should be used.

Note During the installation, a new key is generated, used to sign the **snapapi** module, and registered as a Machine Owner Key (MOK). The restart is mandatory in order to enroll this key. Without enrolling the key, the agent will not be operational. If you enable UEFI Secure Boot after the agent installation, repeat the installation including step 6.

7. After the installation completes, do one of the following:
 - Click **Restart**, if you were prompted to restart the system in the previous step.
During the system restart, opt for MOK (Machine Owner Key) management, choose **Enroll MOK**, and then enroll the key by using the password recommended in the previous step.
 - Otherwise, click **Exit**.

Troubleshooting information is provided in the file:

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

In macOS

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 a proxy server is enabled in your network, click **Protection Agent** in the menu bar, click **Proxy server settings**, and then specify the proxy server host name/IP address, port, and credentials.
6. If prompted, provide administrator credentials.
7. Click **Continue**.
8. Wait until the registration screen appears.
9. Do one of the following:
 - Click **Register the machine**. In the opened browser window, sign in to the service console, review the registration details, and then click **Confirm registration**.
 - Click **Show registration info**. The setup program shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. In this case, you will need to enter the registration code in the registration form. The registration code is valid for one hour.

Alternatively, you can access the registration form by clicking **All devices > Add**, scrolling down to **Registration via code**, and then clicking **Register**.

Tip Do not quit the setup program until you confirm the registration. To initiate the registration again, you will have to restart the setup program and repeat the installation procedure.

As a result, the machine will be assigned to the account that was used to log in to the service console.

- Register the machine manually by using the command line. For more information on how to do this, refer to "Registering machines manually" (p. 53).

10. If your macOS version is Mojave 10.14.x or later, grant full disk access to the protection agent to enable backup operations.

See <https://kb.acronis.com/content/62133> <https://kb.acronis.com/content/62133> for instructions.

8.4.1 Changing the logon account on Windows machines

On the **Select components** screen, define the account under which the services will run by specifying **Logon account for the agent service**. You can select one of the following:

- **Use Service User Accounts** (default for the agent service)
Service User Accounts are Windows system accounts that are used to run services. The advantage of this setting is that the domain security policies do not affect these accounts' user rights. By default, the agent runs under the **Local System** account.
- **Create a new account**
The account name will be Agent User for the agent.
- **Use the following account**
If you install the agent on a domain controller, the system prompts you to specify existing accounts (or the same account) for the agent. For security reasons, the system does not automatically create new accounts on a domain controller.

If you chose the **Create a new account** or **Use the following account** option, ensure that the domain security policies do not affect the related accounts' rights. If an account is deprived of the user rights assigned during the installation, the component may work incorrectly or not work.

Privileges required for the logon account

A protection agent is run as a Managed Machine Service (MMS) on a Windows machine. The account under which the agent will run must have specific rights for the agent to work correctly. Thus, the MMS user should be assigned the following privileges:

1. Included in the **Backup Operators** and **Administrators** groups. On a Domain Controller, the user must be included in the group **Domain Admins**.
2. Granted the **Full Control** permission on the folder **%PROGRAMDATA%\Acronis** (in Windows XP and Server 2003, **%ALLUSERSPROFILE%\Application Data\Acronis**) and on its subfolders.
3. Granted the **Full Control** permission on certain registry keys in the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\Acronis.
4. Assigned the following user rights:
 - Log on as a service
 - Adjust memory quotas for a process
 - Replace a process level token
 - Modify firmware environment values

How to assign the user rights

Follow the instructions below to assign the user rights (this example uses the **Log on as service** user right, the steps are the same for other user rights):

1. Log on to the computer by using an account with administrative privileges.
2. Open **Administrative Tools** from **Control Panel** (or click Win+R, type **control admintools**, and press Enter) and open **Local Security Policy**.
3. Expand **Local Policies** and click on **User Rights Assignment**.
4. In the right pane, right-click **Log on as a service** and select **Properties**.
5. Click on the **Add User or Group...** button to add a new user.
6. In the **Select Users, Computers, Service Accounts, or Groups** window, find the user you wish to enter and click **OK**.
7. Click **OK** in the **Log on as a service Properties** to save the changes.

Important Ensure that the user which you have added to the **Log on as service** user right is not listed in the **Deny log on as a service** policy in **Local Security Policy**.

Note that it is not recommended to change logon accounts manually after the installation is completed.

8.5 Unattended installation or uninstallation

8.5.1 Unattended installation or uninstallation in Windows

This section describes how to install or uninstall protection agents in the unattended mode on a machine running Windows, by using Windows Installer (the **msiexec** program). In an Active Directory domain, another way of performing unattended installation is through Group Policy—see "Deploying agents through Group Policy" (p. 76).

During the installation, you can use a file known as a **transform** (an .mst file). A transform is a file with installation parameters. As an alternative, you can specify installation parameters directly on the command line.

Creating the .mst transform and extracting the installation packages

1. Log on as an administrator and start the setup program.
2. Click **Create .mst and .msi files for unattended installation**.
3. In **What to install**, select the components that you want to install. The installation packages for these components will be extracted from the setup program.
4. In **Registration settings**, select **Use credentials** or **Use registration token**. For more information on how to generate a registration token, refer to "Deploying agents through Group Policy" (p. 76).
5. Review or modify other installation settings that will be added to the .mst file.
6. Click **Proceed**, and then select the folder where the .mst transform will be generated and the .msi and .cab installation packages will be extracted.
7. Click **Generate**.

Installing the product by using the .mst transform

On the command line, run the following command.

Command template:

```
msiexec /i <package name> TRANSFORMS=<transform name>
```

Here:

- <package name> is the name of the .msi file.
- <transform name> is the name of the transform.

Command example:

```
msiexec /i BackupClient64.msi TRANSFORMS=BackupClient64.msi.mst
```

Installing or uninstalling the product by specifying parameters manually

On the command line, run the following command.

Command template (installing):

```
msiexec /i <package name> <PARAMETER 1>=<value 1> ... <PARAMETER N>=<value n>
```

Here, <package name> is the name of the .msi file. All available parameters and their values are described in "Unattended installation or uninstallation parameters" (p. 43).

Command template (uninstalling):

```
msiexec /x <package name> <PARAMETER 1>=<value 1> ... <PARAMETER N>=<value n>
```

The .msi package must be of the same version as the product that you want to uninstall.

8.5.1.1 Unattended installation or uninstallation parameters

This section describes parameters that are used during unattended installation or uninstallation in Windows. In addition to these parameters, you can use other parameters of **msiexec**, as described at [https://msdn.microsoft.com/en-us/library/windows/desktop/aa367988\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/aa367988(v=vs.85).aspx).

Installation parameters

Basic parameters

ADDLOCAL=<list of components>

The components to be installed, separated by commas and without space characters. All of the specified components must be extracted from the setup program prior to installation.

The full list of the components is as follows:

Component	Must be installed together with	Bitness	Component name / description
MmsMspComponents		32-bit/64-bit	Core components for agents
BackupAndRecoveryAgent	MmsMspComponents	32-bit/64-bit	Agent for Windows
ArxAgentFeature	BackupAndRecoveryAgent	32-bit/64-bit	Agent for Exchange
ArsAgentFeature	BackupAndRecoveryAgent	32-bit/64-bit	Agent for SQL
ARADAgentFeature	BackupAndRecoveryAgent	32-bit/64-bit	Agent for Active Directory
ArxOnlineAgentFeature	MmsMspComponents	32-bit/64-bit	Agent for Office 365
OracleAgentFeature	BackupAndRecoveryAgent	32-bit/64-bit	Agent for Oracle

Component	Must be installed together with	Bitness	Component name / description
AcronisESXSupport	MmsMspComponents	64-bit	Agent for VMware ESX(i) (Windows)
HyperVAgent	MmsMspComponents	32-bit/64-bit	Agent for Hyper-V
CommandLineTool		32-bit/64-bit	Command-Line Tool
TrayMonitor	BackupAndRecoveryAgent	32-bit/64-bit	Cyber Protection Monitor

TARGETDIR=<path>

The folder where the product will be installed. By default, this folder is: C:\Program Files\BackupClient.

REBOOT=ReallySuppress

If the parameter is specified, the machine reboot is forbidden.

/1*v <log file>

If the parameter is specified, the installation log in the verbose mode will be saved to the specified file. The log file can be used for analyzing the installation issues.

CURRENT_LANGUAGE=<language ID>


The product language. Available values are as follows: **en, bg, cs, da, de, es, fr, hu, id, it, ja, ko, ms, nb, nl, pl, pt, pt_BR, ru, fi, sr, sv, tr, zh, zh_TW**. If this parameter is not specified, the product language will be defined by your system language on the condition that it is in the list above. Otherwise, the product language will set to English (**en**).

Registration parameters


REGISTRATION_ADDRESS

This is the URL for the Cyber Protection service. You can use this parameter either with the **REGISTRATION_LOGIN** and **REGISTRATION_PASSWORD** parameters, or with the **REGISTRATION_TOKEN** one.

- When you use **REGISTRATION_ADDRESS** with **REGISTRATION_LOGIN** and **REGISTRATION_PASSWORD** parameters, specify the address that you use **to log in** to the Cyber Protection service. For example, <https://cloud.company.com>:



- When you use **REGISTRATION_ADDRESS** with the **REGISTRATION_TOKEN** parameter, specify the exact datacenter address. This is the URL that you see **once you are logged in** to the Cyber Protection service. For example, <https://eu2-cloud.company.com>.



Do not use <https://cloud.company.com> here.

REGISTRATION_LOGIN and REGISTRATION_PASSWORD

Credentials for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.

REGISTRATION_PASSWORD_ENCODED

Password for the account under which the agent will be registered in the Cyber Protection service, encoded in base64. For more information on how to encode your password, refer to "Registering machines manually" (p. 53).

REGISTRATION_TOKEN

The registration token is a series of 12 characters, separated by hyphens in three segments. You can generate one in the service console, as described in "Deploying agents through Group Policy" (p. 76).

REGISTRATION_REQUIRED={0,1}

Defines how the installation will finish if the registration fails. If the value is **1**, the installation also fails. The default value is **0**, so if you don't specify this parameter, the installation completes successfully even though the agent is not registered.

Additional parameters

To define the logon account for the agent service in Windows, use one of the following parameters:

- **MMS_USE_SYSTEM_ACCOUNT={0,1}**
If the value is **1**, the agent will run under the **Local System** account.
- **MMS_CREATE_NEW_ACCOUNT={0,1}**
If the value is **1**, the agent will run under a newly created account named **Acronis Agent User**.
- **MMS_SERVICE_USERNAME=<user name>** and **MMS_SERVICE_PASSWORD=<password>**
Use these parameters to specify an existing account under which the agent will run.

For more information on logon accounts, refer to "Changing the logon account on Windows machines" (p. 41).

SET_ESX_SERVER={0,1}

If the value is **0**, Agent for VMware being installed will not be connected to a vCenter Server or an ESXi host. If the value is **1**, specify the following parameters:

- **ESX_HOST=<host name>**
The host name or IP address of the vCenter Server or the ESXi host.
- **ESX_USER=<user name>** and **ESX_PASSWORD=<password>**
Credentials to access the vCenter Server or ESXi host.

HTTP_PROXY_ADDRESS=<IP address> and **HTTP_PROXY_PORT=<port>**

The HTTP proxy server to be used by the agent. Without these parameters, no proxy server will be used.

HTTP_PROXY_LOGIN=<login> and **HTTP_PROXY_PASSWORD=<password>**

The credentials for the HTTP proxy server. Use these parameters if the server requires authentication.

HTTP_PROXY_ONLINE_BACKUP={0,1}

If the value is **0**, or the parameter is not specified, the agent will use the proxy server only for backup and recovery from the cloud. If the value is **1**, the agent also will connect to the management server through the proxy server.

Uninstallation parameters

REMOVE={<list of components>|ALL}

The components to be removed, separated by commas and without space characters. If the value is **ALL**, all of the product components will be uninstalled.

Additionally, you can specify the following parameter:

DELETE_ALL_SETTINGS={0, 1}

If the value is **1**, the product's logs, tasks, and configuration settings will be removed.

Examples

- Installing Agent for Windows, Command-Line Tool, and Cyber Protection Monitor. Registering the machine in the Cyber Protection service by using a user name and password.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn  
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,TrayMonitor  
TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress  
MMS_USE_SYSTEM_ACCOUNT=1 REGISTRATION_ADDRESS=https://cloud.company.com  
REGISTRATION_LOGIN=johndoe REGISTRATION_PASSWORD=johnspassword
```

- Installing Agent for Windows, Command-Line Tool, and Cyber Protection Monitor. Creating a new logon account for the agent service in Windows. Registering the machine in the Cyber Protection service by using a token.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn  
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,TrayMonitor  
TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress  
MMS_CREATE_NEW_ACCOUNT=1 REGISTRATION_ADDRESS=https://eu2-cloud.company.com  
REGISTRATION_TOKEN=34F6-8C39-4A5C
```

- Installing Agent for Windows, Command-Line Tool, Agent for Oracle and Cyber Protection Monitor. Registering the machine in the Cyber Protection service by using a user name and encoded in base64 password.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn  
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,OracleAgentFeature,TrayMonitor  
TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress  
CURRENT_LANGUAGE=en MMS_USE_SYSTEM_ACCOUNT=1  
REGISTRATION_ADDRESS=https://cloud.company.com REGISTRATION_LOGIN=johndoe  
REGISTRATION_PASSWORD_ENCODED=am9obnNwYXNzd29yZA==
```

- Installing Agent for Windows, Command-Line Tool, and Cyber Protection Monitor. Registering the machine in the Cyber Protection service by using a token. Setting an HTTP proxy.

```
msiexec.exe /i BackupClient64.msi /l*v my_log.txt /qn  
ADDLOCAL=MmsMspComponents,BackupAndRecoveryAgent,CommandLineTool,TrayMonitor  
TARGETDIR="C:\Program Files\BackupClient" REBOOT=ReallySuppress  
CURRENT_LANGUAGE=en MMS_USE_SYSTEM_ACCOUNT=1  
REGISTRATION_ADDRESS=https://eu2-cloud.company.com  
REGISTRATION_TOKEN=34F6-8C39-4A5C  
HTTP_PROXY_ADDRESS=https://my-proxy.company.com HTTP_PROXY_PORT=80  
HTTP_PROXY_LOGIN=tomsmith HTTP_PROXY_PASSWORD=tomspassword
```

- Uninstalling all the agents and deleting their logs, tasks, and configuration settings.

```
msiexec.exe /x BackupClient64.msi /l*v uninstall_log.txt REMOVE=ALL  
DELETE_ALL_SETTINGS=1 REBOOT=ReallySuppress
```

8.5.2 Unattended installation or uninstallation in Linux

This section describes how to install or uninstall protection agents in the unattended mode on a machine running Linux, by using the command line.

To install or uninstall a protection agent

1. Open Terminal.

2. Do one of the following:

- To start the installation by specifying the parameters on the command line, run the following command:

```
<package name> -a <parameter 1> ... <parameter N>
```

Here, <package name> is the name of the installation package (an .i686 or an .x86_64 file). All available parameters and their values are described in "Unattended installation or uninstallation parameters" (p. 47).

- To start the installation with parameters that are specified in a separate text file, run the following command:

```
<package name> -a --options-file=<path to the file>
```

This approach might be useful if you don't want to enter sensitive information on the command line. In this case, you can specify the configuration settings in a separate text file and ensure that only you can access it. Put each parameter on a new line, followed by the desired value, for example:

```
--rain=https://cloud.company.com  
--login=johndoe  
--password=johnspassword  
--auto
```

or

```
-C  
https://cloud.company.com  
-g  
johndoe  
-w  
johnspassword  
-a  
--language  
en
```

If the same parameter is specified both on the command line and in the text file, the command line value precedes.

3. If UEFI Secure Boot is enabled on the machine, you are informed that you need to restart the system after the installation. Be sure to remember what password (that of the root user or "acronis") should be used. During the system restart, opt for MOK (Machine Owner Key) management, choose **Enroll MOK**, and then enroll the key by using the recommended password.

If you enable UEFI Secure Boot after the agent installation, repeat the installation, including step 3. Otherwise, backups will fail.

8.5.2.1 Unattended installation or uninstallation parameters

This section describes parameters that are used during unattended installation or uninstallation in Linux.

The minimal configuration for unattended installation includes **-a** and registration parameters (for example, **--login** and **--password** parameters; **--rain** and **--token** parameters). You can use more parameters to customize you installation.

Installation parameters

Basic parameters

{-i|--id=}<list of components>

The components to be installed, separated by commas and without space characters. The following components are available in the .x86_64 installation package:

Component	Component description
BackupAndRecoveryAgent	Agent for Linux
AgentForPCS	Agent for Virtuozzo
OracleAgentFeature	Agent for Oracle

Without this parameter, all of the above components will be installed.

Both Agent for Virtuozzo and Agent for Oracle require that Agent for Linux is also installed.

The .i686 installation package contains only BackupAndRecoveryAgent.

{-a|--auto}

The installation and registration process will complete without any further user interaction.

When using this parameter, you must specify the account under which the agent will be registered in the Cyber Protection service, either by using the **--token** parameter, or by using the **--login** and **--password** parameters.

{-t|--strict}

If the parameter is specified, any warning that occurs during the installation results in installation failure. Without this parameter, the installation completes successfully even in the case of warnings.

{-n|--nodeps}

The absence of required Linux packages will be ignored during the installation.

{-d|--debug}

Writes the installation log in the verbose mode.

--options-file=<location>

The installation parameters will be read from a text file instead of the command line.

--language=<language ID>

The product language. Available values are as follows: **en, bg, cs, da, de, es, fr, hu, id, it, ja, ko, ms, nb, nl, pl, pt, pt_BR, ru, fi, sr, sv, tr, zh, zh_TW**.

If this parameter is not specified, the product language will be defined by your system language on the condition that it is in the list above. Otherwise, the product language will set to English (**en**).

Registration parameters

Specify one of the following parameters:

- **{-g|--login=}<user name>** and **{-w|--password=}<password>**

Credentials for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.

- **--token=<token>**


The registration token is a series of 12 characters, separated by hyphens in three segments. You can generate one in the service console, as described in "Deploying agents through Group Policy" (p. 76).

You cannot use the **--token** parameter along with **--login**, **--password**, and **--register-with-credentials** parameters.


- **{-C|--rain=}<service address>**

The URL of the Cyber Protection service.

You don't need to include this parameter explicitly when you use **--login** and **--password** parameters for registration, because the installer uses the correct address by default – this would be the address that you use **to log in** to the Cyber Protection service. For example:



However, when you use **{-C|--rain=}** with the **--token** parameter, you must specify the exact datacenter address. This is the URL that you see **once you are logged in** to the Cyber Protection service. For example:



- **--register-with-credentials**

If this parameter is specified, the installer's graphical interface will start. To finish the registration, enter the user name and password for the account under which the agent will be registered in the Cyber Protection service. This cannot be a partner administrator account.

- **--skip-registration**

Use this parameter if you need to install the agent but you plan to register it in the Cyber Protection service later. For more information on how to do this, refer to "Registering machines manually" (p. 53).

Additional parameters

- **--http-proxy-host=<IP address> and --http-proxy-port=<port>**

The HTTP proxy server that the agent will use for backup and recovery from the cloud, and for connection to the management server. Without these parameters, no proxy server will be used.

- **--http-proxy-login=<login> and --http-proxy-password=<password>**

The credentials for the HTTP proxy server. Use these parameters if the server requires authentication.

- **--tmp-dir=<location>**

Specifies the folder where the temporary files are stored during the installation. The default folder is **/var/tmp**.

- **{-s|--disable-native-shared}**

Redistributable libraries will be used during the installation, even though they might have already been present on your system.

--skip-prereq-check

There will be no check of whether the packages required for compiling the snapapi module are already installed.

--force-weak-snapapi

The installer will not compile a snapapi module. Instead, it will use a ready-made module that might not match the Linux kernel exactly. Using this option is not recommended.

--skip-svc-start

The services will not start automatically after the installation. Most often, this parameter is used with the **--skip-registration** one.

Information parameters

{-?|--help}

Shows the description of parameters.

--usage

Shows a brief description of the command usage.

{-v|--version}

Shows the installation package version.

--product-info

Shows the product name and the installation package version.

--snapapi-list

Shows the available ready-made snapapi modules.

--components-list

Shows the installer components.

Parameters for legacy features

These parameters relate to a legacy component, agent.exe.

{-e|--ssl=}<path>

Specifies the path to a custom certificate file for SSL communication.

{-p|--port=}<port>

Specifies the port on which agent.exe listens for connections. The default port is 9876.

Uninstallation parameters

{-u|--uninstall}

Uninstalls the product.

--purge

Uninstalls the product and removes its logs, tasks, and configuration settings. You don't need to specify the **--uninstall** parameter explicitly when you use the **--purge** one.

Examples

- Installing Agent for Linux without registering it.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -i BackupAndRecoveryAgent -a  
--skip-registration
```

- Installing Agent for Linux, Agent for Virtuozzo, and Agent for Oracle, and registering them by using credentials.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -a --login=johndoe
--password=johnspassword
```

- Installing Agent for Oracle and Agent for Linux, and registering them by using a registration token.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -i
BackupAndRecoveryAgent,OracleAgentFeature -a
--rain=https://eu2-cloud.company.com --token=34F6-8C39-4A5C
```

- Installing Agent for Linux, Agent for Virtuozzo, and Agent for Oracle with configuration settings in a separate text file.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -a
--options-file=/home/mydirectory/configuration_file
```

- Uninstalling Agent for Linux, Agent for Virtuozzo, and Agent for Oracle, and removing all its logs, tasks, and configuration settings.

```
./Cyber_Protection_Agent_for_Linux_x86_64.bin -a --purge
```

8.5.3 Unattended installation and uninstallation in macOS

This section describes how to install, register, and uninstall the Acronis Cyber Protection agent in the unattended mode on a machine running macOS, by using the command line.

To download the installation file (.dmg)

1. In the service console, go to **Devices > All devices**.
2. Click **Add**, and then click **Mac**.

To install Agent for Mac

1. Create a temporary directory where you will mount the installation file (.dmg).

```
mkdir <dmg_root>
```

Here, <dmg_root> is a name of your choice.

2. Mount the .dmg file.

```
hdiutil attach <dmg_file> -mountpoint <dmg_root>
```

Here, <dmg_file> is the name of the installation file. For example, **Acronis_Cyber_Protection_Agent_for_MAC_x64.dmg**.

3. Run the installer.

```
sudo installer -pkg <dmg_root>/Install.pkg -target LocalSystem
```

4. Detach the installation file (.dmg).

```
hdiutil detach <dmg_root>
```

Examples

```
mkdir mydirectory
hdiutil attach /Users/JohnDoe/Acronis_Cyber_Protection_Agent_for_MAC_x64.dmg
-mountpoint mydirectory
sudo installer -pkg mydirectory/Install.pkg -target LocalSystem
hdiutil detach mydirectory
```

To register Agent for Mac


Do one of the following:

- Register the agent under a specific account, by using a user name and password.

```
sudo
/Library/Application\ Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent -a <Cyber Protection service address> -t cloud -u <user name> -p <password> -o register
```

Here:

<Cyber Protection service address> is the address that you use **to log in** to the Cyber Protection service. For example:




<user name> and <password> are the credentials for the account under which the agent will be registered. This cannot be a partner administrator account.

- Register the agent by using a registration token.

```
sudo
/Library/Application\ Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent -a <Cyber Protection service address> -t cloud -o register --token <token>
```

The registration token is a series of 12 characters, separated by hyphens in three segments. You can generate one in the service console, as described in "Deploying agents through Group Policy" (p. 76).

When you use a registration token, you must specify the exact datacenter address. This is the URL that you see **once you are logged in** to the Cyber Protection service. For example:



Examples

Registration with a user name and password.

```
sudo
/Library/Application\ Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent -a https://cloud.company.com -t cloud -u johndoe -p johnspassword -o register
```

Registration with a token.

```
sudo
/Library/Application\ Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent -a https://eu2-cloud company.com -t cloud o -register --token D91D-DC46-4F0B
```

Important If you use macOS 10.14 or later, grant the protection agent full disk access. To do so, go to **Applications > Utilities**, and then run **Cyber Protect Agent Assistant**. Then, follow the instructions in the application window.

To uninstall Agent for Mac

Run the following command:

```
sudo
/Library/Application\ Support/BackupClient/Acronis/Cyber\ Protect\ Agent\ Uninstall.app/Contents/MacOS/AgentUninstall /confirm
```

To remove all logs, tasks and configuration settings during the uninstallation, run the following command:


```
sudo  
/Library/Application\ Support/BackupClient/Acronis/Cyber\ Protect\ Agent\ Unins  
tall.app/Contents/MacOS/AgentUninstall /confirm /purge
```

8.6 Registering machines manually

In addition to registering a machine in the Cyber Protection service during the agent installation, you can also register it by using the command line interface. You might need to do so if you have installed the agent but the automatic registration failed, for example, or if you want to register an existing machine under a new account.

To register a machine

To register a machine by using a user name and password, run the following command.

In Windows

Command for registering a machine under the current account:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -s mms  
-t cloud --update
```

Command template for registering a machine under another account:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t  
cloud -a <service address> -u <user name> -p <password>
```

Command example:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t  
cloud -a https://cloud.company.com -u johndoe -p johnspassword
```

In Linux

Command for registering a machine under the current account:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -s mms -t cloud  
--update
```

Command template for registering a machine under another account:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a  
<service address> -u <user name> -p <password>
```

Command example:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a  
https://cloud.company.com -u johndoe -p johnspassword
```

In macOS

Command for registering a machine under the current account:

```
sudo "/Library/Application  
Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent" -o register -s mms -t  
cloud --update
```

Command template for registering a machine under another account:


```
sudo "/Library/Application
Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a
<service address> -u <user name> -p <password>
```

Command example:

```
sudo "/Library/Application
Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a
https://cloud.company.com -u johndoe -p johnspassword
```

Note Use the user name and password for the specific account under which the agent will be registered. This cannot be a partner administrator account.

The service address is the URL that you use **to log in** to the Cyber Protection service. For example, <https://cloud.company.com>:



Alternatively, you can register a machine by using a registration token. To do so, run the following command.

In Windows

Command template:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t
cloud -a <service address> --token <token>
```

Command example:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t
cloud -a https://au1-cloud.company.com --token 3B4C-E967-4FBD
```

In Linux

Command template:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a
<service address> --token <token>
```

Command example:

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a
https://eu2-cloud.company.com --token 34F6-8C39-4A5C
```

In macOS


Command template:

```
sudo "/Library/Application
Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a
<service address> --token <token>
```

Command example:

```
sudo "/Library/Application
Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent" -o register -t cloud -a
https://us5-cloud.company.com --token 9DBF-3DA9-4DAB
```

Note When you use a registration token, you must specify the exact datacenter address. This is the URL that you see **once you are logged in** to the Cyber Protection service. For example, ***https://eu2-cloud.company.com***.



Do not use ***https://cloud.company.com*** here.

The registration token is a series of 12 characters, separated by hyphens in three segments. For more information on how to generate one, refer to "Deploying agents through Group Policy" (p. 76).

To unregister a machine

Run the following command:

In Windows

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o unregister
```

In Linux

```
sudo "/usr/lib/Acronis/RegisterAgentTool/RegisterAgent" -o unregister
```

In macOS

```
sudo "/Library/Application  
Support/BackupClient/Acronis/RegisterAgentTool/RegisterAgent" -o unregister
```

Passwords with special characters or blank spaces

If your password contains special characters or blank spaces, enclose it in quotation marks when you type it on the command line.

For example, in Windows, run this command.

Command template:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t  
cloud -a <service address> -u <user name> -p <"password">
```

Command example:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t  
cloud -a https://cloud.company.com -u johndoe -p "johns password"
```

If you still receive an error:

- Encode your password into base64 format at <https://www.base64encode.org/> (https://www.base64encode.org - https://www.base64encode.org).
- On the command line, specify the encoded password by using the -b or --base64 parameter.

For example, in Windows, run this command.

Command template:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t  
cloud -a <service address> -u <user name> -b -p <encoded password>
```

Command example:

```
"%ProgramFiles%\BackupClient\RegisterAgentTool\register_agent.exe" -o register -t cloud -a https://cloud.company.com -u johndoe -b -p am9obnNwYXNzd29yZA==
```

8.7 Autodiscovery of machines

The discovery of machines functionality allows you to do the following:

- Automate the process of protection agent installation and machine registration, by automatically detecting machines in your Active Directory (AD) domain or local network.
- Install and update the protection agent on a batch of machines.
- Use synchronization with Active Directory, to lower the efforts and overhead for resource provisioning and machine management in a large AD environment.

Important Machine discovery can be performed only by the agents installed on Windows machines. Currently, not only Windows machines can be detected by the discovery agent but remote software installation is possible only on Windows machines.

If there is no machine with the installed agent, then the autodiscovery functionality will be hidden - the **Multiple devices** section will be hidden in the Add new device wizard.

After adding machines to the service console, they are categorized as follows:

- **Discovered** – machines that were discovered, but the protection agent is not installed on them.
- **Managed** – machines on which the protection agent is installed.
- **Unprotected** – machines to which the protection plan is not applied. Unprotected machines include both discovered and managed machines with no protection plan applied.
- **Protected** – machines to which the protection plan is applied.

How it works

During the local network scanning, the discovery agent uses the following technologies: NetBIOS discovery, Web Service Discovery (WSD), and the Address Resolution Protocol (ARP) table. The agent tries to get the following parameters of each machine:

- Name (short/NetBIOS hostname)
- FQDN
- Domain/workgroup
- IPv4/IPv6 addresses
- MAC addresses
- Operating system (name/version/family)
- Machine category (workstation/server/domain controller)

When AD scanning is performed, the agent tries to get almost the same parameters of each machines as listed above. The difference is that it will additionally get the Organizational Unit (OU) parameter, more full information about the name and operating system, and it won't get IP address and MAC address information.

Prerequisites


Before discovering machines, you must install the protection agent on at least one machine in your local network to use it as a discovery agent.

If you are planning to discover machines in the Active Directory domain, you must install the agent on at least one machine in the AD domain. This agent will be used as a discovery agent during scanning of AD.

Note Agent for Windows cannot be installed on a remote machine running Windows XP.
To install Agent for Windows on a machine running Windows Server 2012 R2, you must have Windows update KB2999226 installed on this machine.

Machine discovery process

In the following scheme, you can see the main steps of the machine discovery process:



Generally, the whole autodiscovery process consists of the following steps:

1. Select the method of machine discovery:
 - By scanning Active Directory
 - By scanning the local network
 - Manual – adding a machine by IP address or hostname, or importing a list of machines from a file
2. Select machines to be added from the list received as a result of the previous step.
3. Select how the machines will be added:
 - The protection agent and additional components will be installed on the machines, and they will also be registered in the service console.
 - The machines will be registered in the service console (if they already have the installed agent).

- The machines will be added as **Unmanaged machines** to the service console, without any agent or component installation.

If you selected one of the first two methods to add a machine, you can also select the protection plan from the existing ones and apply to machines.

4. Provide the credentials of the user who has the administrator rights for managing the machines.
5. Verify connectivity to machines by using the provided credentials.

In the next topics, you will get more detailed information about the discovery procedure.

8.7.1 Autodiscovery and manual discovery

Before starting the discovery, ensure that the prerequisites (p. 56) are met.

To discover machines

1. In the service console, go to **Devices > All devices**.
2. Click **Add**.
3. In **Multiple devices**, click **Windows-only**. The discovery wizard opens.
4. [If there are units in your organization] Select a unit. Then, in **Discovery agent** you will be able to select the agents associated with the selected unit and its child units.
5. Select the discovery agent that will perform the scan to detect machines.
6. Select the discovery method:
 - **Search Active Directory**. Ensure that the machine with the discovery agent is the Active Directory domain member.
 - **Scan local network**. If the selected discovery agent could not find any machines, select another discovery agent.
 - **Specify manually or import from file**. Manually define the machines to be added or import them from a text file.
7. [If the Active Directory discovery method is selected] Select how to search for machines:
 - **In organizational unit list**. Select the group of machines to be added.
 - **By LDAP dialect query**. Use the LDAP dialect query to select the machines. **Search base** defines where to search, while **Filter** allows you to specify the criteria for machine selection.
8. [If the Active Directory or local network discovery method is selected] Use a list to select the machines that you want to add.

[If the Manual discovery method is selected] Specify the machine IP addresses or hostnames, or import the machine list from a text file. The file must contain IP addresses/hostnames, one per line. Here is an example of a file:

```
156.85.34.10
156.85.53.32
156.85.53.12
EN-L00000100
EN-L00000101
```

After adding machine addresses manually or importing from a file, the agent tries to ping the added machines and define their availability.

9. Select what actions must be performed after the discovery:
 - **Install agents and register machines**. You can select which components to install on the machines by clicking **Select components**. For more details, refer to "Selecting components for installation".

On the **Select components** screen, define the account under which the services will run by specifying **Logon account for the agent service**. You can select one of the following:

- **Use Service User Accounts** (default for the agent service)
Service User Accounts are Windows system accounts that are used to run services. The advantage of this setting is that the domain security policies do not affect these accounts' user rights. By default, the agent runs under the **Local System** account.
- **Create a new account**
The account name will be Agent User for the agent.
- **Use the following account**
If you install the agent on a domain controller, the system prompts you to specify existing accounts (or the same account) for the agent. For security reasons, the system does not automatically create new accounts on a domain controller.

If you chose the **Create a new account** or **Use the following account** option, ensure that the domain security policies do not affect the related accounts' rights. If an account is deprived of the user rights assigned during the installation, the component may work incorrectly or not work.

- **Register machines with installed agents.** This option is used if the agent is already installed on machines and you need only to register them in Cyber Protection. If no agent is found inside the machines, then they will be added as **Unmanaged** machines.
- **Add as unmanaged machines.** The agent will not be installed on the machines. You will be able to view them in the console and install or register the agent later.

[If the **Install agents and register machines** post-discovery action is selected] **Restart the machine if required** – if the option is enabled, the machine will be restarted as many times as required to complete the installation.

Restart of the machine may be required in one of the following cases:

- Installation of prerequisites is completed and restart is required to continue the installation
- Installation is completed but restart is required as some files are locked during installation
- Installation is completed but restart is required for other previously installed software

[If **Restart the machine if required** is selected] **Do not restart if the user logged in** – if the option is enabled, the machine will not be automatically restarted if the user is logged in to the system. For example, if a user is working while installation requires restart, the system will not be restarted.

If the prerequisites were installed and then the reboot was not done because a user was logged in, then to complete the agent installation you need to reboot the machine and start the installation again.

If the agent was installed but then the reboot was not done, then you need to reboot the machine.

[If there are units in your organization] **User for whom to register the machines** – select the user of your unit or subordinate units for whom the machines will be registered.

If you have selected one of the first two post-discovery actions, then there is also an option to apply the protection plan to the machines. If you have several protection plans, you can select which one to use.

10. Specify the credentials of the user with administrator rights for all of the machines.

Important *Note that remote installation of agent works without any preparations only if you specify the credentials of the built-in administrator account (the first account created when the operating system is installed). If you want to define some custom administrator credentials, then you should do additional manual preparations as described in "Enabling remote installation of an agent for a custom administrator" below.*

11. The system checks connectivity to all of the machines. If the connection to some of the machines fails, you can change the credentials for these machines.

When the discovery of machines is initiated, you will find the corresponding task in **Dashboard > Activities > Discovering machines** activity.

Preparing a machine for remote installation

1. 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* on that machine. For more information on how to disable it, refer to "Requirements on User Account Control (UAC) (p. 60)" > To disable UAC.
3. By default, the credentials of the built-in administrator account are required for remote installation on any Windows machine. To perform remote installation by using the credentials of another administrator account, User Account Control (UAC) remote restrictions must be *disabled*. For more information on how to disable them, refer to "Requirements on User Account Control (UAC) (p. 60)" > To disable UAC remote restrictions.
4. File and Printer Sharing must be *enabled* on the remote machine. To access this option:
 - On a machine running 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**.
5. Cyber Protection uses TCP ports 445, 25001, and 43234 for remote installation.

Port 445 is automatically opened when you enable File and Printer Sharing. Ports 43234 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, port 25001 is automatically closed through Windows Firewall. Ports 445 and 43234 need to remain open if you want to update the agent remotely in the future. Port 25001 is automatically opened and closed through Windows Firewall during each update. If you use a different firewall, keep all the three ports open.

8.7.1.1 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 and UAC remote restrictions 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.

To disable UAC remote restrictions

1. Open Registry Editor.
2. Locate the following registry key:
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System
3. For **LocalAccountTokenFilterPolicy** value, change the setting to **1**.

If the **LocalAccountTokenFilterPolicy** value does not exist, create it as DWORD (32-bit). For more information about this value, refer to the Microsoft documentation:
<https://support.microsoft.com/en-us/help/951016/description-of-user-account-control-and-remote-restrictions-in-windows>.

Note For security reasons, it is recommended that after finishing the management operation – for example, remote installation, both of the settings be reverted to their original state: **EnableLUA=1** and **LocalAccountTokenFilterPolicy = 0**

8.7.1.2 Selecting components for installation

You can find the description of mandatory and additional components in the following table:

Component	Description
Mandatory component	
Agent for Windows	This agent backs up disks, volumes, files and will be installed on Windows machines. It will be always installed, not selectable.
Additional components	
Agent for Hyper-V	This agent backs up Hyper-V virtual machines and will be installed on Hyper-V hosts. It will be installed if selected and detected Hyper-V role on a machine.
Agent for SQL	This agent backs up SQL Server databases and will be installed on machines running Microsoft SQL Server. It will be installed if selected and application detected on a machine.
Agent for Exchange	This agent backs up Exchange databases and mailboxes and will be installed on machines running the Mailbox role of Microsoft Exchange Server. It will be installed if selected and application detected on a machine.
Agent for Active Directory	This agent backs up the data of Active Directory Domain Services and will be installed on domain controllers. It will be installed if selected and application detected on a machine.
Agent for VMware (Windows)	This agent backs up VMware virtual machines and will be installed on Windows machines that have network access to vCenter Server. It will be installed if selected.

Agent for Office 365	This agent backs up Microsoft Office 365 mailboxes to a local destination and will be installed on Windows machines. It will be installed if selected.
Agent for Oracle	This agent backs up Oracle databases and will be installed on machines running Oracle Database. It will be installed if selected.
Cyber Protection Monitor	This component enables a user to monitor execution of running tasks in the notification area and will be installed on Windows machines. It will be installed if selected.
Command-line tool	Cyber Protection supports the command-line interface with the acrocmd utility. acrocmd does not contain any tools that physically execute the commands. It only provides the command-line interface to Cyber Protection components - agents and the management server. It will be installed if selected.

8.7.2 Managing discovered machines

After the discovery process is performed, you can find all of the discovered machines in **Devices > Unmanaged machines**.

This section is divided into subsections by the discovery method used. The full list of machine parameters is shown below (it may vary depending on the discovery method):

Name	Description
Name	The name of the machine. The IP address will be shown if the name of the machine could not be discovered.
IP address	The IP address of the machine.
Discovery type	The discovery method that was used to detect the machine.
Organizational unit	The organizational unit in Active Directory that the machine belongs to. This column is shown if you view the list of machines in Unmanaged machines > Active Directory .
Operating system	The operating system installed in the machine.

There is an **Exceptions** section, where you can add the machines that must be skipped during the discovery process. For example, if you do not need the exact machines to be discovered, you can add them to this list.

To add a machine to **Exceptions**, select it in the list and click **Add to exceptions**. To remove a machine from **Exceptions**, go to **Unmanaged machines > Exceptions**, select the machine, and click **Remove from exceptions**.

You can install the protection agent and register a batch of discovered machines in Cyber Protection by selecting them in the list and clicking **Install and register**. The opened wizard also allows you to assign the protection plan to a batch of machines.


After the protection agent is installed on machines, those machines will be shown in the **Devices > Machines with agents** section.

To check your protection status, go to **Dashboard > Overview** and add the **Protection status** (p. 380) widget or the **Discovered machine** (p. 380) widget.


8.7.3 Troubleshooting

If you have any issues with the autodiscovery functionality, try to check the following:

- Check that NetBIOS over TCP/IP is enabled or set to default.



- In the “Control Panel\Network and Sharing Center\Advanced sharing settings” turn on network discovery.



- Check that the Function Discovery Provider Host service is running on the machine that does discovery and on the machines to be discovered.
- Check that the Function Discovery Resource Publication service is running on the machines to be discovered.

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

8.8.1 Before you start

System requirements for the agent

By default, the virtual appliance is assigned 4 GB of RAM and 2 vCPUs, which is optimal and sufficient for most operations. We recommend increasing these resources to 8 GB of RAM and 4 vCPUs if the backup traffic bandwidth is expected to exceed 100 MB per second (for example, in 10-Gbit networks), in order to improve backup performance.

The appliance's own virtual disks occupy no more than 6 GB. Thick or thin disk format does not matter, it does not affect the appliance performance.

How many agents do I need?

Even though one virtual appliance is able to protect an entire vSphere environment, the best practice is deploying one virtual appliance per vSphere cluster (or per host, if there are no clusters). This makes for faster backups because the appliance can attach the backed-up disks by using the HotAdd transport, and therefore the backup traffic is directed from one local disk to another.

It is normal to use both the virtual appliance and Agent for VMware (Windows) at the same time, as long as they are connected to the same vCenter Server or they are connected to different ESXi hosts. Avoid cases when one agent is connected to an ESXi directly and another agent is connected to the vCenter Server which manages this ESXi.

We do not recommend using locally attached storage (i.e. storing backups on virtual disks added to the virtual appliance) if you have more than one agent. For more considerations, see "Using a locally attached storage".

Disable automatic DRS for the agent

If the virtual appliance is deployed to a vSphere cluster, be sure to disable automatic vMotion for it. In the cluster DRS settings, enable individual virtual machine automation levels, and then set **Automation level** for the virtual appliance to **Disabled**.

8.8.2 Deploying the OVF template

1. Click **All devices > Add > VMware ESXi > Virtual Appliance (OVF)**.

The .zip archive is downloaded to your machine.

2. Unpack the .zip archive. The folder contains one .ovf file and two .vmdk files.
3. Ensure that these files can be accessed from the machine running the vSphere Client.
4. Start the vSphere Client and log on to the vCenter Server.
5. Deploy the OVF template.
 - When configuring storage, select the shared datastore, if it exists. Thick or thin disk format does not matter, as it does not affect the appliance performance.
 - When configuring network connections, be sure to select a network that allows an Internet connection, so that the agent can properly register itself in the cloud.

8.8.3 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.

2. Proxy server

If a proxy server is enabled in your network:

- a. To start the command shell, press CTRL+SHIFT+F2 while in the virtual appliance UI.
- b. Open the file **/etc/Acronis/Global.config** in a text editor.
- c. Find the following section:

```
<key name="HttpProxy">
  <value name="Enabled" type="Tdword">"0"</value>
  <value name="Host" type="TString">"ADDRESS"</value>
  <value name="Port" type="Tdword">"PORT"</value>
  <value name="Login" type="TString">"LOGIN"</value>
  <value name="Password" type="TString">"PASSWORD"</value>
</key>
```

- d. Replace **0** with **1**.
- e. Replace **ADDRESS** with the new proxy server host name/IP address, and **PORT** with the decimal value of the port number.
- f. If your proxy server requires authentication, replace **LOGIN** and **PASSWORD** with the proxy server credentials. Otherwise, delete these lines from the file.
- g. Save the file.
- h. Execute the **reboot** command.

Otherwise, skip this step.

3. 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.

4. 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. 282) on the vCenter Server or ESXi.

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

5. Management server

- a. Under **Agent options**, in **Management Server**, click **Change**.
- b. In **Server name/IP**, select **Cloud**. The software displays the Cyber Protection service address. Do not change this address unless instructed otherwise.
- c. In **User name** and **Password**, specify the user name and password for the Cyber Protection service. The agent and the virtual machines managed by the agent will be registered under this account.

6. 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.

7. [Optional] 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.

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.

8.9 Deploying Agent for Virtuozzo Infrastructure Platform (Virtual Appliance) from a QCOW2 template

8.9.1 Before you start

This appliance is a pre-configured virtual machine that you deploy in Virtuozzo Hybrid Infrastructure. It contains a protection agent that enables you to administer cyber protection for all virtual machines in a Virtuozzo Hybrid Infrastructure cluster.

System requirements for the agent

When deploying the virtual appliance, you can choose between different predefined combinations of vCPUs and RAM (flavors). You can also create your own flavors.

2 vCPUs and 4 GB of RAM (medium flavor) are optimal and sufficient for most operations. We recommend increasing these resources to 4 vCPUs and 8 GB of RAM if the backup traffic bandwidth is expected to exceed 100 MB per second (for example, in 10-Gbit networks), in order to improve backup performance.

How many agents do I need?

One agent can protect the entire cluster. However, you can have more than one agent in the cluster if you need to distribute the backup traffic bandwidth load.

If you have more than one agent in a cluster, the virtual machines are automatically evenly distributed between the agents, so that each agent manages an equal number of machines.

Automatic redistribution takes place when a load imbalance among the agents reaches 20 percent. This may happen, for example, when a machine or an agent is added or removed. For example, you realize that you need more agents to help with throughput and you deploy an additional virtual appliance to the cluster. The management server will assign the most appropriate machines to the new agent. The old agents' load will reduce. When you remove an agent from the management server, the machines assigned to the agent are distributed among the remaining agents. However, this will not happen if an agent gets corrupted or is deleted manually from the Virtuozzo Hybrid Infrastructure node. Redistribution will start only after you remove such an agent from the Cyber Protection web interface.

You can view the result of the automatic distribution:

- In the **Agent** column for each virtual machine in the **All devices** section
- In the **Assigned virtual machines** section of the **Details** panel when an agent is selected in **Settings > Agents**

Limitations

- Virtuozzo Hybrid Infrastructure appliance cannot be deployed remotely.
- Application-aware backup of virtual machines is not supported.

8.9.2 Configuring networks in Virtuozzo Hybrid Infrastructure

Before deploying and configuring the virtual appliance, you need to have your networks in Virtuozzo Hybrid Infrastructure configured. There are two types of networks – the physical ones, in **Infrastructure > Networks**; and the virtual ones in **Compute > Networks**.

The basic configuration consists of only one physical network. Recommended configurations consist of two or more physical networks, and consequently – two or more virtual networks.


Basic configuration

1. Physical network

By default, in **Infrastructure > Networks**, you have two automatically created network templates – **Private** and **Public**. The **Private** network is meant to be a secure or isolated network that is located behind a firewall and without direct access from the Internet.

For this basic configuration, ensure that **VM Private**, **Compute API**, **VM Backups**, **ABGW Public**, and **VM Public** traffic types are assigned to this network.

For more information about the networks, refer to Managing Networks and Traffic Types in the Virtuozzo Hybrid Infrastructure documentation.




	Private 192.168.1.0/24	Public
Exclusive traffic types		
Storage	*	—
Internal management	*	—
OSTOR private	*	—
ABGW private	*	—
VM private	*	—
Compute API	*	—
VM backups	*	—
Regular traffic types		
S3 public	—	*
iSCSI	—	*
NFS	—	*
ABGW public	*	*
Admin panel	*	*
SSH	*	*
VM public	*	—
SNMP	—	—
Self-service panel	—	—

2. Network interfaces of the nodes

Ensure that an IP address is assigned to the network interfaces of all the nodes in the cluster. These addresses depend on your actual network infrastructure.

Then, ensure that the physical network that you configured in step 1 is also assigned to the network interfaces.

For more information on the network interfaces, refer to [Configuring Node Network Interfaces in the Virtuozzo Hybrid Infrastructure documentation](#).




3. Virtual network

By default, in **Compute > Networks**, you have two automatically created networks named **private** and **public**. Their types are *private* and *public*, respectively. You can use them or create your own virtual networks.

For this setup, you need one virtual network of type *public*. It must be created on top of the physical network that you configured in step 1.

Note Only virtual networks of type *public* can be created on top of physical networks. Physical networks must have the VM Public traffic type is assigned.

For more information on how to create and configure a virtual network, refer to [Managing Compute Network in the Virtuozzo Hybrid Infrastructure documentation](#).



When deploying the appliance's virtual machine, you need to select this virtual network.

The screenshot shows the Virtuozzo management interface. On the left is a sidebar with navigation options: TWP, MONITORING, COMPUTE, STORAGE SERVICES, INFRASTRUCTURE, and SETTINGS. The main area is divided into two panels. The left panel, titled 'Virtual machines', contains a table of virtual machines. The right panel, titled 'Agent', shows details for a specific agent, including its status, node, creation time, and various properties. The 'Networks' property is highlighted with a blue box, showing it is set to 'public (unmanaged)'.

Name	Status	IP address	vCPUs	RAM	Storage
Agent	Active		2	4 GB	8 GB
CentOS 8	Active		1	2 GB	64 GB
Cirros	Active		1	512 MB	1 GB
Windows Server 2016	Active		2	4 GB	64 GB

Property	Value
Status	Active
Node	t-vip2
Creation time	June 15, 2020 3:50 PM
SSH key	Not specified
VM ID	3f12a665-c14f-4d9c-9a1b-b96ce849225e
Project ID	dfc42a0509eaf465e9820a54672ca1f74

Property	Value
Name	Agent
Image	ACIAppliance.qcow2
Volumes	disk.1 (8 GB, default) Bootable
Flavor	medium (2 vCPU, 4 GB RAM)
Networks	public (unmanaged) → fa:16:3e:2e:78:76
High availability	Enabled

Recommended configuration with two or more physical networks

1. Physical networks

With this configuration, you need two or more networks. By default, in **Infrastructure > Networks**, you have two automatically created networks – **Private** and **Public**. You can use them or create your own networks. For example, you can have a third network only for the storage traffic. The **Private** network is meant to be a secure or isolated network that is located behind a firewall and without direct access from the Internet.

Ensure that:

- **VM Private**, **VM Backups** and **VM public** traffic types are assigned to the **Private** network.
- **Compute API**, **ABGW Public**, and **VM Public** traffic types are assigned to the **Public** network.
- The other traffic types are distributed according to your needs.

For more information about the networks, refer to Managing Networks and Traffic Types in the Virtuozzo Hybrid Infrastructure documentation.

The screenshot shows the 'Networks' configuration page in the Virtuozzo interface. It displays a table with traffic types and their assignments to Private and Public networks. The 'VM private' traffic type is assigned to the Private network, while 'Compute API', 'VM backups', 'ABGW public', and 'VM public' are assigned to the Public network. The 'Private' network is identified by IP address 192.168.1.0/24, and the 'Public' network is identified by 10.250.40.0/21.


	Private 192.168.1.0/24	Public 10.250.40.0/21
Exclusive traffic types		
Storage	•	—
Internal management	•	—
OSTOR private	•	—
ABGW private	•	—
VM private	•	—
Compute API	—	•
VM backups	•	—
Regular traffic types		
S3 public	—	•
iSCSI	—	•
NFS	—	•
ABGW public	—	•
Admin panel	—	•
SSH	—	•
VM public	•	•
SNMP	—	—
Self-service panel	—	•

2. Network interfaces of the nodes

Ensure that IP addresses are assigned to the network interfaces of all the nodes in the cluster. These addresses depend on your actual network infrastructure. You can configure network interfaces in **Infrastructure > Nodes > Network**.

Then, ensure that the physical networks that you configured in step 1 are also assigned to the network interfaces. You can assign only one network to an interface, so the number of necessary interfaces for each node depends on the number of networks you use.

For more information on the network interfaces, refer to [Configuring Node Network Interfaces](#) in the [Virtuozzo Hybrid Infrastructure](#) documentation.



The screenshot shows the 'Nodes > Network' page in the Virtuozzo interface. A table lists network interfaces for node 't-vip1'. Two interfaces are highlighted with a red box: 'br-ens224' (Private) and 'br-ens192' (Public). Both are status 'OK' and have a speed of '1 Gb / 1 Gb'.

Name	Status	IP addresses	Speed	Network
br-ens224	OK	192.168.1.100/24	1 Gb / 1 Gb	Private
br-ens192	OK	10.250.43.53/21	1 Gb / 1 Gb	Public

3. Virtual networks


By default, in **Compute > Networks**, you have two automatically created networks named **private** and **public**. Their types are *private* and *public*, respectively. You can use them or create your own virtual networks.

For this setup, you need two virtual networks of type *public*.

Ensure that the two virtual networks of type *public* are created on top of the physical networks that you created in step 1. If you have more physical networks in your setup, create more virtual networks of type *public* accordingly.

Note Only virtual networks of type *public* can be created on top of physical networks. Physical networks must have the VM Public traffic type is assigned.


For more information on how to create and configure a virtual network, refer to [Managing Compute Network](#) in the [Virtuozzo Hybrid Infrastructure](#) documentation.



The screenshot shows the 'Networks' page in the Virtuozzo interface. A table lists virtual networks. Three networks are highlighted with red boxes: 'ForAgent' (Public), 'private' (Private), and 'public' (Public). All have IP address management disabled and DHCP disabled.

Name	IP address management	Type	Subnet	Gateway	DHCP	Project
ForAgent	Disabled	Public			Disabled	admin
private	Enabled	Private	192.168.128.0/24		Enabled	admin
public	Disabled	Public			Disabled	admin

When deploying the appliance's virtual machine, you need to select these virtual networks.



8.9.3 Configuring user accounts in Virtuozzo Hybrid Infrastructure

To configure the virtual appliance, you need a Virtuozzo Hybrid Infrastructure user account. This account must have the **Administrator** role in the **Default** domain. For more information about users, refer to Managing Domain Users in the Virtuozzo Hybrid Infrastructure documentation. Ensure that you granted this account access to all projects in the **Default** domain.

To grant access to all projects in the Default domain

Run the following script in the Virtuozzo Hybrid Infrastructure cluster via the OpenStack Command-Line Interface. For more information on how to connect to this interface, refer to Connecting to OpenStack Command-Line Interface in the Virtuozzo Hybrid Infrastructure documentation.

```
su - vstoradmin
kolla-ansible post-deploy
exit
. /etc/kolla/admin-openrc.sh
openstack --insecure user set --project admin --project-domain Default --domain Default <username>
openstack --insecure role add --domain Default --user <username> --user-domain Default compute --inherited
```

Here, <username> is the Virtuozzo Hybrid Infrastructure account with the **Administrator** role in the **Default** domain. The virtual appliance will use this account in order to back up and restore the virtual machines in any child project under the **Default** domain.

Example

```
su - vstoradmin
kolla-ansible post-deploy
exit
. /etc/kolla/admin-openrc.sh
openstack --insecure user set --project admin --project-domain Default --domain Default johndoe
openstack --insecure role add --domain Default --user johndoe --user-domain Default compute --inherited
```

To manage backups for virtual machines in a domain that is different from the **Default** domain, run the following script as well.

To grant access to all projects in a different domain

```
su - vstoradmin
kolla-ansible post-deploy
exit
. /etc/kolla/admin-openrc.sh
openstack --insecure role add --domain <domain name> --inherited --user <username>
--user-domain Default admin
```

Here, <domain name> is the domain to the projects in which the <username> account will have access.

Example

```
su - vstoradmin
kolla-ansible post-deploy
exit
. /etc/kolla/admin-openrc.sh
openstack --insecure role add --domain MyNewDomain --inherited --user johndoe
--user-domain Default admin
```

8.9.4 Deploying the QCOW2 template

1. Log in to your Cyber Protection account.
2. Click **Devices > All devices > Add > Virtioozzo Hybrid Infrastructure**.
The .zip archive is downloaded to your machine.
3. Unpack the .zip archive. It contains a .qcow2 image file.
4. Log in to your Virtioozzo Hybrid Infrastructure account.
5. Add the .qcow2 image file to the Virtioozzo Hybrid Infrastructure compute cluster as follows:
 - On the **Compute > Virtual machines > Images** tab, click **Add image**.
 - In the **Add image** window, click **Browse**, and then select the .qcow2 file.
 - Specify the image name, select the **Generic Linux OS** type, and then click **Add**.
6. In the **Compute > Virtual machines > Virtual machines** tab, click **Create virtual machine**. A window will open where you need to specify the following parameters:
 - A name for the new virtual machine.
 - In **Deploy from**, choose **Image**.
 - In the **Images** window, select the .qcow2 image file of the appliance, and then click **Done**.
 - In the **Volumes** window, you don't need to add any volumes. The volume that is added automatically for the system disk is sufficient.
 - In the **Flavor** window, choose your desired combination of vCPUs and RAM, and then click **Done**. Usually, 2 vCPUs and 4 GiB of RAM are enough.
 - In the **Network interfaces** window, click **Add**, select the virtual network of type *public*, and then click **Add**. It will appear in the **Network interfaces** list.
If you use a setup with more than one physical network (and thus, with more than one virtual network of type public), repeat this step and select the virtual networks that you need. For an example of a possible network configuration, refer to Network configuration in Virtioozzo Hybrid Infrastructure.


7. Click **Done**.
8. Back in the **Create virtual machine** window, click **Deploy** to create and boot the virtual machine.

8.9.5 Configuring the virtual appliance

After deploying the virtual appliance, you need to configure it so that it can reach both the Virtuozzo Hybrid Infrastructure cluster that it will protect and the Cyber Protection cloud service.


To configure the virtual appliance

1. Log in to your Virtuozzo Hybrid Infrastructure account.
2. On the **Compute > Virtual machines > Virtual Machines** tab, select the virtual machine that you created. Then, click **Console**.
3. Configure the network interfaces of the appliance. There may be one or more interfaces to configure – it depends on the number of virtual networks that the appliance uses. Ensure that automatically assigned DHCP addresses (if any) are valid within the networks that your virtual machine uses, or assign them manually.




4. Specify the Virtuozzo cluster address and credentials:
 - DNS name or IP address of the Virtuozzo Hybrid Infrastructure cluster – this is the address of the management node of the cluster. The default port 5000 will be automatically set. If you use a different port, you need to specify it manually.
 - In the **User domain name** field, specify your domain in Virtuozzo Hybrid Infrastructure. For example, **Default**.

- In the **User name** and **Password** fields, enter the credentials for Virtuozzo Hybrid Infrastructure user account with **Administrator** role in the specified domain. For more information about users, roles, and domains, refer to Configuring user accounts in Virtuozzo Hybrid Infrastructure (p. 71).




- Specify the Cyber Protection management server address and credentials for accessing it.



To protect the virtual machines in the Virtuozone Hybrid Infrastructure cluster

- Log in to your Cyber Protection account.
- Navigate to **Devices > Virtuozone Hybrid Infrastructure > <your cluster> > Default project > admin** or find your machines in **Devices > All devices**.
- Select the desired machines and apply a protection plan for them.



8.10 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 Windows** setup program. The download link is available on the **Add devices** page in the service console.

Step 1: Generating a registration token

A registration token passes your identity to the setup program without storing your login and password for the service console. This enables you to register any number of machines under your account. For more security, a token has limited lifetime.

To generate a registration token

1. Sign in to the service console by using the credentials of the account to which the machines should be assigned.
2. Click **All devices > Add**.
3. Scroll down to **Registration token**, and then click **Generate**.
4. Specify the token lifetime, and then click **Generate token**.
5. Copy the token or write it down. Be sure to save the token if you need it for further use.

You can click **Manage active tokens** to view and manage the already generated tokens. Please be aware that for security reasons, this table does not display full token values.

Step 2: 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. Start the setup program.
4. Click **Create .mst and .msi files for unattended installation**.
5. Click **Specify** next to **Registration settings**, and then enter the token you generated.

You can change the method of registering the machine in the Cyber Protection service from **Use registration token** (default) to **Use credentials** or **Skip registration**. The **Skip registration** option presumes that you will register the machine at a later time.

6. Review or modify the installation settings that will be added to the .mst file, and then click **Proceed**.
7. In **Save the files to**, specify the path to the folder you created.
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.

Step 3: 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 or later).
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 or later:
 - 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 or later, 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 or later:
 - 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.

8.11 Updating agents

Agents with the following versions can be updated by using the web interface:

- Agent for Windows, Agent for VMware (Windows), Agent for Hyper-V: version 11.9.191 and later
- Agent for VMware (Virtual Appliance): version 12.5.23094 and later
- Agent for Virtuozzo Hybrid Infrastructure (Virtual Appliance): version 12.5.23094 and later
- Agent for Linux: version 11.9.179 and later
- Other agents: any version can be updated

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

To update from earlier agent versions, download and install the newest agent manually. To find the download links, click **All devices** > **Add**.

Prerequisites

On Windows machines, Cyber Protect features require Microsoft Visual C++ 2017 Redistributable. Please ensure that it is already installed on your machine or install it before updating the agent. After the installation, a restart may be required. The Microsoft Visual C++ Redistributable package can be found here

<https://support.microsoft.com/help/2999226/update-for-universal-c-runtime-in-windows>.

To update an agent by using the web interface

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**.

Note During the update, any backups that are in progress will fail.

To update Agent for VMware (Virtual Appliance) whose version is below 12.5.23094

1. Click **Settings** > **Agents** > the agent that you want to update > **Details**, and then examine the **Assigned virtual machines** section. You will need to re-enter these settings after the update.
 - a. Make note of the position of the **Automatic assignment** switch.
 - b. To find out what virtual machines are manually assigned to the agent, click the **Assigned:** link. The software displays the list of assigned virtual machines. Make note of the machines that have **(M)** after the agent name in the **Agent** column.
2. Remove Agent for VMware (Virtual Appliance), as described in "Uninstalling agents (p. 79)". In step 5, delete the agent from **Settings** > **Agents**, even though you are planning to install the agent again.
3. Deploy Agent for VMware (Virtual Appliance), as described in "Deploying the OVF template" (p. 64).
4. Configure Agent for VMware (Virtual Appliance), as described in "Configuring the virtual appliance" (p. 65).

If you want to reconstruct the locally attached storage, in step 7 do the following:

- a. Add the disk containing the local storage to the virtual appliance.
 - b. Click **Refresh** > **Create storage** > **Mount**.
 - c. The software displays the original **Letter** and **Label** of the disk. Do not change them.
 - d. Click **OK**.
5. Click **Settings** > **Agents** > the agent that you want to update > **Details**, and then reconstruct the settings that you made note of in step 1. If some virtual machines were manually assigned to the agent, assign them again as described in "Virtual machine binding" (p. 279).

Once the agent configuration is completed, the protection plans that were applied to the old agent are re-applied automatically to the new agent.
 6. The plans with application-aware backup enabled require the guest OS credentials to be re-entered. Edit these plans and re-enter the credentials.
 7. The plans that back up ESXi configuration require the "root" password to be re-entered. Edit these plans and re-enter the password.

To update the Cyber Protection definitions on a machine

1. Click **Settings > Agents**.
2. Select the machine on which you want to update the Cyber Protection definitions and click **Update definitions**. The machine must be online.

To assign the Updater role to an agent

1. Click **Settings > Agents**.
2. Select the machine to which you want to assign the Updater role (p. 80), click **Details**, then in the **Cyber Protection definitions** section, enable **Use this agent to download and distribute patches and updates**.

To clear cached data on an agent

1. Click **Settings > Agents**.
2. Select the machine on which you want to clear the cached data (outdated update files and patch management data) and click **Clear cache**.

8.12 Uninstalling agents

In Windows

If you want to remove individual product components (for example, one of the agents or Cyber Protection Monitor), run the **All agents for Windows** setup program, choose to modify the product, and clear the selection of the components that you want to remove. The link to the setup program is 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.

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 Cyber Protection Agent > Uninstall**.
3. [Optional] Select the **Remove the logs and configuration settings** check box.
If you are planning to install the agent again, keep this check box cleared. If you select the check box, the machine may be duplicated in the service 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 service 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.
If you are planning to install the agent again, keep this check box cleared. If you select the check box, the machine may be duplicated in the service 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 service console, click **Settings > Agents**, select the machine where the agent was installed, and then click **Delete**.

In macOS

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 service 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 service console, do the following:
 - a. Click **Settings > Agents**, select the virtual appliance, and then click **Delete**.
 - b. Click **Backup storage > Locations**, and then delete the location corresponding to the locally attached storage.

8.13 Security settings

To configure the general protection settings for Cyber Protection, go to **Settings > Protection** in the service console.

Automatic updates for components

Cyber Protection uses peer-to-peer technology for component updates to minimize network bandwidth traffic. You can choose one or more dedicated agents that will download updates from the Internet and distribute them among other agents in the network as peer-to-peer agents.

By default the **Use this agent to download and distribute patches and updates** option is disabled for agents, this means that all of the registered agents check for the latest updates and distribute them. If a user turns on the **Use this agent to download and distribute patches and updates** option for a certain agent, then this agent receives the Updater role and all other agents use this one for checking updates and their distribution. You must ensure that the agents with the Updater role are powerful enough, with high-speed stable Internet access, and they have enough disk space.

The update workflow is the following:

1. The agent with the Updater role checks by schedule the index file provided by the service provider to update the core components.
2. The agent with the Updater role starts to download and distribute updates to all agents.

To assign the Updater role to a protection agent

1. In the service console, go to **Settings > Agents**.
2. Select the machine to which you want to assign the Updater role.
3. Click **Details**, and then enable the **Use this agent to download and distribute patches and updates** option.

Updating the Cyber Protection definitions by schedule

On the **Schedule** tab, you can set up the schedule for automatic update of the Cyber Protection definitions for each of the components:

- Antimalware
- Vulnerability assessment
- Patch management

Schedule type:

- **Daily** – define on which days of a week to update definitions.
Start at – you can select at what time to update definitions.
- **Hourly** – define more granular hourly schedule for definition updates.
Run every – define the periodicity for running definition updates.
From ... To – define a specific time range within which the automatic definition updates will be performed.

Updating the Cyber Protection definitions on-demand

To update Cyber Protection Definitions for a particular machine on-demand

1. In the service console, go to **Settings > Agents**.
2. Select the machines on which you want to update the Cyber Protection definitions and click **Update definitions**.

Cache storage

Location of cached data:

- On Windows machines: **C:\ProgramData\Acronis\Agent\var\atp-downloader\Cache**
- On Linux machines: **/opt/acronis/var/atp-downloader/Cache**
- On MacOS machines: **/Library/Application Support/Acronis/Agent/var/atp-downloader/Cache**

In **Outdated update files and patch management data**, specify after what period to remove cached data.

Maximum cache storage size (GB) for agents:

- **Updater role** – define storage size for cache on the machines with the Updater role.
- **Other roles** – define storage size for cache on other machines.

Remote connection

Click **Remote desktop connection**, to enable the remote connection to machines via RDP client or HTML5 client. If it is disabled, then the **Connect via RDP client / Connect via HTML5 client** options will be hidden in the service console, and users will not be able to connect to machines remotely. This option affects all users of your organization.

Click **Share remote desktop connection**, to enable sharing the remote connection with users. As a result, the new option **Share remote connection** will appear in the right menu when you select a machine, and you will be able to generate a link to be shared with users for accessing the remote machine.

8.14 Changing quota for a device

The Cyber Protection service matches quotas and devices automatically. To find out which quota the device is assigned to, view the **Service quota** section in the device **Details**.

In rare cases, there may be a need to change this configuration manually. For example, a virtual machine is assigned to the **Workstations** quota. After you purchase the **Virtual machines** quota, you may want to assign the machine to this cheaper quota.

To change quota for a device

1. Select the device.
2. Click **Details**.
3. In **Service quota**, click **Change**.
4. Select the new quota or select **No quota**. **No quota** means that the currently used quota will be released, and you can assign another device to this quota.
5. Confirm your decision.

8.15 Cyber Protection services installed in your environment

Cyber Protection installs some or all of the following services, depending on the Cyber Protection options that you use.

Services installed in Windows

Service name	Purpose
Acronis Managed Machine Service	Provides backup, recovery, replication, retention, validation functionality
Acronis Scheduler2 Service	Executes scheduled tasks on certain events
Acronis Active Protection Service	Provides protection against ransomware
Acronis Cyber Protection Service	Provides antimalware protection


Services installed in macOS

Service name and location	Purpose
/Library/LaunchDaemons/com.acronis.aakore.plist	Serves for communication between the agent and management components
/Library/LaunchDaemons/com.acronis.cyber-protect-service.plist	Provides detection of malware
/Library/LaunchDaemons/com.acronis.mms.plist	Provides backup and recovery functionality
/Library/LaunchDaemons/com.acronis.schedule.plist	Executes scheduled tasks


9 Service console

The service console allows you to manage your devices and protection plans, and provides you with a nifty dashboard where you can find the most important information about your protection.

In the service console, you can change your settings, configure your reports, or check your backup storage. The console also gives you access to additional Cyber Protection services or features, such as File Sync & Share or Antivirus and Antimalware protection, Patch management, and Vulnerability assessment. Depending on the Cyber Protection edition, their type and number may vary.



For the **Devices** section, you can choose between the simple and the table view. To switch between them, click the corresponding icon in the top right corner.



The simple view shows only a few machines.

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

All devices							+ Add	Table View Icon	Grid View Icon	Help Icon	User Icon
Search							Selected: 1 / Loaded: 2 / Total: 2				
☐	Type	Name ↑	Account	Status	Last backup	Next backup					
☐	VM	exw.win8.dcon.local	an3	Not protected	Never	Not scheduled					
☐	VM	WIN-IIS48KGULO3	an3	Not protected	Never	Not scheduled					

Protect
 Recovery
 Connect via RDP client
 Connect via HTML5 client

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

To delete a machine from the service console

1. Select the check box next to the desired machine.
2. Click **Delete**, and then confirm your choice.

Important Deleting a machine from the service console does not uninstall the protection agent on it and does not delete the protection plans applied to this machine. The backups of the deleted machine will also be kept.

VMware or Hyper-V virtual machines and ESXi hosts can be backed up by an agent that is not installed on them. You cannot delete such machines individually. To delete them, you need to find and delete the machine on which the respective Agent for VMware or Agent for Hyper-V is installed.

To delete a virtual machine or ESXi host without an agent

1. Under **Devices**, select **All devices**.
2. Click the gear icon in the upper right corner and enable the **Agent** column.

All devices							+ Add	Table View Icon	Grid View Icon	Help Icon	User Icon
Search							Loaded: 2 / Total: 2				
☐	Type	Name ↑	Account	Status	Last backup	Next backup					
☐	VM	WIN-2A1NUKBHD7U	an3	OK	Nov 08, 2019 03:29:1...						
☐	VM	WIN-SN20102R9Q7	an3	OK	Nov 08, 2019 03:29:1...						

✓ Account
 ✓ Status
 ✓ Last backup
 ✓ Next backup
 Plan
 Comment
Agent
 IP addresses
 Operating system


3. In the **Agent** column, check the name of the machine where the respective agent is installed.
4. Delete this machine from the service console. This will also delete all of the machines that are backed up by its agent.
5. Uninstall the agent from the deleted machine as described in "Uninstalling agents" (p. 79).

10 Voice control for operations in the console


You can enable voice control to perform different operations in the service console by using your voice. This functionality is available only in the Cyber Protect editions.

Before turning on voice control, ensure that access to the microphone is enabled.

To turn on voice control, click the user icon in the upper right corner in the service console, click **Voice control**, and then turn on the switch. After that, every time you pronounce a command, you will see the recognized text in the lower left corner in the service console.



To turn off voice control, either hover you mouse on the microphone icon in the lower left corner and click the red button or click the user icon in the upper right corner in the service console, click **Voice control**, and then turn off the switch.



Limitations:

- Voice control functionality is supported only in English. This option is not available for other languages.
- Voice control functionality is supported only in the Google Chrome browser for desktop. It is not supported in mobile browsers.
- Voice control functionality will only work with the https protocol. If a user tries to access the service console by using the http protocol, the user will not be able to enable the microphone, therefore will not be able to turn on the **Voice control** option at all.

To open the help and view all of the available voice commands, click the question mark icon in the upper right corner, and then select **Voice control help**.

Voice command	Description
General	
Help	Open a help modal window
Scroll up	To scroll up the page
Scroll down	To scroll down the page
Up (or Previous)	Move to the upper row in a table
Down (or Next)	Move to the lower row in a table
Close	Close a modal window
Dashboard > Overview	
Overview	Go to the Overview section

[Status name]	Drill down by the machine status in the Protection status widget
Dashboard > Alerts	
Alerts	Go to the Alerts section
[Name of alert]	Filter by the alert type by saying the corresponding alert name. For example, "Backup Failed", "License Expired", "Machine Offline"
Clear all	Clear all alerts
Dashboard > Activities	
Activities	Go to the Activities section
Dashboard > Threat feed	
Threat feed	Go to the Threat feed section
Devices > All devices	
Devices	Go to the All devices section
Protect	Open the list of protection plans for a device
[Plan name]	Select the protection plan for a device. For example, "Entire machine to cloud", "Cloud to cloud backup"
Apply	Apply the certain protection plan to a device
Cancel backup	Stop running a backup for a certain device
Back up now	Run backup on a certain device
Recovery	Open recovery points for a device
Details	View the device details
Activities	Open activities for a device
Alerts	Open alerts for a device
Antivirus and Antimalware protection	
Quarantine	Go to the Quarantine section
Allowlist	Go to the Allowlist section
Software management	
Patches	Go to the Patches section
Vulnerabilities	Go to the Vulnerabilities section

Common scenarios:

- To select a protection plan for a device, say Protect > <Plan_name>
- To apply a specific protection plan to a device, say Protect > <Plan_name> > Apply
- To run backup on a specific device, say Protect > <Plan_name> > Back up now
- To cancel backup on a specific device, say Protect > <Plan_name> > Cancel backup

11 Device groups

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

Device groups are designed for convenient management of a large number of registered devices.

You can apply a protection plan to a group. Once a new device appears in the group, the device becomes protected by the plan. If a device is removed from the group, the device will no longer be protected by the plan. A plan that is applied to a group cannot be revoked from a member of the group, only from the group itself.

Only devices of the same type can be added to a group. For example, under **Hyper-V** you can create a group of Hyper-V virtual machines. Under **Machines with agents**, you can create a group of machines with installed agents. Under **All devices**, you cannot create a group.

A single device can be a member of more than one group.

Built-in groups

Once a device is registered, it appears in one of the built-in root groups on the **Devices** tab.

Root groups cannot be edited or deleted. You cannot apply plans to root groups.

Some of the root groups contain built-in sub-root groups. These groups cannot be edited or deleted. However, you can apply plans to sub-root built-in groups.

Custom groups

Protecting all devices in a built-in group with a single protection plan may not be satisfactory because of the different roles of the machines. The backed-up data is specific for each department; some data has to be backed up frequently, other data is backed up twice a year. Therefore, you may want to create various protection plans applicable to different sets of machines. In this case, consider creating custom groups.

A custom group can contain one or more nested groups. Any custom group can be edited or deleted. There are the following types of custom groups:

- **Static groups**

Static groups contain the machines that were manually added to them. The static group content never changes unless you explicitly add or delete a machine.

Example: You create a custom group for the accounting department and manually add the accountants' machines to this group. Once you apply a protection plan to the group, the accountants' machines become protected. If a new accountant is hired, you will have to add the new machine to the group manually.

- **Dynamic groups**

Dynamic groups contain the machines added automatically according to the search criteria specified when creating a group. The dynamic group content changes automatically. A machine remains in the group while it meets the specified criteria.

Example 1: The host names of the machines that belong to the accounting department contain the word "accounting". You specify the partial machine name as the group membership criterion and apply a protection plan to the group. If a new accountant is hired, the new machine will be added to the group as soon as it is registered, and thus will be protected automatically.

Example 2: The accounting department forms a separate Active Directory organizational unit (OU). You specify the accounting OU as the group membership criterion and apply a protection plan to the group. If a new accountant is hired, the new machine will be added to the group as soon as it is registered and added to the OU (regardless of which comes first), and thus will be protected automatically.

11.1 Creating a static group

1. Click **Devices**, and then select the built-in group which contains the devices for which you want to create a static group.
2. Click the gear icon next to the group in which you want to create a group.
3. Click **New group**.
4. Specify the group name, and then click **OK**.
The new group appears in the groups tree.

11.2 Adding devices to static groups

1. Click **Devices**, and then select one or more devices that you want to add to a group.
2. Click **Add to group**.
The software displays a tree of groups to which the selected device can be added.
3. If you want to create a new group, do the following. Otherwise, skip this step.
 - a. Select the group in which you want to create a group.
 - b. Click **New group**.
 - c. Specify the group name, and then click **OK**.
4. Select the group to which you want to add the device, and then click **Done**.

Another way to add devices to a static group is to select the group and click **Add devices**.

11.3 Creating a dynamic group

1. Click **Devices**, and then select the group which contains the devices for which you want to create a dynamic group.

Note You cannot create dynamic groups for the All devices group.

2. Search for devices by using the search field. You can use multiple search criteria and operators described below.
3. Click **Save as** next to the search field.

Note Some search criteria are not supported for group creation. See the table in section Search criteria below.

4. Specify the group name, and then click **OK**.

Search criteria

The following table summarizes the available search criteria.

Criterion	Meaning	Search query examples	Supported
name	▪ Host name for physical machines ▪ Name for virtual machines ▪ Database name ▪ Email address for mailboxes	name = 'en-00'	Yes

Criterion	Meaning	Search query examples	Sup
comment	Comment for a device. Default value: - For physical machines running Windows, the computer description taken from the computer properties in Windows. This value is updated automatically every 15 minutes. - Empty for other devices. To view the comment, under Devices, select the device, click Details, and then locate the Comment section. To add or change the comment manually, click Add or Edit. In this case, automatic update will stop working. To allow automatic updates again, clear the comment that you have added. To refresh the comment field for your devices, restart the Managed Machine Service in Windows Services or run the following commands at the command prompt: <pre>net stop mms net start mms</pre>	comment = 'important machine' comment = '' (all machines without a comment)	Yes
ip	IP address (only for physical machines)	ip RANGE ('10.250.176.1', '10.250.176.50')	Yes
memorySize	RAM size in megabytes (MiB)	memorySize < 1024	Yes
insideVm	Virtual machine with an agent inside. Possible values: true false	insideVm = true	Yes
osName	Operating system name.	osName LIKE '%Windows XP%'	Yes
osType	Operating system type. Possible values: 'windows' 'linux' 'macosx'	osType IN ('linux', 'macosx')	Yes

Criterion	Meaning	Search query examples	Sup
osProductType	The operating system product type. Possible values: ▪ 'dc' Stands for Domain Controller. ▪ 'server' ▪ 'workstation' Note When the domain controller role is assigned on a Windows server, the osProductType changes from "server" to "dc". Such machines will be not included in search results for filter "osProductType='server'".	osProductType = 'server'	Yes
tenant	The name of the unit to which the device belongs.	tenant = 'Unit 1'	Yes
tenantId	The identifier of the unit to which device belongs. To get the unit ID, under Devices, select the device, click Details > All properties. The ID is shown in the ownerId field.	tenantId = '3bfe6ca9-9c6a-4953-9cb2-a1323f454fc9'	Yes
state	Device state. Possible values: ▪ 'idle' ▪ 'interactionRequired' ▪ 'canceling' ▪ 'backup' ▪ 'recover' ▪ 'install' ▪ 'reboot' ▪ 'failback' ▪ 'testReplica' ▪ 'run_from_image' ▪ 'finalize' ▪ 'failover' ▪ 'replicate' ▪ 'createAsz' ▪ 'deleteAsz' ▪ 'resizeAsz'	state = 'backup'	No

Criterion	Meaning	Search query examples	Sup
protectedByPlan	Devices that are protected by a protection plan with a given ID. To get the plan ID, click Plans > Backup , select the plan, click on the diagram in the Status column, and then click on a status. A new search with the plan ID will be created.	protectedByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
okByPlan	Devices that are protected by a protection plan with a given ID and have an OK status.	okByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
errorByPlan	Devices that are protected by a protection plan with a given ID and have an Error status.	errorByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
warningByPlan	Devices that are protected by a protection plan with a given ID and have a Warning status.	warningByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
runningByPlan	Devices that are protected by a protection plan with a given ID and have a Running status.	runningByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
interactionByPlan	Devices that are protected by a protection plan with a given ID and have an Interaction Required status.	interactionByPlan = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	No
ou	Machines that belong to the specified Active Directory organizational unit.	ou IN ('RnD', 'Computers')	Yes
id	Device ID. To get the device ID, under Devices , select the device, click Details > All properties . The ID is shown in the id field.	id != '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	Yes
lastBackupTime*	The date and time of the last successful backup. The format is 'YYYY-MM-DD HH:MM'.	lastBackupTime > '2016-03-11' lastBackupTime <= '2016-03-11 00:15' lastBackupTime is null	No
lastBackupTryTime*	The time of the last backup attempt. The format is 'YYYY-MM-DD HH:MM'.	lastBackupTryTime >= '2016-03-11'	No
nextBackupTime*	The time of the next backup. The format is 'YYYY-MM-DD HH:MM'.	nextBackupTime >= '2016-03-11'	No
agentVersion	Version of the installed protection agent.	agentVersion LIKE '12.0.*'	Yes
hostId	Internal ID of the protection agent. To get the protection agent ID, under Devices , select the machine, click Details > All properties . Use the "id" value of the agent property.	hostId = '4B2A7A93-A44F-4155-BDE3-A023C57C9431'	Yes

Criterion	Meaning	Search query examples	Sup
resourceType	Resource type. Possible values: 'machine' 'virtual_machine.vmwesx' 'virtual_machine.mshyperv' 'virtual_machine.rhev' 'virtual_machine.kvm' 'virtual_machine.xen'	resourceType = 'machine' resourceType in ('mssql_aag_database', 'mssql_database')	Yes

Note If you skip the hour and minutes value, the start time is considered to be YYYY-MM-DD 00:00, and the end time is considered to be YYYY-MM-DD 23:59:59. For example, lastBackupTime = 2020-02-20, means that the search results will include all backups from the interval
lastBackupTime >= 2020-02-20 00:00 and lastBackup time <= 2020-02-20 23:59:59

Operators

The following table summarizes the available operators.

Operator	Meaning	Examples
AND	Logical conjunction operator.	name like 'en-00' AND tenant = 'Unit 1'
OR	Logical disjunction operator.	state = 'backup' OR state = 'interactionRequired'
NOT	Logical negation operator.	NOT(osProductType = 'workstation')
LIKE 'wildcard pattern'	This operator is used to test if an expression matches the wildcard pattern. This operator is case-insensitive. The following wildcard operators can be used: * or % The asterisk and the percent sign represent zero, one, or multiple characters _ The underscore represents a single character	name LIKE 'en-00' name LIKE '*en-00' name LIKE '*en-00*' name LIKE 'en-00_'
IN (<value1>, ... <valueN>)	This operator is used to test if an expression matches any value in a list of values. This operator is case-sensitive.	osType IN ('windows', 'linux')
RANGE(<startin g_value>, <ending_value>)	This operator is used to test if an expression is within a range of values (inclusive).	ip RANGE('10.250.176.1', '10.250.176.50')

11.4 Applying a protection plan to a group

- Click **Devices**, and then select the built-in group that contains the group to which you want to apply a protection plan.
The software displays the list of child groups.
- Select the group to which you want to apply a protection plan.

3. Click **Group backup**.

The software displays the list of protection plans that can be applied to the group.

4. Do one of the following:

- Expand an existing protection plan, and then click **Apply**.
- Click **Create new**, and then create a new protection plan as described in "Protection plan (p. 375)".

12 Multitenancy support

Cyber Protection supports multitenancy. This means that a tenant administrator/user can manage objects that are related to their tenant or its sub-tenants (units). An administrator/user from a unit cannot manage objects of the parent tenant.

For example, a customer administrator created a protection plan and applied it to a machine. A customer administrator can also manage protection plans created by a unit administrator. However, a unit administrator cannot manage the protection plan created by the customer administrator. A unit administrator can create its own protection plan that does not conflict with the plan of the customer administrator.

Multitenancy also implies that a tenant administrator/user can see all objects that are related to this tenant or its sub-tenants (units). An administrator/user from a unit cannot see objects of the parent tenant.


For example, the data shown in the patch list, quarantine, threat feed, alerts, and activities will be shown only for the current tenant and its sub-tenants. The data related to the parent tenant is not shown.

13 Protection plan and modules

The protection plan is a plan that combines several data protection modules including

- Backup (p. 109) – allows you to back up your data sources to local or cloud storage.
- Antivirus and Antimalware protection (p. 328) – allows you to check your machines with the built-in antimalware solution.
- URL filtering (p. 338) – allows you to protect your machines from threats coming from the Internet by blocking access to malicious URLs and content to be downloaded.
- Windows Defender Antivirus (p. 336) – allows you to manage the settings of Windows Defender Antivirus to protect your environment.
- Microsoft Security Essentials (p. 338) – allows you to manage the settings of Microsoft Security Essentials to protect your environment.
- Vulnerability assessment (p. 350) – automatically checks the Microsoft and third-party products installed on your machines for vulnerabilities and notifies you about them.
- Patch management (p. 354) – allows you to install patches and updates for the Microsoft and third-party products on your machines to close the discovered vulnerabilities.
- Data protection map (p. 372) – allows you to discover the data in order to monitor the protection status of important files.

The protection plan allows you to protect your data sources completely from external and internal threats. By enabling and disabling different modules and setting up the module settings, you can build flexible plans satisfying various business needs.



13.1 Creating a protection plan

A protection plan can be applied to multiple machines at the time of its creation, or later. When you create a plan, the system checks the operating system and the device type (for example, workstation, virtual machine, etc.) and shows only those plan modules that are applicable to your devices.

A protection plan can be created in two ways:


- In the **Devices** section – when you select the device or devices to be protected and then create a plan for them.

- In the **Plans** section – when you create a plan and then select the machines to be applied to (p. 375).

Let's consider the first way.

To create the first protection plan

1. In the service console, go to **Devices > All devices**.
2. Select the machines that you want to protect.
3. Click **Protect**, and then click **Create plan**. You will see the protection plan with the default settings.



4. [Optional] To modify the protection plan name, click on the pencil icon next to the name.
5. [Optional] To enable or disable the plan module, click the switch next to the module name.


6. [Optional] To configure the module parameters, click the corresponding section of the protection plan.
7. When ready, click **Create**.

The Backup, Antivirus and Antimalware protection, Vulnerability assessment, Patch management, and Data protection map modules can be performed on demand by clicking **Run now**.

13.2 Default protection plans

Three preconfigured plans, available by default, ensure quick protection for specific workloads:

- **Office workers (Acronis Antivirus)**
This plan is optimized for users working in the office and having a preference to use the Acronis antivirus software.
- **Office workers (third-party Antivirus)**
This plan is optimized for users working in the office and having a preference to use a third-party antivirus software. The main difference is that this plan has the **Antivirus and Antimalware protection** module and **Active Protection** disabled.
- **Remote workers**
This plan is optimized specifically for users working remotely. It has more frequent tasks (such as backup, antimalware protection, vulnerability assessment), stricter protection actions, and optimized performance and power options.



To apply a default protection plan

1. In the service console, go to **Devices > All devices**.
2. Select the machines that you want to protect.
3. Click **Protect**.
4. Select one of the default plans, and then click **Apply**.

Tip You can also configure your own protection plan (p. 95) by clicking **Create plan**.

To modify an applied default protection plan

1. In the service console, go to **Plans > Protection**.
2. Select the plan that you want to modify, and then click **Edit**.
3. Modify the modules that are included in this plan, or their options, and then click **Save**.

Important Some settings cannot be changed for an existing protection plan.

Default plan options

The preconfigured plans use the default options for each module*, with the following modifications:

Modules and options/Plan	Office workers (Acronis Antivirus)	Office workers (third-party Antivirus)	Remote workers
Backup (p. 109)			
What to back up	Entire machine	Entire machine	Entire machine
Continuous data protection (CDP)	Disabled	Disabled	Enabled
Where to back up	Cloud storage	Cloud storage	Cloud storage
Backup scheme	Always incremental (single-file)	Always incremental (single-file)	Always incremental (single-file)
Schedule	Default daily schedule	Default daily schedule	Daily: Monday to Friday at 12:00 PM Additionally enabled options and start conditions: ▪ If the machine is turned off, run missed tasks at the machine startup ▪ Wake up from the sleep or hibernate mode to start a scheduled backup ▪ Save battery power: Do not start when on battery ▪ Do not start when on metered connection
How long to keep	Monthly: 12 months Weekly: 4 weeks Daily: 7 days	Monthly: 12 months Weekly: 4 weeks Daily: 7 days	Monthly: 12 months Weekly: 4 weeks Daily: 7 days
Backup options	Default options	Default options	Default options, plus: Performance and backup window (the green set): ▪ CPU priority: Low ▪ Output speed: 50%
Antivirus and Antimalware protection			

Schedule scan	Scan type: Quick	n/a	Scan type: Full Additionally enabled options and start conditions: ▪ If the machine is turned off, run missed tasks at the machine startup ▪ Wake up from the sleep or hibernate mode to start a scheduled backup ▪ Save battery power: Do not start when on battery
URL filtering (p. 338)			
Malicious websites access	Always ask user	Always ask user	Block
Vulnerability assessment (p. 350)			
	Default	Default	Default
Patch management (p. 354)			
Schedule	Default	Default	Daily: Monday to Friday at 02:20PM
Pre-update backup	Off	Off	On
Data protection map (p. 372)			

Extensions	Default options	Default options	Default options, plus:
			Images: ▪ .bmp ▪ .png ▪ .ico ▪ .wbmp ▪ .gif ▪ .bmp ▪ .xcf ▪ .psd ▪ .tiff ▪ .jpeg, .jpg ▪ .dwg Audio: ▪ .wav ▪ .aif, .aifc, .aiff ▪ .au, .snd ▪ .mid, .midi ▪ .mid ▪ .mpga, .mp3 ▪ .oga ▪ .flac ▪ .oga ▪ .oga ▪ .opus ▪ .oga ▪ .spx ▪ .oga ▪ .ogg ▪ .ogx ▪ .ogx ▪ .mp4

* The number of modules in the default protection plan may vary between the different editions of the Cyber Protection service.

13.3 Resolving plan conflicts

A protection plan can be in the following statuses:

- **Active** - a plan that is assigned to devices and executed on them.
- **Inactive** - a plan that is assigned to devices but disabled and not executed on them.

Applying several plans to a device

You can apply several protection plans to a single device. As a result, you will get a combination of different protection plans assigned on a single device. For example, you may apply a plan that has only the Antivirus and Antimalware protection module enabled in the plan and another plan that contains only the backup module. The protection plans can be combined only if they do not have intersecting modules. If there are similar enabled modules in the applied protection plans, you must resolve conflicts between such modules.

Resolving plan conflicts

Plan conflicts with already applied plans

When you create a new plan on a device or devices with already applied plans that conflict with the new plan, you can resolve a conflict with one of the following ways:

- Create a new plan, apply it, and disable all already applied conflicting plans.
- Create a new plan and disable it.

When you edit a plan on a device or devices with already applied plans that conflict with the changes made, you can resolve a conflict with one of the following ways:

- Save changes to the plan and disable all already applied conflicting plans.
- Save changes to the plan and disable it.

A device plan conflicts with a group plan

If a device is included in a group of devices with an assigned group plan, and you try to assign a new plan to a device, then the system will ask you to resolve the conflict by doing one of the following:

- Remove a device from the group and apply a new plan to the device.
- Apply a new plan to the whole group or edit the current group plan.

License issue

The assigned quota on a device must be appropriate for the protection plan to be performed, updated, or applied. To resolve the license issue, do one of the following:

- Disable the modules that are unsupported by the assigned quota and continue using the protection plan.
- Change the assigned quota manually: go to **Devices > <particular_device> > Details > Service quota**, then revoke the existing quota and assign a new one.

13.4 Operations with protection plans

Available actions with a protection plan

You can perform the following actions with a protection plan:

- Rename a plan.
- Enable/disable modules and edit each module setting.

- Enable/disable a plan.

A disabled plan will not be carried out on the device to which it is applied.

This action is convenient for administrators who intend to protect the same device with the same plan later. Since the plan is not revoked from the device, to restore its protection, an administrator must only re-enable the plan.

- Apply a plan to a device or a group of devices.
- Revoke a plan from a device.

A revoked plan is not applied to a device anymore.

This action is convenient for administrators who do not need to protect quickly the same device with the same plan again. To restore the protection a revoked plan, an administrator must know the name of the plan, select it from the list of available plans, and then re-apply the plan to the desired device.

- Import/export a plan.

Note You can import protection plans created only in Cyber Protection 9.0. Plans created in previous product versions are incompatible with version 9.0.

- Delete a plan.

To apply an existing protection plan

1. Select the machines that you want to protect.
2. Click **Protect**. If a protection plan is already applied to the selected machines, click **Add plan**.
3. The software displays previously created protection plans.
4. Select a protection plan to apply and click **Apply**.

To edit a protection plan

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

To revoke a protection plan from machines

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

To delete a protection plan

1. Select any machine to which the protection plan that you want to delete is applied.
2. Click **Protect**.
3. If several protection plans are applied to the machine, select the protection plan that you want to delete.

4. Click the ellipsis icon next to the protection plan name, and then click **Delete**.
As a result, the protection plan is revoked from all of the machines and completely removed from the web interface.

14 #CyberFit Score for machines

#CyberFit Score provides you with a security assessment and scoring mechanism that evaluates the security posture of your machine. It identifies security gaps in the IT environment and open attack vectors to endpoints and provides recommended actions for improvements in the form of a report. This feature is available in all three Cyber Protect editions.

The #CyberFit Score functionality is supported on:

- Windows 7 (first version) and later versions
- Windows Server 2008 R2 and later versions

How it works

The protection agent that is installed on a machine performs a security assessment and calculates the #CyberFit Score for the machine. The #CyberFit Score of a machine is automatically periodically recalculated.

#CyberFit scoring mechanism

The #CyberFit Score for a machine is calculated, based on the following metrics:

- Antimalware protection 0-275
- Backup protection 0-175
- Firewall 0-175
- Virtual private network (VPN) 0-75
- Full disk encryption 0-125
- Network security 0-25

The maximum #CyberFit Score for a machine is 850.

Metric	What is assessed?	Recommendations to users	Scoring
Antimalware	The agent checks whether antimalware software is installed on a machine.	Findings: ▪ You have antimalware protection enabled (+275 points) ▪ You don't have antimalware protection, your system may be at risk (0 points) Recommendations provided by #CyberFit Score: You should have an antimalware solution installed and enabled on your machine to stay protected from security risks. You should refer to websites such as AV-Test or AV-Comparatives for a list of recommended antimalware solutions.	275 - antimalware software is installed on a machine 0 - no antimalware software is installed on a machine

Backup	The agent checks that a backup solution is installed on a machine.	Findings: ■ You have a backup solution protecting your data (+175 points) ■ No backup solution was found, your data may be at risk (0 points) Recommendations provided by #CyberFit Score: You are recommended to back up your data regularly to prevent data loss or ransomware attacks. Below are some backup solutions that you should consider using: ■ Acronis Cyber Protect / Cyber Backup / True Image ■ Windows Server Backup (Windows Server 2008 R2 and later)	175 - a backup solution is installed on a machine 0 - no backup solution is installed on a machine
Firewall	The agent checks whether a firewall is available and enabled in your environment. The agent does the following: 1. Checks Windows Firewall and Network Protection whether a public firewall is turned on. 2. Checks Windows Firewall and Network Protection whether a private firewall is turned on. 3. Checks for a 3-rd party firewall solution/agent if Windows public and private firewalls are disabled.	Findings: ■ You have a firewall enabled for public and private networks, or a 3-rd party firewall solution is found (+175 points) ■ You have a firewall enabled only for public networks (+100 points) ■ You have a firewall enabled only for private networks (+75 points) ■ You have no firewall enabled, your network connection is not secure (0 points) Recommendations provided by #CyberFit Score: It is recommended to enable firewall for your public and private networks to improve your security protection against malicious attacks on your system. Below are provided detailed guides on setting-up your Windows firewall, depending on your security needs and network architecture: Guides for end-users/employees: How to set up Windows Defender Firewall on your PC How to set up Windows Firewall on your PC Guides for system administrators and engineers: How to deploy Windows Defender Firewall with Advanced Security How to create Advanced Rules in Windows Firewall	100 - Windows public firewall is enabled 75 - Windows private firewall is enabled 175 - Windows public and private firewall are enabled OR a third-party firewall solution is enabled 0 - neither a Windows firewall, nor a third-party firewall solution are enabled


Virtual Private Network (VPN)	The agent checks whether a VPN solution is installed on a machine and whether the VPN is enabled and running.	Findings: ■ You have a VPN solution and can safely receive and send data across public and shared networks (+75 points) ■ No VPN solution was found, your connection to public and shared networks is not secure (0 points) Recommendations provided by #CyberFit Score: It is recommended to use VPN to access your corporate network and confidential data. It is critical to use a VPN to keep your communications safe and private, especially if you use complimentary Internet access from a cafe, library, airport, or elsewhere. Below are some VPN solutions that you should consider using: ■ Acronis Business VPN ■ OpenVPN ■ Cisco AnyConnect ■ NordVPN ■ TunnelBear ■ ExpressVPN ■ PureVPN ■ CyberGhost VPN ■ Perimeter 81 ■ VyprVPN ■ IPVanish VPN ■ Hotspot Shield VPN ■ Fortigate VPN ■ ZYXEL VPN ■ SonicWall GVPN ■ LANCOM VPN	75 - VPN is enabled and running 0 - VPN is not enabled

Disk encryption	The agent checks whether a machine has disk encryption enabled. The agent checks whether Windows BitLocker is turned on.	Findings: ■ You have full disk encryption enabled, your machine is protected against physical tampering (+125 points) ■ Only some hard drives are encrypted, your machine may be at risk from physical tampering (+75 points) ■ No disk encryption was found, your machine is at risk from physical tampering (0 points) Recommendations provided by #CyberFit Score: It is recommended to turn on Windows BitLocker to improve protection of your data and files. Guide: How to turn on device encryption on Windows	125 - all disks are encrypted 75 - at least one of your disks is encrypted but there are also unencrypted disks 0 - no disks are encrypted
Network security (outgoing NTLM traffic to remote servers)	The agent checks whether a machine has restricted outgoing NTLM traffic to remote servers.	Findings: ■ Outgoing NTLM traffic to remote servers is denied, your credentials are protected (+25 points) ■ Outgoing NTLM traffic to remote servers is not denied, your credentials may be vulnerable to exposure (0 points) Recommendations provided by #CyberFit Score: It is recommended to deny all outgoing NTLM traffic to remote servers for better security protection. You can find information on how to change the NTLM settings and add exceptions by following the link below. Guide: Restrict outgoing NTLM traffic to remote servers	25 - outgoing NTLM traffic is set to DenyAll 0 - outgoing NTLM traffic is set to another value

Based on the summed points awarded to each metric, the total #CyberFit Score of a machine can fit one of the following ratings that reflect the endpoint's level of protection:

- 0 - 579 - Poor
- 580 - 669 - Fair
- 670 - 739 - Good
- 740 - 799 - Very good
- 800 - 850 - Excellent

You can see the #CyberFit Score for your machines in the service console: go to **Devices > All devices**. In the list of devices, you can see the **#CyberFit Score** column. You can also run the #CyberFit Score scan for a machine to check its security posture.




You can also get information about the #CyberFit Score in the corresponding widget (p. 381) and report (p. 393) pages.

14.1 Running a #CyberFit Score scan


To run a #CyberFit Score scan

1. In the service console, go to **Devices**.
2. Select the machine and click **#CyberFit Score**.
3. If the machine has never been scanned before, then click **Run a first scan**.

4. After the scan is completed, you will see the total #CyberFit Score for the machine along with the scores of each of the six assessed metrics - Antimalware, Backup, Firewall, Virtual Private Network (VPN), Disk encryption, and NT LAN Manager (NTLM) traffic.



5. To check how to increase the score of each metric for which the security configurations could be improved, expand the corresponding section and read the recommendations.



6. After addressing the recommendations, you can always recalculate the #CyberFit Score of the machine by clicking on the arrow button right under the total #CyberFit Score.

15 Backup and recovery

The backup module enables backup and recovery of physical and virtual machines, files, and databases to local or cloud storage.

15.1 Backup

A protection plan with the Backup module enabled is a set of rules that specify how the given data will be protected on a given machine.

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

To create the first protection plan with the Backup module enabled

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

The software displays protection plans that are applied to the machine. If the machine does not have any plans already assigned to it, then you will see the default protection plan that can be applied. You can adjust the settings as needed and apply this plan or create a new one.

- To create a new plan, click **Create plan**. Enable the **Backup** module and unroll the settings.

New protection plan (2)

Cancel

Create

Backup

Entire machine to Cloud storage, Monday to Friday at 05:45 PM

What to back up

Entire machine

Continuous data protection (CDP)

Where to back up

Cloud storage

Schedule

Monday to Friday at 05:45 PM

How long to keep

Monthly: 6 months

Weekly: 4 weeks

Daily: 7 days

Encryption

Application backup

Disabled

Backup options


Change

- [Optional] To modify the protection plan name, click the default name.
- [Optional] To modify the Backup module parameters, click the corresponding setting of the protection plan panel.
- [Optional] To modify the backup options, click **Change** next to **Backup options**.
- Click **Create**.

To apply an existing protection plan

- Select the machines that you want to back up.
- Click **Protect**. If a common protection plan is already applied to the selected machines, click **Add plan**.

The software displays previously created protection plans.



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

15.2 Protection plan cheat sheet

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

WHAT TO BACK UP	ITEMS TO BACK UP Selection methods	WHERE TO BACK UP	SCHEDULE Backup schemes	HOW LONG TO KEEP
Disks/volumes (physical machines (p. 398))	Direct selection Policy rules File filters (p. 153)	Cloud (p. 123) Local folder (p. 123) Network folder (p. 123) NFS (p. 123)* Secure Zone (p. 123)**	Always incremental (Single-file) (p. 127) Always full (p. 127) Weekly full, Daily incremental (p. 127) Monthly full, Weekly differential, Daily incremental (GFS) (p. 127) Custom (F-D-I) (p.	By backup age (single rule/per backup set) (p. 136) By number of backups (p. 136) By total size of backups (p. 136)**** Keep indefinitely

WHAT TO BACK UP		ITEMS TO BACK UP Selection methods	WHERE TO BACK UP	SCHEDULE Backup schemes	HOW LONG TO KEEP
Disks/volumes (virtual machines (p. 400))		Policy rules File filters (p. 153)	Cloud (p. 123) Local folder (p. 123) Network folder (p. 123) NFS (p. 123)*	127)	(p. 136)
Files (physical machines only (p. 398))		Direct selection Policy rules File filters (p. 153)	Cloud (p. 123) Local folder (p. 123) Network folder (p. 123) NFS (p. 123)* Secure Zone (p. 123)**	Always incremental (Single-file) (p. 127)*** Always full (p. 127) Weekly full, Daily incremental (p. 127) Monthly full, Weekly differential, Daily incremental (GFS) (p. 127) Custom (F-D-I) (p. 127)	
ESXi configuration		Direct selection	Local folder (p. 123) Network folder (p. 123) NFS (p. 123)*		
Websites (files and MySQL databases)		Direct selection (p. 264)	Cloud (p. 123)	—	
System state		Direct selection (p. 117)	Cloud (p. 123) Local folder (p. 123) Network folder (p. 123)	Always full (p. 127) Weekly full, daily incremental (p. 127) Custom (F-I) (p. 127)	
SQL databases		Direct selection			
Exchange databases		Direct selection (p. 206)			
Microsoft Office 365	Mailboxes (local Agent for Office 365)	Direct selection	Cloud (p. 123) Local folder (p. 123) Network folder (p. 123)	Always incremental (Single-file) (p. 127)	
	Mailboxes (cloud Agent for Office 365)	Direct selection (p. 236)	Cloud (p. 123)	—	
	Public folders	Direct selection (p. 236)			
	OneDrive files	Direct selection (p. 240) Policy rules (p. 240)			
	SharePoint Online data	Direct selection (p. 244) Policy rules (p. 244)			
	Teams	Direct selection (p. 246)			
G Suite	Gmail mailboxes	Direct selection (p. 255)	Cloud (p. 123)	—	

WHAT TO BACK UP		ITEMS TO BACK UP Selection methods	WHERE TO BACK UP	SCHEDULE Backup schemes	HOW LONG TO KEEP
	Google Drive files	Direct selection (p. 258) Policy rules (p. 258)			
	Shared drive files	Direct selection (p. 261) Policy rules (p. 261)			

* Backup to NFS shares is not available in Windows.

** Secure Zone cannot be created on a Mac.

*** The option Always incremental (Single-file) is available only if the primary backup destination is Cloud.

**** The **By total size of backups** retention rule is not available with the **Always incremental (single-file)** backup scheme or when backing up to the cloud storage.

15.3 Selecting data to back up

15.3.1 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 non-removable disks.

Disks connected via the iSCSI protocol to a physical machine can also be backed up though there are limitations (p. 20) if you use Agent for VMware or Agent for Hyper-V for backing up the iSCSI-connected 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. 153).

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 protection 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 protection 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 macOS

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

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 path as it appears after running the **ls /dev/mapper** command under the root account. For example:

```
[root@localhost ~]# ls /dev/mapper/  
control vg_1-lv1 vg_1-lv2
```

This output shows two logical volumes, **lv1** and **lv2**, that belong to the volume group **vg_1**. To back up these volumes, enter:

```
/dev/mapper/vg_1-lv1  
/dev/mapper/vg_1-lv2
```

Rules for macOS

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

15.3.1.1 What does a disk or volume backup store?

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

With the **sector-by-sector (raw mode)** backup option (p. 171) enabled, a disk backup stores all the disk sectors. The sector-by-sector backup can be used for backing up disks with unrecognized or unsupported file systems and other proprietary data formats.

Windows

A volume backup stores all files and folders of the selected volume independent of their attributes (including hidden and system files), the boot record, the file allocation table (FAT) if it exists, the root and the zero track of the hard disk with the master boot record (MBR).

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

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

- The swap file (pagefile.sys) and the file that keeps the RAM content when the machine goes into hibernation (hiberfil.sys). After recovery, the files will be re-created in the appropriate place with the zero size.
- If the backup is performed under the operating system (as opposed to bootable media or backing up virtual machines at a hypervisor level):
 - Windows shadow storage. The path to it is determined in the registry value **VSS Default Provider** which can be found in the registry key **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup**. This means that in operating systems starting with Windows Vista, Windows Restore Points are not backed up.
 - If the **Volume Shadow Copy Service (VSS)** backup option (p. 172) is enabled, files and folders that are specified in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot** registry key.

Linux

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

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

Mac

A disk or volume backup stores all files and directories of the selected disk or volume, plus a description of the volume layout.

The following items are excluded:

- System metadata, such as the file system journal and Spotlight index
- The Trash
- Time machine backups

Physically, disks and volumes on a Mac are backed up at a file level. Bare metal recovery from disk and volume backups is possible, but the sector-by-sector backup mode is not available.

15.3.2 Selecting files/folders

File-level backup is available for physical machines and virtual machines backed up by an agent installed in the guest system. Files and folders located on disks connected via the iSCSI protocol to a

physical machine can also be backed up though there are limitations (p. 20) if you use Agent for VMware or Agent for Hyper-V for backing up data on the iSCSI-connected disks.

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. 153).

Direct selection

1. In **What to back up**, select **Files/folders**.
2. Specify **Items to back up**.
3. In **Select items for backup**, select **Directly**.
4. For each of the machines included in the protection 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.
Backing up a folder with anonymous access is not supported.
 - d. Select the required files/folders.
 - e. Click **Done**.

Using policy rules

1. In **What to back up**, select **Files/folders**.
2. Specify **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 protection 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.
- Templates:
 - **[All Profiles Folder]** selects **/home**. This is the folder where all user profiles are located by default.

Selection rules for macOS

- Full path to a file or directory.
- Templates:
 - **[All Profiles Folder]** selects **/Users**. This is the folder where all user profiles are located by default.

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**.

15.3.3 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

15.3.4 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.

Limitations

- ESXi configuration backup is not supported for VMware vSphere 6.7.
- An ESXi configuration cannot be backed up to the cloud storage.

To select an ESXi configuration

1. Click **Devices > All devices**, and then select the ESXi hosts that you want to back up.
2. Click **Protect**.
3. In **What to back up**, select **ESXi configuration**.
4. 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.

15.4 Continuous data protection (CDP)

Backups are usually performed with the regular but quite long time intervals due to performance reasons. If the system is suddenly damaged, the data changes between the last backup and the system failure will be lost.

The **Continuous data protection** functionality allows you to back up changes of the selected data between the scheduled backups on the continuous basis:

- By tracking changes in the specified files/folders
- By tracking changes of the files modified by the specified applications

You can select particular files for continuous data protection from the data selected for a backup. The system will back up every change of these files. You can recover these files to the last change time.

Currently, the **Continuous data protection** functionality is supported for the following operating systems:

- Windows 7 and later
- Windows Server 2008 R2 and later

The supported file system: NTFS only, local folders only (shared folders are not supported).

The **Continuous data protection** option is not compatible with the **Application backup** option.


How it works

Let's call the backup that is created on continuous basis the CDP backup. For the CDP backup to be created, a full backup or incremental backup have to be created preliminarily.

When you first run the protection plan with the backup module and **Continuous data protection** enabled, a full backup is created first. Right after that the CDP backup for the selected or changed files/folders will be created. The CDP backup always contains data selected by you in the latest state. When you make changes to the selected files/folders, no new CDP backup is created, all changes are recorded to the same CDP backup.

When the time comes for a scheduled incremental backup, the CDP backup is dropped, and a new CDP backup is created after the incremental backup is done.

Thus, the CDP backup always stays as the latest backup in the backup chain having the latest actual state of the protected files/folders.



If you already have a protection plan with the backup module enabled and you decided to enable **Continuous data protection**, then the CDP backup will be created right after enabling the option as the backup chain already has full backups.

Supported data sources and destinations for continuous data protection

For continuous data protection proper work, you need to specify the following items for the following data sources:

What to back up	Items to back up
Entire machine	Either files/folders or applications must be specified
Disks/volumes	Disks/volumes and either files/folders or applications must be specified
Files/folders	Files/folders must be specified Applications can be specified (not mandatory)

The following backup destinations are supported for continuous data protection:

- Local folder
- Network folder
- Location defined by a script

- Cloud storage
- Acronis Cyber Infrastructure

To protect the devices with continuous data protection

1. In the service console, create a protection plan (p. 95) with the **Backup** module enabled.
2. Enable the **Continuous data protection (CDP)** option.
3. Specify **Items to protect continuously**:

- **Applications** (any file modified by the selected applications will be backed up). We recommend to use this option to protect your Office documents with the CDP backup.

X
Items to protect continuously

Choose files for continuous protection out of the data selected for backup. The software will back up every change of these files. You will be able to revert these files to the last change time.

Applications

Files/folders

Every file modified by the selected applications will be backed-up

Predefined application categories

☒ Office documents

▼

☒ Engineering

▼

☒ Imaging and video

▼

Other applications

To add more applications, specify their paths in the format: C:\Program Files\Microsoft Office\Office16\WINWORD.EXE or *:\Program Files (x86)\Microsoft Office\Office16\WINWORD.EXE

Add applications

OK

Cancel

You can select the applications from the predefined categories or specify other applications by defining the path to the application executable file. Use one of the following formats:

C:\Program Files\Microsoft Office\Office16\WINWORD.EXE

OR

121

Copyright © Acronis International GmbH, 2003-2020

*:\Program Files (x86)\Microsoft Office\Office16\WINWORD.EXE

- **Files/folders** (any file modified in the specified location(s) will be backed up). We recommend to use this option to protect those files and folders that are constantly changing.

×

Items to protect continuously

Choose files for continuous protection out of the data selected for backup. The software will back up every change of these files. You will be able to revert these files to the last change time.

Applications

Files/folders

Every change of the selected files, and of files in the selected folders, will be backed up.

Machine to browse from: NIKITATIKHOB524

⊕ Select files and folders

Add files/folders

OK

Cancel

Machine to browse from – specify the machine whose files/folders you want to select for continuous data protection.

Click **Select files and folders** to select files/folders on the specified machine.

Important If you manually specify a whole folder whose files will be continuously backed up, use the mask, for example:

Correct path: D:\Data*

Incorrect path: D:\Data\

In the text field, you can also specify rules for selecting files/folders that will be backed up. For more details how to define rules, refer to "Selecting files/folders". When ready, click **Done**.


4. Click **Create**.

As a result, the protection plan with continuous data protection enabled will be assigned to the selected machine. After the first regular backup, the backups with the latest copy of the protected by CDP data will be created on the continuous basis. Both, the data defined via Applications and Files/folders, will be backed up.

Continuously backed-up data are retained according to the retention policy defined for the backup module.

How to distinguish backups that are protected on continuous basis

The backups that are backed up on continuous basis have the CDP prefix.



How to recover your entire machine to the latest state

If you want to be able to recover an entire machine to the latest state, you can use the **Continuous data protection (CDP)** option in the backup module of a protection plan.

You can recover either an entire machine or files/folders from a CDP backup. In first case, you will get an entire machine in the latest state, in the second case – files/folders in the latest state.

15.5 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/DFS.

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

- For SMB/CIFS shares: \\<host name>\<path>\ or smb://<host name>/<path>/
- For DFS shares: \\<full DNS domain name>\<DFS root>\<path>
For example, \\example.company.com\shared\files

Then, click the arrow button. If prompted, specify the user name and password for the shared folder. You can change these credentials at any time by clicking the key icon next to the folder name.

Backing up to a folder with anonymous access is not supported.

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

Verify that the nfs-utils package is installed on the Linux server where the Agent for Linux is installed.

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.

Note 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".

Advanced storage option

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

- **Defined by a script** (available for machines running Windows)

You can store each machine's backups in a folder defined by a script. The software supports scripts written in JScript, VBScript, or Python 3.5. When deploying the protection plan, the software runs the script on each machine. The script output for each machine should be a local or network folder path. If a folder does not exist, it will be created (limitation: scripts written in Python cannot create folders on network shares). On the **Backup storage** tab, each folder is shown as a separate backup location.

In **Script type**, select the script type (**JScript**, **VBScript**, or **Python**), and then import, or copy and paste the script. For network folders, specify the access credentials with the read/write permissions.

Example. The following JScript script outputs the backup location for a machine in the format \\bkpsrv\<machine name>:

```
WScript.echo("\\\\bkpsrv\\" +  
WScript.CreateObject("WScript.Network").ComputerName);
```

As a result, the backups of each machine will be saved in a folder of the same name on the server **bkpsrv**.

15.5.1 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. 399). When you change the destination to Secure Zone in a protection plan that has the **Always incremental (Single-file)** backup scheme, the scheme is changed to **Weekly full, daily incremental**.

How creating Secure Zone transforms the disk

- Secure Zone is always created at the end of the hard disk.
- 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 of the disk.
- When all unallocated space is collected but it is still not enough, the software will take free space from the volumes you select, proportionally reducing the volumes' size.
- However, there should be free space on a volume, so that the operating system and applications can operate; for example, create temporary files. The software 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 software 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.

Important *Moving or resizing the volume from which the system is booted requires a reboot.*

How to create Secure Zone

1. Select the machine that you want to create Secure Zone on.
2. Click **Details > Create Secure Zone**.
3. Under **Secure Zone disk**, click **Select**, and then select a hard disk (if several) on which to create the zone.

The software calculates the maximum possible size of Secure Zone.

4. Enter the Secure Zone size or drag the slider to select any size between the minimum and the maximum ones.

The minimum size is approximately 50 MB, depending on the geometry of the hard disk. The maximum size is equal to the disk's unallocated space plus the total free space on all of the disk's volumes.

5. If all unallocated space is not enough for the size you specified, the software will take free space from the existing volumes. By default, all volumes are selected. If you want to exclude some volumes, click **Select volumes**. Otherwise, skip this step.

Create Secure Zone

Secure Zone disk

Disk 1, 60.0 GB

Maximum possible size of Secure Zone: 35.9 GB

Secure Zone size:

20 GB

There is not enough unallocated space. Free space will be taken from all volumes where it is present.

[Select volumes](#)

Password protection

☐ Off

6. [Optional] Enable the **Password protection** switch and specify a password.
The password will be required to access the backups located in Secure Zone. Backing up to Secure Zone does not require a password, unless the backup is performed under bootable media.
7. Click **Create**.
The software displays the expected partition layout. Click **OK**.
8. Wait while the software creates Secure Zone.

You can now choose Secure Zone in **Where to back up** when creating a protection plan.

How to delete Secure Zone

1. Select a machine with Secure Zone.
2. Click **Details**.
3. Click the gear icon next to **Secure Zone**, and then click **Delete**.
4. [Optional] Specify the volumes to which the space freed from the zone will be added. By default, all volumes are selected.

The space will be distributed equally among the selected volumes. If you do not select any volumes, the freed space will become unallocated.

Resizing the volume from which the system is booted requires a reboot.
5. Click **Delete**.

As a result, Secure Zone will be deleted along with all backups stored in it.

15.6 Schedule

The schedule employs the time settings (including the time zone) of the operating system where the agent is installed. The time zone of Agent for VMware (Virtual Appliance) can be configured in the agent's interface (p. 65).

For example, if a protection plan is scheduled to run at 21:00 and applied to several machines located in different time zones, the backup will start on each machine at 21:00 local time.

Backup schemes

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

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

- **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 single-file backup format (p. 399).

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

This scheme is highly recommended if the backup location is cloud storage. Other backup schemes may include multiple full backups that consume much time and network traffic.

This scheme is not available when backing up to Secure Zone.

- **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**).

- **Monthly full, Weekly differential, Daily incremental (GFS)**

By default, incremental backups are performed on a daily basis, Monday to Friday; differential backups are performed every Saturday; full backups are performed on the first day of each month. You can modify these schedules and the time to run the backup.

This backup scheme is displayed as a **Custom** scheme on the protection plan panel.

- **Custom**

Specify schedules for full, differential, and incremental backups.

Differential backup is not available when backing up SQL data, Exchange data, or system state.

With any backup scheme, you can schedule the backup to run by events, instead of by time. To do this, select the event type in the schedule selector. For more information, refer to "Schedule by events".

Additional scheduling options

With any destination, you can do the following:


- Specify the backup start conditions, so that a scheduled backup is performed only if the conditions are met. For more information, refer to "Start conditions".
- 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.

In the protection plan in the Backup module settings, go to **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 protection plan is applied to the machine and remains the same until you edit the protection plan and change the maximum delay value.

Note This option is enabled by default, with the maximum delay set to 30 minutes.

- Click **Show more** to access the following options:
 - **If the machine is turned off, run missed tasks at the machine startup** (disabled by default)
 - **Prevent the sleep or hibernate mode during backup** (enabled by default)
This option is effective only for machines running Windows.
 - **Wake up from the sleep or hibernate mode to start a scheduled backup** (disabled by default)

This option is effective only for machines running Windows whose power plan has the **Allow wake timers** setting enabled.



This option is not effective when the machine is powered off, i.e. the option does not employ the Wake-on-LAN functionality.

15.6.1 Schedule by events

When setting up a schedule for the Backup module of the protection plan, you can select the event type in the schedule selector. The backup will be launched as soon as the event occurs.

You can choose one of the following events:

- **Upon time since last backup**
This is the time since the completion of the last successful backup within the same protection plan. You can specify the length of time.
- **When a user logs on to the system**
By default, logging on of any user will initiate a backup. You can change any user to a specific user account.
- **When a user logs off the system**
By default, logging off of any user will initiate a backup. You can change any user to a specific user account.

Note The backup will not run at a system shutdown because shutting down is not the same as logging off.

- **On the system startup**
- **On the system shutdown**

- **On Windows Event Log event**

You must specify the event properties.

The table below lists the events available for various data under Windows, Linux, and macOS.

WHAT TO BACK UP	Upon time since last backup	When a user logs on to the system	When a user logs off the system	On the system startup	On the system shutdown	On Windows Event Log event
Disks/volumes or files (physical machines)	Windows, Linux, macOS	Windows	Windows	Windows, Linux, macOS	Windows	Windows
Disks/volumes (virtual machines)	Windows, Linux	–	–	–	–	–
ESXi configuration	Windows, Linux	–	–	–	–	–
Office 365 mailboxes	Windows	–	–	–	–	Windows
Exchange databases and mailboxes	Windows	–	–	–	–	Windows
SQL databases	Windows	–	–	–	–	Windows

15.6.1.1 On Windows Event Log event

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

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

To browse the events and view the event properties, use the **Event Viewer** snap-in available in the **Computer Management** console. To be able to open the **Security** log, you must be a member of the **Administrators** group.

Event properties

Log name

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

Event source

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

Any event source that contains the specified string will trigger the scheduled backup. This option is not case sensitive. Thus, if you specify the string **service**, both **Service Control Manager** and **Time-Service** event sources will trigger a backup.

Event type

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

Event ID

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

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

Example: "Bad block" emergency backup

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

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

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

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

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

15.6.2 Start conditions

These settings add more flexibility to the scheduler, enabling it to execute a backup with respect to certain conditions. With multiple conditions, all of them must be met simultaneously to enable a backup to start. Start conditions are not effective when a backup is started manually.

To access these settings, click **Show more** when setting up a schedule for a protection plan.

The scheduler behavior, in case the condition (or any of multiple conditions) is not met, is defined by the Backup start conditions (p. 172) backup option. To handle the situation when the conditions are not met for too long and further delaying the backup is becoming risky, you can set the time interval after which the backup will run irrespective of the condition.

The table below lists the start conditions available for various data under Windows, Linux, and macOS.

WHAT TO BACK UP	Disks/volumes or files (physical machines)	Disks/volumes (virtual machines)	ESXi configuration	Office 365 mailboxes	Exchange databases and mailboxes	SQL databases
User is idle (p. 132)	Windows	–	–	–	–	–
The backup location's host is available (p. 132)	Windows, Linux, macOS	Windows, Linux	Windows, Linux	Windows	Windows	Windows
Users logged off (p. 133)	Windows	–	–	–	–	–

Fits the time interval (p. 133)	Windows, Linux, macOS	Windows, Linux	–	–	–	–
Save battery power (p. 134)	Windows	–	–	–	–	–
Do not start when on metered connection (p. 134)	Windows	–	–	–	–	–
Do not start when connected to the following Wi-Fi networks (p. 135)	Windows	–	–	–	–	–
Check device IP address (p. 136)	Windows	–	–	–	–	–

15.6.2.1 User is idle

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

Example

Run the backup on the machine every day at 21:00, preferably when the user is idle. If the user is still active by 23:00, run the backup anyway.

- Schedule: Daily, Run every day. Start at: **21:00**.
- Condition: **User is idle**.
- Backup start conditions: **Wait until the conditions are met, Start the backup anyway after 2 hour(s)**.

As a result,

- (1) If the user becomes idle before 21:00, the backup will start at 21:00.
- (2) If the user becomes idle between 21:00 and 23:00, the backup will start immediately after the user becomes idle.
- (3) If the user is still active at 23:00, the backup will start at 23:00.

15.6.2.2 The backup location's host is available

"The backup location's host is available" means that the machine hosting the destination for storing backups is available over the network.

This condition is effective for network folders, the cloud storage, and locations managed by a storage node.

This condition does not cover the availability of the location itself — only the host availability. For example, if the host is available, but the network folder on this host is not shared or the credentials for the folder are no longer valid, the condition is still considered met.

Example

Data is backed up to a network folder every workday at 21:00. If the machine that hosts the folder is not available at that moment (for instance, due to maintenance work), you want to skip the backup and wait for the scheduled start on the next workday.

- Schedule: Daily, Run Monday to Friday. Start at: **21:00**.
- Condition: **The backup location's host is available**.
- Backup start conditions: **Skip the scheduled backup**.

As a result:

- (1) If 21:00 comes and the host is available, the backup will start immediately.
- (2) If 21:00 comes but the host is unavailable, the backup will start on the next workday if the host is available.
- (3) If the host is never available on workdays at 21:00, the backup will never start.

15.6.2.3 Users logged off

Enables you to put a backup on hold until all users log off from Windows.

Example

Run the backup at 20:00 every Friday, preferably when all users are logged off. If one of the users is still logged on at 23:00, run the backup anyway.

- Schedule: Weekly, on Fridays. Start at: **20:00**.
- Condition: **Users logged off**.
- Backup start conditions: **Wait until the conditions are met, Start the backup anyway after 3 hour(s)**.

As a result:

- (1) If all users are logged off at 20:00, the backup will start at 20:00.
- (2) If the last user logs off between 20:00 and 23:00, the backup will start immediately after the user logs off.
- (3) If any user is still logged on at 23:00, the backup will start at 23:00.

15.6.2.4 Fits the time interval

Restricts a backup start time to a specified interval.

Example

A company uses different locations on the same network-attached storage for backing up users' data and servers. The workday starts at 08:00 and ends at 17:00. Users' data should be backed up as soon as the users log off, but not earlier than 16:30. Every day at 23:00 the company's servers are backed up. So, all the users' data should preferably be backed up before this time, in order to free network bandwidth. It is assumed that backing up user's data takes no more than one hour, so the latest backup start time is 22:00. If a user is still logged on within the specified time interval, or logs off at any other time – do not back up the users' data, i.e., skip backup execution.

- Event: **When a user logs off the system**. Specify the user account: **Any user**.

- Condition: **Fits the time interval** from **16:30** to **22:00**.
- Backup start conditions: **Skip the scheduled backup**.

As a result:

(1) if the user logs off between 16:30 and 22:00, the backup will start immediately following the logging off.

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

15.6.2.5 Save battery power

Prevents a backup if the device (a laptop or a tablet) is not connected to a power source. Depending on the value of the Backup start conditions (p. 172) backup option, the skipped backup will or will not be started after the device is connected to a power source. The following options are available:

- **Do not start when on battery**
A backup will start only if the device is connected to a power source.
- **Start when on battery if the battery level is higher than**
A backup will start if the device is connected to a power source or if the battery level is higher than the specified value.

Example

Data is backed up every workday at 21:00. If the device is not connected to a power source (for instance, the user is attending a late meeting), you want to skip the backup to save the battery power and wait until the user connects the device to a power source.

- Schedule: Daily, Run Monday to Friday. Start at: 21:00.
- Condition: **Save battery power, Do not start when on battery**.
- Backup start conditions: **Wait until the conditions are met**.

As a result:

(1) If 21:00 comes and the device is connected to a power source, the backup will start immediately.

(2) If 21:00 comes and the device is running on battery power, the backup will start as soon as the device is connected to a power source.

15.6.2.6 Do not start when on metered connection

Prevents a backup (including a backup to a local disk) if the device is connected to the Internet by using a connection that is set as metered in Windows. For more information about metered connections in Windows, refer to <https://support.microsoft.com/en-us/help/17452/windows-metered-internet-connections-faq>.

As an additional measure to prevent backups over mobile hotspots, when you enable the **Do not start when on metered connection** condition, the condition **Do not start when connected to the following Wi-Fi networks** is enabled automatically. The following network names are specified by default: "android", "phone", "mobile", and "modem". You can delete these names from the list by clicking on the X sign.

Example

Data is backed up every workday at 21:00. If the device is connected to the Internet by using a metered connection (for instance, the user is on a business trip), you want to skip the backup to save the network traffic and wait for the scheduled start on the next workday.

- Schedule: Daily, Run Monday to Friday. Start at: 21:00.
- Condition: **Do not start when on metered connection.**
- Backup start conditions: **Skip the scheduled backup.**

As a result:

(1) If 21:00 comes and the device is not connected to the Internet by using a metered connection, the backup will start immediately.

(2) If 21:00 comes and the device is connected to the Internet by using a metered connection, the backup will start on the next workday.

(3) If the device is always connected to the Internet by using a metered connection on workdays at 21:00, the backup will never start.

15.6.2.7 Do not start when connected to the following Wi-Fi networks

Prevents a backup (including a backup to a local disk) if the device is connected to any of the specified wireless networks. You can specify the Wi-Fi network names, also known as service set identifiers (SSID).

The restriction applies to all networks that contain the specified name as a substring in their name, case-insensitive. For example, if you specify "phone" as the network name, the backup will not start when the device is connected to any of the following networks: "John's iPhone", "phone_wifi", or "my_PHONE_wifi".

This condition is useful to prevent backups when the device is connected to the Internet by using a mobile phone hotspot.

As an additional measure to prevent backups over mobile hotspots, the **Do not start when connected to the following Wi-Fi** condition is enabled automatically when you enable the **Do not start when on metered connection** condition. The following network names are specified by default: "android", "phone", "mobile", and "modem". You can delete these names from the list by clicking on the X sign.

Example

Data is backed up every workday at 21:00. If the device is connected to the Internet by using a mobile hotspot (for example, a laptop is connected in the tethering mode), you want to skip the backup and wait for the scheduled start on the next workday.

- Schedule: Daily, Run Monday to Friday. Start at: 21:00.
- Condition: **Do not start when connected to the following networks, Network name: <SSID of the hotspot network>.**
- Backup start conditions: **Skip the scheduled backup.**

As a result:

(1) If 21:00 comes and the machine is not connected to the specified network, the backup will start immediately.

(2) If 21:00 comes and the machine is connected to the specified network, the backup will start on the next workday.

(3) If the machine is always connected to the specified network on workdays at 21:00, the backup will never start.

15.6.2.8 Check device IP address

Prevents a backup (including a backup to a local disk) if any of the device IP addresses are within or outside of the specified IP address range. The following options are available:

- **Start if outside IP range**
- **Start if within IP range**

With either option, you can specify several ranges. Only IPv4 addresses are supported.

This condition is useful in the event of a user being overseas, to avoid large data transit charges. Also, it helps to prevent backups over a Virtual Private Network (VPN) connection.

Example

Data is backed up every workday at 21:00. If the device is connected to the corporate network by using a VPN tunnel (for instance, the user is working from home), you want to skip the backup and wait until the user brings the device to the office.

- Schedule: Daily, Run Monday to Friday. Start at: 21:00.
- Condition: **Check device IP address, Start if outside IP range, From:** <beginning of the VPN IP address range>, **To:** <end of the VPN IP address range>.
- Backup start conditions: **Wait until the conditions are met.**

As a result:

(1) If 21:00 comes and the machine IP address is not in the specified range, the backup will start immediately.

(2) If 21:00 comes and the machine IP address is in the specified range, the backup will start as soon as the device obtains a non-VPN IP address.

(3) If the machine IP address is always in the specified range on workdays at 21:00, the backup will never start.

15.7 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 protection plan. By default, the retention rules are specified for each backup set (p. 397) 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.
 - **By total size of backups**
Specify the maximum total size of backups to keep.

This setting is not available with the **Always incremental (single-file)** backup scheme or when backing up to the cloud storage.

- **Keep backups indefinitely**

3. Select when to start the cleanup:

- **After backup** (default)

The retention rules will be applied after a new backup is created.

- **Before backup**

The retention rules will be applied before a new backup is created.

This setting is not available when backing up Microsoft SQL Server clusters or Microsoft Exchange Server clusters.

What else you need to know

- If, according to the backup scheme and backup format, each backup is stored as a separate file, this file cannot be deleted until the lifetime of all its dependent (incremental and differential) backups expires. This requires extra space for storing backups whose deletion is postponed. Also, the backup age, number, or size of backups may exceed the values you specify.
This behavior can be changed by using the "Backup consolidation" (p. 144) backup option.
- Retention rules are a part of a protection plan. They stop working for a machine's backups as soon as the protection plan is revoked from the machine, or deleted, or the machine itself is deleted from the Cyber Protection service. If you no longer need the backups created by the plan, delete them as described in "Deleting backups" (p. 201).

15.8 Replication

If you enable backup replication, each backup will be copied to another 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. If backup replication is interrupted in the middle of a process, then on the next replication start the already replicated data will not be replicated again which allows reducing time loss.

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

- Secure Zone

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

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

To enable replication of backups

1. On the protection plan panel, click **Add location**.
The **Add location** control is shown only if replication is supported *from* the last selected location.
2. Specify the location where the backups will be replicated.
3. [Optional] In **How long to keep**, change the retention rules for the chosen location, as described in "Retention rules" (p. 136).
4. [Optional] Click the gear icon > **Performance and backup window**, and then set the backup window for the chosen location, as described in "Performance and backup window" (p. 164).
These settings will define the replication performance.
5. [Optional] Repeat steps 1-4 for all locations where you want to replicate the backups. Up to five consecutive locations are supported, including the primary one.

15.9 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 protection plan

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

To specify the encryption settings in a protection plan

1. On the protection plan panel in the Backup module settings, 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 protection plan encryption settings, save the encryption settings on each machine individually. The backups will be encrypted using the AES algorithm with a 256-bit key.

Saving the encryption settings on a machine affects the protection plans in the following way:

- **Protection plans that are already applied to the machine.** If the encryption settings in a protection plan are different, the backups will fail.
- **Protection plans that will be applied to the machine later.** The encryption settings saved on a machine will override the encryption settings in a protection plan. Any backup will be encrypted, even if encryption is disabled in the Backup module settings.

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 type of load balancing among them.

After the encryption settings are saved, they can be changed or reset as described below.

Important *If a protection plan that runs on this machine has already created backups, changing the encryption settings will cause this plan to fail. To continue backing up, create a new plan.*

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 protection agent installation path. By default, it is `%ProgramFiles%\BackupClient`.
 - In Linux: `/usr/sbin/acropsh -m manage_creds --set-password <encryption_password>`

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 protection agent installation path. By default, it is `%ProgramFiles%\BackupClient`.
 - In Linux: `/usr/sbin/acropsh -m manage_creds --reset`

To change the encryption settings by using Cyber Protection Monitor

1. Log on as an administrator in Windows or macOS.
2. Click the Cyber Protection Monitor icon in the notification area (in Windows) or the menu bar (in macOS).
3. Click the gear icon.
4. Click **Encryption**.
5. Do one of the following:
 - Select **Set a specific password for this machine**. Specify and confirm the encryption password.
 - Select **Use encryption settings specified in the protection plan**.
6. Click **OK**.

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.

15.10 Notarization

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

Notarization enables you to prove that a file is authentic and unchanged since it was backed up. We recommend that you enable notarization when backing up your legal document files or other files that require proved authenticity.

Notarization is available only for file-level backups. Files that have a digital signature are skipped, because they do not need to be notarized.

Notarization is *not* available:

- If the backup format is set to **Version 11**
- If the backup destination is Secure Zone

How to use notarization

To enable notarization of all files selected for backup (except for the files that have a digital signature), enable the **Notarization** switch when creating a protection plan.

When configuring recovery, the notarized files will be marked with a special icon, and you can verify the file authenticity (p. 189).

How it works

During a backup, the agent calculates the hash codes of the backed-up files, builds a hash tree (based on the folder structure), saves the tree in the backup, and then sends the hash tree root to the notary service. The notary service saves the hash tree root in the Ethereum blockchain database to ensure that this value does not change.

When verifying the file authenticity, the agent calculates the hash of the file, and then compares it with the hash that is stored in the hash tree inside the backup. If these hashes do not match, the file is considered not authentic. Otherwise, the file authenticity is guaranteed by the hash tree.

To verify that the hash tree itself was not compromised, the agent sends the hash tree root to the notary service. The notary service compares it with the one stored in the blockchain database. If the hashes match, the selected file is guaranteed to be authentic. Otherwise, the software displays a message that the file is not authentic.

15.11 Starting a backup manually

1. Select a machine that has at least one applied protection plan.
2. Click **Protect**.
3. If more than one protection plans are applied, select the protection plan.
4. Do one of the following:
 - Click **Run now**. An incremental backup will be created.
 - If the backup scheme includes several backup methods, you can choose the method to use. Click the arrow on the **Run now** button, and then select **Full**, **Incremental**, or **Differential**.

The first backup created by a protection plan is always full.

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

15.12 Default backup options

The default values of backup options (p. 141) exist at the company, unit, and user level. When a unit or a user account is created within a company or within a unit, it inherits the default values set for the company or for the unit.

Company administrators, unit administrators, and every user without the administrator rights can change a default option value against the pre-defined one. The new value will be used by default in all protection plans created at the respective level after the change takes place.

When creating a protection plan, a user can override a default value with a custom value that will be specific for this plan only.

To change a default option value

1. Do one of the following:
 - To change the default value for the company, sign in to the service console as a company administrator.
 - To change the default value for a unit, sign in to the service console as an administrator of the unit.
 - To change the default value for yourself, sign in to the service console by using an account without the administrator rights.
2. Click **Settings > System settings**.
3. Expand the **Default backup options** section.
4. Select the option, and then make the necessary changes.
5. Click **Save**.

15.13 Backup options

To modify the backup options, click **Change** next to **Backup options** in the Backup module of the protection plan.

Availability of the backup options

The set of available backup options depends on:

- The environment the agent operates in (Windows, Linux, macOS).
- 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	macOS	Windows	Linux	macOS	ESXi	Hyper-V	Virtuozzo	Windows
Alerts (p. 144)	+	+	+	+	+	+	+	+	+	+
Backup consolidation (p. 144)	+	+	+	+	+	+	+	+	+	-
Backup file name (p. 145)	+	+	+	+	+	+	+	+	+	+
Backup format (p. 148)	+	+	+	+	+	+	+	+	+	+
Backup validation (p. 149)	+	+	+	+	+	+	+	+	+	+
Changed block tracking (CBT) (p. 149)	+	-	-	-	-	-	+	+	-	-
Cluster backup mode (p. 150)	-	-	-	-	-	-	-	-	-	+
Compression level (p. 151)	+	+	+	+	+	+	+	+	+	+
Error handling (p. 151)										
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. 152)	+	+	+	-	-	-	-	-	-	-
File-level backup snapshot (p. 154)	-	-	-	+	+	+	-	-	-	-
File filters (p. 153)	+	+	+	+	+	+	+	+	+	-

	Disk-level backup			File-level backup			Virtual machines			SQL and Exchange
	Windows	Linux	macOS	Windows	Linux	macOS	ESXi	Hyper-V	Virtuozzo	Windows
Forensic data (p. 154)	+	-	-	-	-	-	-	-	-	-
Log truncation (p. 162)	-	-	-	-	-	-	+	+	-	SQL only
LVM snapshotting (p. 162)	-	+	-	-	-	-	-	-	-	-
Mount points (p. 163)	-	-	-	+	-	-	-	-	-	-
Multi-volume snapshot (p. 163)	+	+	-	+	+	-	-	-	-	-
Performance and backup window (p. 164)	+	+	+	+	+	+	+	+	+	+
Physical Data Shipping (p. 166)	+	+	+	+	+	+	+	+	+	-
Pre/Post commands (p. 167)	+	+	+	+	+	+	+	+	+	+
Pre/Post data capture commands (p. 169)	+	+	+	+	+	+	-	-	-	+
Scheduling (p. 170)										
Distribute start times within a time window	+	+	+	+	+	+	+	+	+	+
Limit the number of simultaneously running backups	-	-	-	-	-	-	+	+	+	-
Sector-by-sector backup (p. 171)	+	+	-	-	-	-	+	+	+	-
Splitting (p. 171)	+	+	+	+	+	+	+	+	+	+
Task failure handling (p. 172)	+	+	+	+	+	+	+	+	+	+

	Disk-level backup			File-level backup			Virtual machines			SQL and Exchange
	Windows	Linux	macOS	Windows	Linux	macOS	ESXi	Hyper-V	Virtuozzo	Windows
Task start conditions (p. 172)	+	+	-	+	+	-	+	+	+	+
Volume Shadow Copy Service (VSS) (p. 172)	+	-	-	+	-	-	-	+	-	+
Volume Shadow Copy Service (VSS) for virtual machines (p. 173)	-	-	-	-	-	-	+	+	-	-
Weekly backup (p. 174)	+	+	+	+	+	+	+	+	+	+
Windows event log (p. 174)	+	-	-	+	-	-	+	+	-	+

15.13.1 Alerts

No successful backups for a specified number of consecutive days

The preset is: **Disabled**.

This option determines whether to generate an alert if no successful backups were performed by the protection plan for a specified period of time. In addition to failed backups, the software counts backups that did not run on schedule (missed backups).

The alerts are generated on a per-machine basis and are displayed on the **Alerts** tab.

You can specify the number of consecutive days without backups after which the alert is generated.

15.13.2 Backup consolidation

This option defines whether to consolidate backups during cleanup or to delete entire backup chains.

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.

This option is *not* effective if any of the following is true:

- The backup destination is the cloud storage.
- The backup scheme is set to **Always incremental (single-file)**.
- The backup format (p. 148) is set to **Version 12**.

Backups stored in the cloud storage, as well as single-file backups (both version 11 and 12 formats), are always consolidated because their inner structure makes for fast and easy consolidation.

However, if version 12 format is used, and multiple backup chains are present (every chain being stored in a separate .tibx file), consolidation works only within the last chain. Any other chain is deleted as a whole, except for the first one, which is shrunk to the minimum size to keep the meta information (~12 KB). This meta information is required to ensure the data consistency during simultaneous read and write operations. The backups included in these chains disappear from the GUI as soon as the retention rule is applied, although they physically exist until the entire chain is deleted.

In all other cases, backups whose deletion is postponed are marked with the trash can icon () in the GUI. If you delete such a backup by clicking the X sign, consolidation will be performed.

15.13.3 Backup file name

This option defines the names of the backup files created by the protection plan.

These names can be seen in a file manager when browsing the backup location.

What is a backup file?

Each protection plan creates one or more files in the backup location, depending on which backup scheme and which backup format (p. 148) is used. The following table lists the files that can be created per machine or mailbox.

	Always incremental (single-file)	Other backup schemes
Version 11 backup format	One .tib file and one .xml metadata file	Multiple .tib files and one .xml metadata file
Version 12 backup format	One .tibx file per backup chain (a full or differential backup, and all incremental backups that depend on it). If the size of a file stored in a local or network (SMB) folder exceeds 200 GB, the file is split to 200-GB files by default.	

All files have the same name, with or without the addition of a timestamp or a sequence number. You can define this name (referred to as the backup file name) when creating or editing a protection plan.

Note Timestamp is added to the backup file name only in the **Version 11** backup format.

After you change a backup file name, the next backup will be a full backup, unless you specify a file name of an existing backup of the same machine. If the latter is the case, a full, incremental, or differential backup will be created according to the protection plan schedule.

Note that it is possible to set backup file names for locations that cannot be browsed by a file manager (such as the cloud storage). This makes sense if you want to see the custom names on the **Backup storage** tab.

Where can I see backup file names?

Select the **Backup storage** tab, and then select the group of backups.

- The default backup file name is shown on the **Details** panel.
- If you set a non-default backup file name, it will be shown directly on the **Backup storage** tab, in the **Name** column.

Limitations for backup file names

- A backup file name cannot end with a digit.
In the default backup file name, to prevent the name from ending with a digit, the letter "A" is appended. When creating a custom name, always make sure that it does not end with a digit. When using variables, the name must not end with a variable, because a variable might end with a digit.
- A backup file name cannot contain the following symbols: `()&?*${<>":\|/ #`, line endings (`\n`), and tabs (`\t`).

Default backup file name

The default backup file name for backups of entire physical and virtual machines, disks/volumes, files/folders, Microsoft SQL Server databases, Microsoft Exchange Server databases, and ESXi configuration is **[Machine Name]-[Plan ID]-[Unique ID]A**.


The default name for Exchange mailbox backups and Office 365 mailbox backups created by a local Agent for Office 365 is **[Mailbox ID]_mailbox_[Plan ID]A**.

The default name for cloud application backups created by cloud agents is **[Resource Name]_[Resource Type]_[Resource ID]_[Plan ID]A**.


The default name consists of the following variables:

- **[Machine Name]** This variable is replaced with the name of the machine (the same name that is shown in the service console).
- **[Plan ID]**, **[Plan Id]** These variables are replaced with the unique identifier of the protection plan. This value does not change if the plan is renamed.
- **[Unique ID]** This variable is replaced with the unique identifier of the selected machine. This value does not change if the machine is renamed.
- **[Mailbox ID]** This variable is replaced with the mailbox user's principal name (UPN).
- **[Resource Name]** This variable is replaced with the cloud data source name, such as the user's principal name (UPN), SharePoint site URL, or Shared drive name.
- **[Resource Type]** This variable is replaced with the cloud data source type, such as **mailbox**, **0365Mailbox**, **0365PublicFolder**, **OneDrive**, **SharePoint**, **GDrive**.
- **[Resource ID]** This variable is replaced with the unique identifier of the cloud data source. This value does not change if the cloud data source is renamed.
- **"A"** is a safeguard letter that is appended to prevent the name from ending with a digit.

The diagram below shows the default backup file name.



The diagram below shows the default backup file name for Office 365 mailbox backups performed by a local agent.



Names without variables

If you change the backup file name to **MyBackup**, the backup files will look like the following examples. Both examples assume daily incremental backups scheduled at 14:40, starting from September 13, 2016.

For the **Version 12** format with the **Always incremental (single-file)** backup scheme:

```
MyBackup.tibx
```

For the **Version 12** format with other backup schemes:

```
MyBackup.tibx
MyBackup-0001.tibx
MyBackup-0002.tibx
...
```

Using variables

Besides the variables that are used by default, you can use the following variables:

- The **[Plan name]** variable, which is replaced with the name of the protection plan.
- The **[Virtualization Server Type]** variable, which is replaced with "vmwesx" if virtual machines are backed up by Agent for VMware or with "mshyperv" if virtual machines are backed up by Agent for Hyper-V.

If multiple machines or mailboxes are selected for backup, the backup file name must contain the **[Machine Name]**, the **[Unique ID]**, the **[Mailbox ID]**, the **[Resource Name]**, or the **[Resource Id]** variable.

Usage examples

▪ View user-friendly file names

You want to easily distinguish backups when browsing the backup location with a file manager.

▪ Continue an existing sequence of backups

Let's assume a protection plan is applied to a single machine, and you have to remove this machine from the service console or to uninstall the agent along with its configuration settings. After the machine is re-added or the agent is reinstalled, you can force the protection plan to continue backing up to the same backup or backup sequence. Just go this option, click **Select**, and select the required backup.

The **Select** button shows the backups in the location selected in the **Where to back up** section of the protection plan panel. It cannot browse anything outside this location.

File name template

[Machine Name]-[Plan ID]-[Unique ID]A **SELECT**

If the file name template is changed, the next backup will be a full backup.

The following variables can be used:

- [Machine Name]
- [Plan ID]
- [Plan name]
- [Unique ID]

Note The **Select** button is only available for protection plans that are created for and applied to a single device.

15.13.4 Backup format

The **Backup format** option defines the format of the backups created by the protection plan. This option is available only for protection plans that already use the **Version 11** backup format. If this is the case, you can change the backup format to **Version 12**. After you switch the backup format to Version 12, the option becomes unavailable.

■ Version 11

The legacy format preserved for backward compatibility.

Note You cannot back up Database Availability Groups(DAG) by using archive format Version 11. Backing up of DAG is supported only in archive format Version 12.

■ Version 12

The backup format that was introduced in Acronis Backup 12 for faster backup and recovery. Each backup chain (a full or differential backup, and all incremental backups that depend on it) is saved to a single .tibx file.

Backup format and backup files

For backup locations that can be browsed with a file manager (such as local or network folders), the backup format determines the number of files and their extension. The following table lists the files that can be created per machine or mailbox.

	Always incremental (single-file)	Other backup schemes
Version 11 backup format	One .tib file and one .xml metadata file	Multiple .tib files and one .xml metadata file
Version 12 backup format	One .tibx file per backup chain (a full or differential backup, and all incremental backups that depend on it). If the size of a file stored in a local or network (SMB) folder exceeds 200 GB, the file is split to 200-GB files by default.	

Changing the backup format to version 12 (.tibx)

If you change the backup format from version 11 (.tib format) to version 12 (.tibx format):

- The next backup will be full.
- In backup locations that can be browsed with a file manager (such as local or network folders), a new .tibx file will be created. The new file will have the name of the original file, appended with the **_v12A** suffix.
- Retention rules and replication will be applied only to the new backups.
- The old backups will not be deleted and will remain available on the **Backup storage** tab. You can delete them manually.
- The old cloud backups will not consume the **Cloud storage** quota.
- The old local backups will consume the **Local backup** quota until you delete them manually.

In-archive deduplication

The backup format of version 12 supports in-archive deduplication that brings the following advantages:

- Reduced backup size in tens of times, with built-in block-level deduplication for any type of data
- Efficient handling of hard links ensures that there are no storage duplicates
- Hash-based chunking

Note *In-archive deduplication is enabled by default for all backups in .tibx format. You do not have to enable it in the backup options, and you cannot disable it.*

15.13.5 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 protection 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. 268) in the ESXi or Hyper-V environment.

Note *Validation is available for cloud storage located in an Acronis data center and provided by Acronis partners.*

15.13.6 Changed block tracking (CBT)

This option is effective for disk-level backups of virtual machines and of physical machines running Windows. It is also effective for backups of Microsoft SQL Server databases and Microsoft Exchange Server databases.

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 or database content are continuously tracked at the block level. When a backup starts, the changes can be immediately saved to the backup.

15.13.7 Cluster backup mode

Note *This functionality is available only in the Advanced edition of the Cyber Protection service.*

These options are effective for database-level backup of Microsoft SQL Server and Microsoft Exchange Server.

These options are effective only if the cluster itself (Microsoft SQL Server Always On Availability Groups (AAG) or Microsoft Exchange Server Database Availability Group (DAG)) is selected for backup, rather than the individual nodes or databases inside of it. If you select individual items inside the cluster, the backup will not be cluster-aware and only the selected copies of the items will be backed up.

Microsoft SQL Server

This option determines the backup mode for SQL Server Always On Availability Groups (AAG). For this option to be effective, Agent for SQL must be installed on all of the AAG nodes. For more information about backing up Always On Availability Groups, refer to "Protecting Always On Availability Groups (AAG)" (p. 206).

The preset is: **Secondary replica if possible.**

You can choose one of the following:

- **Secondary replica if possible**
If all secondary replicas are offline, the primary replica is backed up. Backing up the primary replica may slow down the SQL Server operation, but the data will be backed up in the most recent state.
- **Secondary replica**
If all secondary replicas are offline, the backup will fail. Backing up secondary replicas does not affect the SQL server performance and allows you to extend the backup window. However, passive replicas may contain information that is not up-to-date, because such replicas are often set to be updated asynchronously (lagged).
- **Primary replica**
If the primary replica is offline, the backup will fail. Backing up the primary replica may slow down the SQL Server operation, but the data will be backed up in the most recent state.

Regardless of the value of this option, to ensure the database consistency, the software skips databases that are *not* in the **SYNCHRONIZED** or **SYNCHRONIZING** states when the backup starts. If all databases are skipped, the backup fails.

Microsoft Exchange Server

This option determines the backup mode for Exchange Server Database Availability Groups (DAG). For this option to be effective, Agent for Exchange must be installed on all of the DAG nodes. For more information about backing up Database Availability Groups, refer to "Protecting Database Availability Groups (DAG)".

The preset is: **Passive copy if possible**.

You can choose one of the following:

- **Passive copy if possible**
If all passive copies are offline, the active copy is backed up. Backing up the active copy may slow down the Exchange Server operation, but the data will be backed up in the most recent state.
- **Passive copy**
If all passive copies are offline, the backup will fail. Backing up passive copies does not affect the Exchange Server performance and allows you to extend the backup window. However, passive copies may contain information that is not up-to-date, because such copies are often set to be updated asynchronously (lagged).
- **Active copy**
If the active copy is offline, the backup will fail. Backing up the active copy may slow down the Exchange Server operation, but the data will be backed up in the most recent state.

Regardless of the value of this option, to ensure the database consistency, the software skips databases that are *not* in the **HEALTHY** or **ACTIVE** states when the backup starts. If all databases are skipped, the backup fails.

15.13.8 Compression level

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

The preset is: **Normal**.

A higher compression level means that the backup process takes longer, but the resulting backup occupies less space. Currently, the High and Maximum levels work similarly.

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.

15.13.9 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.

Cloud storage

If the cloud storage is selected as a backup destination, the option value is automatically set to **Enabled**. **Number of attempts: 300**. **Interval between attempts: 30 seconds**.

In this case, the actual number of attempts is unlimited, but the timeout before the backup failure is calculated as follows: $(300 \text{ seconds} + \text{Interval between attempts}) * (\text{Number of attempts} + 1)$.

Examples:

- With the default values, the backup will fail after $(300 \text{ seconds} + 30 \text{ seconds}) * (300 + 1) = 99330$ seconds, or ~27.6 hours.
- If you set **Number of attempts** to 1 and **Interval between attempts** to 1 second, the backup will fail after $(300 \text{ seconds} + 1 \text{ second}) * (1 + 1) = 602$ seconds, or ~10 minutes.

If the calculated timeout exceeds 30 minutes, and the data transfer has not started yet, the actual timeout is set to 30 minutes.

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.

15.13.10 Fast incremental/differential backup

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

This option is not effective (always disabled) for volumes formatted with the JFS, ReiserFS3, ReiserFS4, ReFS, or XFS file systems.

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.

15.13.11 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, unless stated otherwise.

To enable file filters

1. Select the data to back up.
2. Click **Change** next to **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.

Note This filter is not effective for file-level backup if **Version 11** is selected in **Backup format** (p. 148) and the backup destination is NOT cloud storage.

- **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 macOS).

Both in Windows and Linux/macOS, 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**

[Only for backups in the **Version 12** format] The double asterisk (**) substitutes for zero or more characters in a file name and path, including the slash character. For example, the criterion ****/Docs/**/*.txt** matches all txt files in all subfolders of all folders **Docs**.

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*.

15.13.12 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:

- If only machines running Linux are selected for backup: **Do not create a snapshot.**
- Otherwise: **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.

15.13.13 Forensic data

Malicious activities on a machine can be carried out by viruses, malware, and ransomware. The other case that may require investigations is stealing or changing data on a machine by means of different programs. Such activities may need to be investigated but it is possible only if you keep digital

evidence on a machine to investigate. Unfortunately, evidence (files, traces, and so on) may be deleted or a machine may become unavailable.

The backup option called **Forensic data** allows you to collect digital evidence that can be used in forensic investigations. The following items can be used as digital evidence: a snapshot of the unused disk space, memory dumps, and a snapshot of running processes. The **Forensic data** functionality is available only for an entire machine backup.

Currently, the **Forensic data** option is available only for Windows machines with the following OS versions:

- Windows 8.1, Windows 10
- Windows Server 2012 R2 – Windows Server 2019

Note

- *After a protection plan with the Backup module is applied to a machine, the forensic data settings cannot be modified. To use different forensic data settings, create a new protection plan.*
 - *Backups with forensic data collection are not supported for machines that are connected to your network through VPN and do not have direct access to the Internet.*
-

The supported locations for backups with forensic data are:

- Cloud storage
- Local folder

Note

1. *The local folder is supported only on an external hard disk connected via USB.*
 2. *Local dynamic disks are not supported as a location for forensic backups.*
-

- Network folder

Backups with forensic data are automatically notarized. Forensic backups allow investigators to analyze disk areas that are usually not included in a regular disk backup.

Forensic backup process

The system performs the following during a forensic backup process:

1. Collects raw memory dump and the list of running processes.
2. Automatically reboots a machine into the bootable media.
3. Creates the backup that includes both the occupied and unallocated space.
4. Notarizes the backed-up disks.
5. Reboots into the live operating system and continues plan execution (for example, replication, retention, validation and other).

To configure forensic data collection

1. In the service console, go to **Devices > All devices**. Alternatively, the protection plan can be created from the **Plans** tab.
2. Select the device and click **Protect**.
3. In the protection plan, enable the **Backup** module.
4. In **What to back up**, select **Entire machine**.
5. In **Backup options**, click **Change**.
6. Find the **Forensic data** option.
7. Enable **Collect forensic data**. The system will automatically collect a memory dump and create a snapshot of running processes.


Note Full memory dump may contain sensitive data such as passwords.

8. Specify the location.
9. Click **Run Now** to perform a backup with forensic data right away or wait until the backup is created according to the schedule.
10. Go to **Dashboard > Activities**, verify that the backup with forensic data was successfully created.


As a result, backups will include forensic data and you will be able to get them and analyze. Backups with forensic data are marked and can be filtered among other backups in **Backup storage > Locations** by using the **Only with forensic data** option.

How to get forensic data from a backup?

1. In the service console, go to **Backup storage**, select the location with backups that include forensic data.
2. Select the backup with forensic data and click **Show backups**.
3. Click **Recover** for the backup with forensic data.
 - To get only the forensic data, click **Forensic data**.



The system will show a folder with forensic data. Select a memory dump file or any other forensic file and click **Download**.



- To recover a full forensic backup, click **Entire machine**. The system will recover the backup without the boot mode. Thus, it will be possible to check that the disk was not changed.

You can use the provided memory dump with several of third-party forensic software, for example, use Volatility Framework at <https://www.volatilityfoundation.org/> for further memory analysis.

15.13.13.1 Notarization of backups with forensic data

To ensure that a backup with forensic data is exactly the image that was taken and it was not compromised, the backup module provides the notarization of backups with forensic data.

How it works


Notarization enables you to prove that a disk with forensic data is authentic and unchanged since it was backed up.

During a backup, the agent calculates the hash codes of the backed-up disks, builds a hash tree, saves the tree in the backup, and then sends the hash tree root to the notary service. The notary service saves the hash tree root in the Ethereum blockchain database to ensure that this value does not change.

When verifying the authenticity of the disk with forensic data, the agent calculates the hash of the disk, and then compares it with the hash that is stored in the hash tree inside the backup. If these hashes do not match, the disk is considered not authentic. Otherwise, the disk authenticity is guaranteed by the hash tree.

To verify that the hash tree itself was not compromised, the agent sends the hash tree root to the notary service. The notary service compares it with the one stored in the blockchain database. If the hashes match, the selected disk is guaranteed to be authentic. Otherwise, the software displays a message that the disk is not authentic.

The scheme below shows shortly the notarization process for backups with forensic data.



To verify the notarized disk backup manually, you can get the certificate for it and follow the verification procedure shown with the certificate by using the tibxread (p. 158) tool.

Getting the certificate for backups with forensic data

To get the certificate for a backup with forensic data from the console, do the following:

1. Go to **Backup storage** and select the backup with forensic data.
2. Recover the entire machine.
3. The system opens the **Disk Mapping** view.
4. Click the **Get certificate** icon for the disk.
5. The system will generate the certificate and open a new window in the browser with the certificate. Below the certificate you will see the instruction for manual verification of notarized disk backup.

15.13.13.2 The tool "tibxread" for getting the backed-up data

Cyber Protection provides the tool, called **tibxread**, for manual check of the backed-up disk integrity. The tool allows you to get data from a backup and calculate hash of the specified disk. The tool is installed automatically with the following components: Agent for Windows, Agent for Linux, and Agent for Mac.

The installation path: the same folder as the agent has (for example, **C:\Program Files\BackupClient\BackupAndRecovery**).

The supported locations are:

- The local disk
- The network folder (CIFS/SMB) that can be accessed without the credentials.
In case of a password-protected network folder, you can mount the network folder to the local folder by using the OS tools and then the local folder as the source for this tool.
- The cloud storage

You should provide the URL, port, and certificate. The URL and port can be obtained from the Windows registry key or configuration files on Linux/Mac machines.

For Windows:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\BackupAndRecovery\Settings\OnlineBackup\FesAddressCache\Default\<tenant_login>\FesUri
```

For Linux:

```
/etc/Acronis/BackupAndRecovery.config
```

For macOS:

```
/Library/Application Support/Acronis/Registry/BackupAndRecovery.config
```

The certificate can be found in the following locations:

For Windows:

```
%allusersprofile%\Acronis\BackupAndRecovery\OnlineBackup\Default
```

For Linux:

```
/var/lib/Acronis/BackupAndRecovery/OnlineBackup/Default
```

For macOS:

```
/Library/Application Support/Acronis/BackupAndRecovery/OnlineBackup/Default
```

The tool has the following commands:

- list backups
- list content
- get content
- calculate hash

list backups

Lists recovery points in a backup.

SYNOPSIS:

```
tibxread list backups --loc=URI --arc=BACKUP_NAME --raw
```

Options

```
--loc=URI  
--arc=BACKUP_NAME  
--raw  
--utc  
--log=PATH
```

Output template:

```
GUID    Date    Date timestamp  
----    -  
<guid> <date> <timestamp>
```

<guid> – a backup GUID.

<date> – a creation date of the backup. Format is “DD.MM.YYYY HH24:MM:SS”. In local timezone by default (can be changed by using the --utc option).

Output example:

GUID	Date	Date timestamp
516FCE73-5E5A-49EF-B673-A9EACB4093B8	18.12.2019	16:01:05 1576684865
516FCE73-5E5A-49EF-B673-A9EACB4093B9	18.12.2019	16:02:05 1576684925

list content

Lists content in a recovery point.

SYNOPSIS:

```
tibxread list content --loc=URI --arc=BACKUP_NAME --password
--backup=RECOVERY_POINT_ID --raw --log=PATH
```

Options

```
--loc=URI
--arc=BACKUP_NAME
--password
--backup=RECOVERY_POINT_ID
--raw
--log=PATH
```

Output template:

Disk	Size	Notarization status
<number>	<size>	<notarization_status>

<number> – identifier of the disk.

<size> – size in bytes.

<notarization_status> – the following statuses are possible: Without notarization, Notarized, Next backup.

Output example:

Disk	Size	Notary status
1	123123465798	Notarized
2	123123465798	Notarized

get content

Writes content of the specified disk in the recovery point to the standard output (stdout).

SYNOPSIS:

```
tibxread get content --loc=URI --arc=BACKUP_NAME --password --backup=RECOVERY_POINT_ID
--disk=DISK_NUMBER --raw --log=PATH --progress
```

Options

```
--loc=URI
--arc=BACKUP_NAME
--password
--backup=RECOVERY_POINT_ID
--disk=DISK_NUMBER
--raw
--log=PATH
--progress
```

calculate hash

Calculates the hash of the specified disk in the recovery point by using the SHA-256 algorithm and writes it to the stdout.

SYNOPSIS:

```
tibxread calculate hash --loc=URI --arc=BACKUP_NAME --password  
--backup=RECOVERY_POINT_ID --disk=DISK_NUMBER --raw --log=PATH --progress
```

Options

```
--loc=URI  
--arc=BACKUP_NAME  
--password  
--backup=RECOVERY_POINT_ID  
--disk=DISK_NUMBER  
--raw  
--log=PATH
```

Options description

Option	Description
--arc=BACKUP_NAME	The backup file name that you can get from the backup properties in the web console. The backup file must be specified with the extension .tibx.
--backup=RECOVERY_POINT_ID	The recovery point identifier
--disk=DISK_NUMBER	Disk number (the same as was written to the output of the "get content" command)
--loc=URI	A backup location URI. The possible formats of the "--loc" option are: - Local path name (Windows) c:/upload/backups - Local path name (Linux) /var/tmp - SMB/CIFS \\server\folder - Cloud storage --loc=<IP_address>:443 --cert=<path_to_certificate> [--storage_path=/1] <IP_address> – you can find it in the registry key in Windows: HKEY_LOCAL_MACHINE\SOFTWARE\Acronis\BackupAndRecovery\Settings\OnlineBackup\FesAddressCache\Default\<tenant_login>\FesUri <path_to_certificate> – a path to the certificate file to access Cyber Cloud. For example, in Windows this certificate is located in C:\ProgramData\Acronis\BackupAndRecovery\OnlineBackup\Default\<username>.crt where <username> – is your account name to access Cyber Cloud.
--log=PATH	Enables writing the logs by the specified PATH (local path only, format is the same as for --loc=URI parameter). Logging level is DEBUG.

<code>--password=PASSWORD</code>	An encryption password for your backup. If the backup is not encrypted, leave this value empty.
<code>--raw</code>	Hides the headers (2 first rows) in the command output. It is used when the command output should be parsed. Output example without "<code>--raw</code>": <pre> GUID Date Date timestamp ---- - 516FCE73-5E5A-49EF-B673-A9EACB4093B8 18.12.2019 16:01:05 1576684865 516FCE73-5E5A-49EF-B673-A9EACB4093B9 18.12.2019 16:02:05 1576684925 </pre> Output with "<code>--raw</code>": <pre> 516FCE73-5E5A-49EF-B673-A9EACB4093B8 18.12.2019 16:01:05 1576684865 516FCE73-5E5A-49EF-B673-A9EACB4093B9 18.12.2019 16:02:05 1576684925 </pre>
<code>--utc</code>	Shows dates in UTC
<code>--progress</code>	Shows progress of the operation. For example: <pre> 1% 2% 3% 4% ... 100% </pre>

15.13.14 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.

15.13.15 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.

15.13.16 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. 197) 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. 197).

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 protection 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. 197).

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.

15.13.17 Multi-volume snapshot

This option is effective for backups of physical machines running Windows or Linux.

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. 154) 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:

- If at least one machine running Windows is selected for backup: **Enabled**.
- Otherwise: **Disabled**.

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.

15.13.18 Performance and backup window

This option enables you to set one of three levels of backup performance (high, low, prohibited) for every hour within a week. This way, you can define a time window when backups are allowed to start and run. The high and low performance levels are configurable in terms of the process priority and output speed.

This option is not available for backups executed by the cloud agents, such as website backups or backups of servers located on the cloud recovery site.

You can configure this option separately for each location specified in the protection plan. To configure this option for a replication location, click the gear icon next to the location name, and then click **Performance and backup window**.

This option is effective only for the backup and backup replication processes. Post-backup commands and other operations included in a protection plan (for example, validation) will run regardless of this option.

The preset is: **Disabled**.

When this option is disabled, backups are allowed to run at any time, with the following parameters (no matter if the parameters were changed against the preset value):

- CPU priority: **Low** (in Windows, corresponds to **Below normal**).
- Output speed: **Unlimited**.

When this option is enabled, scheduled backups are allowed or blocked according to the performance parameters specified for the current hour. At the beginning of an hour when backups are blocked, a backup process is automatically stopped and an alert is generated.

Even if scheduled backups are blocked, a backup can be started manually. It will use the performance parameters of the most recent hour when backups were allowed.

Backup window

Each rectangle represents an hour within a week day. Click a rectangle to cycle through the following states:

- **Green:** backup is allowed with the parameters specified in the green section below.
- **Blue:** backup is allowed with the parameters specified in the blue section below.
This state is not available if the backup format is set to **Version 11**.
- **Gray:** backup is blocked.

You can click and drag to change the state of multiple rectangles simultaneously.

Performance and backup window settings

	AM				PM				AM			
	00	03	06	09	12	03	06	09	00			
Sun	■	■	■	■	■	■	■	■	■			
Mon	■	■	■	■	■	■	■	■	■			
Tue	■	■	■	■	■	■	■	■	■			
Wed	■	■	■	■	■	■	■	■	■			
Thu	■	■	■	■	■	■	■	■	■			
Fri	■	■	■	■	■	■	■	■	■			
Sat	■	■	■	■	■	■	■	■	■			

■ CPU priority Low

■ Output speed - 100 + % ▼

■ CPU priority Low

■ Output speed - 25 + % ▼

■ No backing up

CPU priority

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

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

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

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.



When this option is enabled, you can specify the maximum allowed output speed:

- As a percentage of the estimated writing speed of the destination hard disk (when backing up to a local folder) or of the estimated maximum speed of the network connection (when backing up to a network share or cloud storage).
This setting works only if the agent is running in Windows.
- In KB/second (for all destinations).

This option is effective if the backup destination is the cloud storage and the backup format (p. 148) is set to **Version 12**.

This option is effective for disk-level backups and file backups created by Agent for Windows, Agent for Linux, Agent for Mac, Agent for VMware, Agent for Hyper-V, and Agent for Virtuozzo.

This option determines whether the first full backup created by the protection plan will be sent to the cloud storage on a hard disk drive by using the Physical Data Shipping service. The subsequent incremental backups can be performed over the network.

The preset is: **Disabled**.

About the Physical Data Shipping service

The Physical Data Shipping service web interface is available only to administrators.

For detailed instructions about using the Physical Data Shipping service and the order creation tool, refer to the Physical Data Shipping Administrator's Guide. To access this document in the Physical Data Shipping service web interface, click the question mark icon.

Overview of the physical data shipping process

1. Create a new protection plan. In this plan, enable the **Physical Data Shipping** backup option. You can back up directly to the drive or back up to a local or a network folder, and then copy/move the backup(s) to the drive.

Important Once the initial full backup is done, the subsequent backups must be performed by the same protection plan. Another protection plan, even with the same parameters and for the same machine, will require another Physical Data Shipping cycle.

2. After the first backup is complete, use the Physical Data Shipping service web interface to download the order creation tool and create the order.
To access this web interface, log in to the management portal, click **Overview > Usage**, and then click **Manage service** under **Physical Data Shipping**.
3. Package the drives and ship them to the data center.


Important Ensure that you follow the packaging instructions provided in the Physical Data Shipping Administrator's Guide.

4. Track the order status by using the Physical Data Shipping service web interface. Note that the subsequent backups will fail until the initial backup is uploaded to the cloud storage.

15.13.20 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 protection 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").

15.13.20.1 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			
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.

15.13.20.2 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**.

15.13.21 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. 172) 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.

15.13.21.1 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.

15.13.21.2 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.

15.13.22 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: **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 protection plan is applied to the machine and remains the same until you edit the protection 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 protection plan is applied to multiple virtual machines. This option defines how many virtual machines an agent can back up simultaneously when executing the given protection plan.

If, according to the protection 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. However, if the agent executes multiple protection plans that overlap in time, the numbers specified in their options are added up. You can limit the total number of virtual machines (p. 286) that an agent can back up simultaneously, no matter how many protection plans are running.

Backups of physical machines will start exactly as scheduled.

15.13.23 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. 151) 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.

Note: *It will be impossible to perform a recovery of application data from the backups which were created in the sector-by-sector mode.*

15.13.24 Splitting

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

The preset is:

- If the backup location is a local or network (SMB) folder, and the backup format is Version 12: **Fixed size - 200 GB**

This setting allows the backup software to work with large volumes of data on the NTFS file system, without negative effects caused by file fragmentation.

- Otherwise: **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.

15.13.25 Task failure handling

This option determines the program behavior when a scheduled execution of a protection plan fails. This option is not effective when a protection plan is started manually.

If this option is enabled, the program will try to execute the protection 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**.

15.13.26 Task start conditions

This option is effective in Windows and Linux operating systems.

This option determines the program behavior in case a task is about to start (the scheduled time comes or the event specified in the schedule occurs), but the condition (or any of multiple conditions) is not met. For more information about conditions refer to "Start conditions".

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

Wait until the conditions from the schedule are met

With this setting, the scheduler starts monitoring the conditions and launches the task as soon as the conditions are met. If the conditions are never met, the task will never start.

To handle the situation when the conditions are not met for too long and further delaying the task is becoming risky, you can set the time interval after which the task will run irrespective of the condition. Select the **Run the task anyway after** check box and specify the time interval. The task will start as soon as the conditions are met OR the maximum time delay lapses, depending on which comes first.

Skip the task execution

Delaying a task might be unacceptable, for example, when you need to execute a task strictly at the specified time. Then it makes sense to skip the task rather than wait for the conditions, especially if the tasks occur relatively often.

15.13.27 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. 169) 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.

Note *If this option is enabled, files and folders that are specified in the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot` registry key are not backed up. In particular, offline Outlook Data Files (.ost) are not backed up because they are specified in the `OutlookOST` value of this key.*

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. 162) backup option.

15.13.28 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, Hyper-V Integration Services, or Virtio Guest Tools.

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. 151) 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.

15.13.29 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**.

15.13.30 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**.

15.14 Recovery

15.14.1 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. 178) Using bootable media (p. 183)
Physical machine (Mac)		Using bootable media (p. 183)
Virtual machine (VMware or Hyper-V)		Using the web interface (p. 182) Using bootable media (p. 183)
Virtual machine or container (Virtuozzo)		Using the web interface (p. 182)
ESXi configuration		Using bootable media (p. 192)
Files/Folders		Using the web interface (p. 187) Downloading files from the cloud storage (p. 188) Using bootable media (p. 190) Extracting files from local backups (p. 191)
System state		Using the web interface (p. 192)
SQL databases		Using the web interface (p. 212)
Exchange databases		Using the web interface (p. 215)
Exchange mailboxes		Using the web interface (p. 217)
Websites		Using the web interface (p. 266)
Microsoft Office 365	Mailboxes (local Agent for Office 365)	Using the web interface (p. 233)
	Mailboxes (cloud Agent for Office 365)	Using the web interface (p. 237)
	Public folders	Using the web interface (p. 239)
	OneDrive files	Using the web interface (p. 241)
	SharePoint Online data	Using the web interface (p. 244)
G Suite	Mailboxes	Using the web interface (p. 256)
	Google Drive files	Using the web interface (p. 259)
	Shared drive files	Using the web interface (p. 261)

What to recover		Recovery method

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.
- Starting with macOS Sierra 10.12, rarely used files can be moved to iCloud by the Store in Cloud feature. Small footprints of these files are kept on the file system. These footprints are backed up instead of the original files.
When you recover a footprint to the original location, it is synchronized with iCloud and the original file becomes available. When you recover a footprint to a different location, it cannot be synchronized and the original file will be unavailable.

15.14.2 Safe recovery

A backed-up OS image can have malware that can reinfect a machine after recovery.

The safe recovery functionality allows you to prevent recurrence of infections by using the integrated antimalware scanning (p. 347) and malware deletion during the recovery process.

Limitations:


- Safe recovery is supported only for physical or virtual Windows machines with Agent for Windows installed inside the machine.
- The supported backup types are "Entire machine" or "Disks/volumes" backups.
- Safe recovery is supported only for the volumes with NTFS file system. Non-NTFS partitions will be recovered without anti-malware scanning.
- Safe recovery is not supported for CDP backups (p. 118). The machine will be recovered based on the last regular backup without the data in the CDP backup. To recover the CDP data, start a **Files/folders** recovery.

How it works

If you enable the Safe recovery option during the recovery process, then the system will perform the following:


- Scan the image backup for malware and mark the infected files. One of the following statuses is assigned to a backup:
 - No malware** – no malware was found in a backup during scanning.
 - Malware detected** – malware was found in a backup during scanning.

- **Not scanned** – backup was not scanned for malware.



2. Recover the backup to the selected machine.
3. Delete the detected malware.

You can filter backups by using the **Status** parameter.



15.14.3 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 protection agent.

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

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. [Optional] Copy and print, or write down the registration token displayed by the service console. This token allows access to the cloud storage from bootable media without entering a login and password. It is necessary if you do not have a direct login to the cloud, but use third-party authentication instead.
3. 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.

To create bootable media in macOS

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.

15.14.4 Startup Recovery Manager

Startup Recovery Manager is a bootable component residing on the system disk in Windows, or on the /boot partition in Linux and configured to start at boot time on pressing F11. It eliminates the need for a separate media or network connection to start the bootable rescue utility.

Startup Recovery Manager is especially useful for traveling users. If a failure occurs, reboot the machine, wait for the prompt "Press F11 for Acronis Startup Recovery Manager..." to appear, and then press F11. The program will start and you can perform recovery.

You can also back up using Startup Recovery Manager, while on the move.

On machines with the GRUB boot loader installed, you select the Startup Recovery Manager from the boot menu instead of pressing F11.

A machine booted with Startup Recovery Manager can be registered on the management server similarly to a machine booted from bootable media. To do this, click **Tools > Register media on the management server**, and then follow the step-by-step procedure described in "Registering media on the management server".

Activating Startup Recovery Manager

On a machine running Agent for Windows or Agent for Linux, Startup Recovery Manager can be activated by using the service console.

To activate Startup Recovery Manager in the service console

1. Select the machine that you want to activate Startup Recovery Manager on.
2. Click **Details**.
3. Enable the **Startup Recovery Manager** switch.
4. Wait while the software activates Startup Recovery Manager.

To activate Startup Recovery Manager on a machine without an agent

1. Boot the machine from bootable media.
2. Click **Tools > Activate Startup Recovery Manager**.
3. Wait while the software activates Startup Recovery Manager.

What happens when you activate Startup Recovery Manager

Activation enables the boot-time prompt "Press F11 for Acronis Startup Recovery Manager..." (if you do not have the GRUB boot loader) or adds the "Startup Recovery Manager" item to GRUB's menu (if you have GRUB).

The system disk (or, the /boot partition in Linux) should have at least 100 MB of free space to activate Startup Recovery Manager.

Unless you use the GRUB boot loader and it is installed in the Master Boot Record (MBR), Startup Recovery Manager activation overwrites the MBR with its own boot code. Thus, you may need to reactivate third-party boot loaders if they are installed.

Under Linux, when using a boot loader other than GRUB (such as LILO), consider installing it to a Linux root (or boot) partition boot record instead of the MBR before activating Startup Recovery Manager. Otherwise, reconfigure the boot loader manually after the activation.

Deactivating Startup Recovery Manager

Deactivation is performed similarly to activation.

Deactivation disables the boot time prompt "Press F11 for Acronis Startup Recovery Manager..." (or, the menu item in GRUB). If Startup Recovery Manager is not activated, you will need one of the following to recover the system when it fails to boot:

- boot the machine from a separate bootable media
- use network boot from a PXE server or Microsoft Remote Installation Services (RIS)

15.14.5 Recovering a machine

Related topics

Running a virtual machine from a backup (Instant Restore) 268

15.14.5.1 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:

- macOS
- Any operating system to bare metal or to an offline machine
- The structure of logical volumes (volumes created by Logical Volume Manager in Linux). The media enables you to recreate the logical volume structure automatically.

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 location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select a target machine that is online, and then select a recovery point.
 - Select a recovery point on the Backup storage tab (p. 199).
 - Recover the machine as described in "Recovering disks by using bootable media" (p. 183).
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.

× Recover machine ?

RECOVER TO
Physical machine ▼

TARGET MACHINE
ssd-win2016

DISK MAPPING
Disk 1 → Disk 1
Disk 2 → Disk 2
Disk 3 → Disk 3


SAFE RECOVERY
☐ Off i

START RECOVERY

⚙ RECOVERY OPTIONS

5. If you are unsatisfied with the mapping result or if the disk mapping fails, click **Volume mapping** to re-map the disks manually.

The mapping section also enables you to choose individual disks or volumes for recovery. You can switch between recovering disks and volumes by using the **Switch to...** link in the top-right corner.



6. [Optional] Enable **Safe recovery** to scan the backup for malware. If malware is detected, it will be marked in the backup and deleted right after the recovery process is completed.
7. Click **Start recovery**.
8. 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.

15.14.5.2 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.

Note You cannot recover MacOS virtual machines to Hyper-V hosts, because Hyper-V does not support MacOS. You can recover MacOS virtual machines to a VMware host that is installed on Mac hardware.

For more information about P2V migration, refer to "Machine migration" (p. 287).

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 location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select a machine that is online, and then select a recovery point.

- Select a recovery point on the Backup storage tab (p. 199).
 - Recover the machine as described in "Recovering disks by using bootable media" (p. 183).
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 **Disk mapping** to select the datastore (storage), interface, and provisioning mode for each virtual disk. The mapping section also enables you to choose individual disks for recovery.
 - Click **VM settings** to change the memory size, the number of processors, and the network connections of the virtual machine.

RECOVER TO
Virtual machine

TARGET MACHINE
New machine on 10.250.22.17 New

DATASTORE
datastore1 (1)

DISK MAPPING
Disk 1 → datastore1 (1), 50.0 GB
Disk 2 → datastore1 (1), 50.0 GB

VM SETTINGS
Memory: 2.00 GB
Virtual processors: 2
Network adapters: 2

START RECOVERY ⚙️ RECOVERY OPTIONS

8. [Optional] Enable **Safe recovery** to scan the backup for malware. If malware is detected, it will be marked in the backup and deleted right after the recovery process is completed.
9. Click **Start recovery**.
10. When recovering to an existing virtual machine, confirm that you want to overwrite the disks.

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

15.14.5.3 Virtual machine

You can recover virtual machines from their backups.

Prerequisites

- A virtual machine must be stopped during the recovery to this machine. By default, the software stops the machine without a prompt. When the recovery is completed, you have to start the machine manually. You can change the default behavior by using the VM power management recovery option (click **Recovery options > VM power management**).

Procedure

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 Backup storage tab (p. 199).
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. 178). Otherwise, we recommend that you perform the V2P migration by using bootable media (p. 183).
4. [Optional] By default, 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, Hyper-V, Virtuozzo, or Virtuozzo Hybrid Infrastructure**).

Only Virtuozzo virtual machines can be recovered to Virtuozzo. For more information about V2V migration, refer to "Machine migration" (p. 287).
 - 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. Setup up the additional recovery options that you need.
 - [Not available for Virtuozzo Hybrid Infrastructure] To select the datastore for the virtual machine, click **Datastore** for ESXi or **Path** for Hyper-V and Virtuozzo, and then select the datastore (storage) for the virtual machine.
 - [Optional] To view the datastore (storage), interface, and the provisioning mode for each virtual disk, click **Disk mapping**. You can change these settings, unless you are recovering a Virtuozzo container or Virtuozzo Hybrid Infrastructure virtual machine.

For Virtuozzo Hybrid Infrastructure, you can only select the storage policy for the target disks. To do so, select the desired target disk, and then click **Change**. In the blade that opens, click the gear icon, select the storage policy, and then click **Done**.

The mapping section also enables you to choose individual disks for recovery.

- [Optional for VMware ESXi, Hyper-V, and Virtuozzo] Click **VM settings** to change the memory size and the number of processors (for Virtuozzo Hybrid Infrastructure: select **Flavor**), or the network connections of the virtual machine.

Note For Virtuozzo Hybrid Infrastructure, selecting flavor is a required step.

RECOVER TO
Virtual machine

TARGET MACHINE

New machine on 10.250.22.17 New

DATASTORE

datastore1 (1)

DISK MAPPING

Disk 1 → datastore1 (1), 50.0 GB

Disk 2 → datastore1 (1), 50.0 GB

VM SETTINGS

Memory: 2.00 GB

Virtual processors: 2

Network adapters: 2

START RECOVERY

 RECOVERY OPTIONS

6. [Optional] Enable **Safe recovery** to scan the backup for malware. If malware is detected, it will be marked in the backup and deleted right after the recovery process is completed.
7. Click **Start recovery**.
8. When recovering to an existing virtual machine, confirm that you want to overwrite the disks.

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

15.14.5.4 Recovering disks by using bootable media

For information about how to create bootable media, refer to "Creating bootable media" (p. 176).

To recover disks by using bootable media

1. Boot the target machine by using bootable media.
2. [Only when recovering a Mac] If you are recovering APFS-formatted disks/volumes to a non-original machine or to bare metal, re-create the original disk configuration manually:
 - a. Click **Disk Utility**.
 - b. Erase and format the target disk into APFS. For instructions, refer to <https://support.apple.com/en-us/HT208496#erasedisk>.

- c. Re-create the original disk configuration. For instructions, refer to <https://support.apple.com/guide/disk-utility/add-erase-or-delete-apfs-volumes-dskua9e6a110/19.0/mac/10.15>.
 - d. Click **Disk Utility > Quit Disk Utility**.
 3. Click **Manage this machine locally** or click **Rescue Bootable Media** twice, depending on the media type you are using.
 4. If a proxy server is enabled in your network, click **Tools > Proxy server**, and then specify the proxy server host name/IP address, port, and credentials. Otherwise, skip this step.
 5. [Optional] When recovering Windows or Linux, click **Tools > Register media in the Cyber Protection service**, and then specify the registration token that you obtained when downloading the media. If you do this, you will not need to enter credentials or a registration code to access the cloud storage, as described in step 8.
 6. On the welcome screen, click **Recover**.
 7. Click **Select data**, and then click **Browse**.
 8. 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.
When recovering Windows or Linux, you have the option to request a registration code and use it instead of the credentials. Click **Use registration code > Request the code**. The software shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. The registration code is valid for one hour.
 - 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.
 9. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
 10. In **Backup contents**, select the disks that you want to recover. Click **OK** to confirm your selection.
 11. 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.*
-
12. [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.
 13. [Optional] Click **Recovery options** to specify additional settings.
 14. Click **OK** to start the recovery.

15.14.5.5 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. 185).
5. Click **OK**.

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.

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.

15.14.6 Recovering files

15.14.6.1 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. Do any of the following:

- [Recommended] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select a target machine that is online, and then select a recovery point.
 - Select a recovery point on the Backup storage tab (p. 199).
 - Download the files from the cloud storage (p. 188).
 - Use bootable media (p. 190).
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. 153).

Note Search is not available for disk-level backups that are stored in the cloud storage.

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.

Downloading is not available if your selection contains folders or the total size of the selected files exceeds 100 MB.

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 an agent is installed on this machine).
- The machine where Agent for VMware, Agent for Hyper-V, or Agent for Virtuozzo is installed (if the files originate from an ESXi, Hyper-V, or Virtuozzo virtual machine).

This is the target machine for the recovery. You can select another machine, if necessary.

9. 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

Note *Symbolic links are not supported.*

- A network folder that is accessible from the target machine.
1. Click **Start recovery**.
 2. 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.

15.14.6.2 Downloading files from the cloud storage


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

Limitations

- Backups of system state, SQL databases, and Exchange databases cannot be browsed.
- For a better downloading experience, download no more than 100 MB at a time. To quickly retrieve larger amounts of data from the cloud, use the file recovery procedure (p. 187).

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.




.. > ABR11MMS > ABR11MMS-New Backup Plan

Versions ^

 Backup #10	14/01/15 08:43	Size: 21.52 MB
 Backup #1	14/01/15 07:32	Size: 3.05 GB

[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.



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**.

15.14.6.3 Verifying file authenticity with Notary Service

If notarization was enabled during backup (p. 140), you can verify the authenticity of a backed-up file.

To verify the file authenticity

1. Select the file as described in steps 1-6 of the "Recovering files by using the web interface" (p. 187) section, or steps 1-5 of the "Downloading files from the cloud storage (p. 188)" section.
2. Ensure that the selected file is marked with the following icon: . This means that the file is notarized.
3. Do one of the following:
 - Click **Verify**.
The software checks the file authenticity and displays the result.
 - Click **Get certificate**.
A certificate that confirms the file notarization is opened in a web browser window. The window also contains instructions that allow you to verify the file authenticity manually.

15.14.6.4 Signing a file with ASign

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

ASign is a service that allows multiple people to sign a backed-up file electronically. This feature is available only for file-level backups stored in the cloud storage.

Only one file version can be signed at a time. If the file was backed up multiple times, you must choose the version to sign, and only this version will be signed.

For example, ASign can be used for electronic signing of the following files:

- Rental or lease agreements
- Sales contracts
- Asset purchase agreements
- Loan agreements
- Permission slips
- Financial documents
- Insurance documents
- Liability waivers
- Healthcare documents
- Research papers
- Certificates of product authenticity
- Nondisclosure agreements
- Offer letters
- Confidentiality agreements
- Independent contractor agreements

To sign a file version

1. Select the file as described in steps 1-6 of the "Recovering files by using the web interface" (p. 187) section, or steps 1-5 of the "Downloading files from the cloud storage (p. 188)" section.
2. Ensure that the correct date and time is selected on the left panel.
3. Click **Sign this file version**.
4. Specify the password for the cloud storage account under which the backup is stored. The login of the account is displayed in the prompt window.
The ASign service interface is opened in a web browser window.
5. Add other signees by specifying their email addresses. It is not possible to add or remove signees after sending invitations, so ensure that the list includes everyone whose signature is required.
6. Click **Invite to sign** to send invitations to the signees.
Each signee receives an email message with the signature request. When all the requested signees sign the file, it is notarized and signed through the notary service.
You will receive notifications when each signee signs the file and when the entire process is complete. You can access the ASign web page by clicking **View details** in any of the email messages that you receive.
7. Once the process is complete, go to the ASign web page and click **Get document** to download a .pdf document that contains:
 - The Signature Certificate page with the collected signatures.
 - The Audit Trail page with history of activities: when the invitation was sent to the signees, when each signee signed the file, and so on.

15.14.6.5 Recovering files by using bootable media

For information about how to create bootable media, refer to "Creating bootable media" (p. 176).

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, port, and credentials. Otherwise, skip this step.
4. [Optional] When recovering Windows or Linux, click **Tools > Register media in the Cyber Protection service**, and then specify the registration token that you obtained when downloading the media. If you do this, you will not need to enter credentials or a registration code to access the cloud storage, as described in step 7.
5. On the welcome screen, click **Recover**.
6. Click **Select data**, and then click **Browse**.
7. 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.
When recovering Windows or Linux, you have the option to request a registration code and use it instead of the credentials. Click **Use registration code > Request the code**. The software shows the registration link and the registration code. You can copy them and perform the registration steps on a different machine. The registration code is valid for one hour.
 - 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.
8. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
9. In **Backup contents**, select **Folders/files**.
10. Select the data that you want to recover. Click **OK** to confirm your selection.
11. Under **Where to recover**, specify a folder. Optionally, you can prohibit overwriting of newer versions of files or exclude some files from recovery.
12. [Optional] Click **Recovery options** to specify additional settings.
13. Click **OK** to start the recovery.

15.14.6.6 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 protection 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, FAT32, NTFS, ReFS, Ext2, Ext3, Ext4, XFS, or HFS+.
- The backup must be stored in a local folder or on a network share (SMB/CIFS).

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> - <protection 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.

15.14.7 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.

15.14.8 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. 176).

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. On the welcome screen, click **Recover**.
4. Click **Select data**, and then click **Browse**.
5. Specify the backup location:
 - Browse to the folder under **Local folders** or **Network folders**.Click **OK** to confirm your selection.
6. In **Show**, select **ESXi configurations**.
7. Select the backup from which you want to recover the data. If prompted, type the password for the backup.
8. Click **OK**.
9. 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.
10. 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**.
11. [Optional] In **Network mapping**, change the result of automatic mapping of the virtual switches present in the backup to the physical network adapters.

12. [Optional] Click **Recovery options** to specify additional settings.
13. Click **OK** to start the recovery.

15.14.9 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, macOS, 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	macOS	Bootable media	ESXi, Hyper-V, and Virtuozzo	Windows
Backup validation (p. 194)	+	+	+	+	+	+	+	+	+
Boot mode (p. 194)	+	-	-	-	-	-	-	+	-
Date and time for files (p. 195)	-	-	-	+	+	+	+	-	-
Error handling (p. 195)	+	+	+	+	+	+	+	+	+
File exclusions (p. 196)	-	-	-	+	+	+	+	-	-
File-level security (p. 196)	-	-	-	+	-	-	-	-	-
Flashback (p. 196)	+	+	+	-	-	-	-	+	-
Full path recovery (p. 197)	-	-	-	+	+	+	+	-	-
Mount points (p. 197)	-	-	-	+	-	-	-	-	-
Performance (p. 197)	+	+	-	+	+	+	-	+	+
Pre/post commands (p. 197)	+	+	-	+	+	+	-	+	+
SID changing (p. 199)	+	-	-	-	-	-	-	-	-

	Disks			Files				Virtual machines	SQL and Exchange
	Windows	Linux	Bootable media	Windows	Linux	macOS	Bootable media	ESXi, Hyper-V, and Virtuozzo	Windows
VM power management (p. 199)	-	-	-	-	-	-	-	+	-
Windows event log (p. 199)	+	-	-	+	-	-	-	Hyper-V only	+

15.14.9.1 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.

Note Validation is available for cloud storage located in an Acronis data center and provided by Acronis partners.

15.14.9.2 Boot mode

This option is effective when recovering a physical or a virtual machine from a disk-level backup that contains a Windows operating system.

This option enables you to select the boot mode (BIOS or UEFI) that Windows will use after the recovery. If the boot mode of the original machine is different from the selected boot mode, the software will:

- Initialize the disk to which you are recovering the system volume, according to the selected boot mode (MBR for BIOS, GPT for UEFI).
- Adjust the Windows operating system so that it can start using the selected boot mode.

The preset is: **As on the target machine**.

You can choose one of the following:

- **As on the target machine**

The agent that is running on the target machine detects the boot mode currently used by Windows and makes the adjustments according to the detected boot mode.

This is the safest value that automatically results in bootable system unless the limitations listed below apply. Since the **Boot mode** option is absent under bootable media, the agent on media always behaves as if this value is chosen.

- **As on the backed-up machine**

The agent that is running on the target machine reads the boot mode from the backup and makes the adjustments according to this boot mode. This helps you recover a system on a different machine, even if this machine uses another boot mode, and then replace the disk in the backed-up machine.

- **BIOS**

The agent that is running on the target machine makes the adjustments to use BIOS.

- **UEFI**

The agent that is running on the target machine makes the adjustments to use UEFI.

Once a setting is changed, the disk mapping procedure will be repeated. This will take some time.

Recommendations

If you need to transfer Windows between UEFI and BIOS:

- Recover the entire disk where the system volume is located. If you recover only the system volume on top of an existing volume, the agent will not be able to initialize the target disk properly.
- Remember that BIOS does not allow using more than 2 TB of disk space.

Limitations

- Transferring between UEFI and BIOS is supported for:
 - 64-bit Windows operating systems starting with Windows Vista SP1
 - 64-bit Windows Server operating systems starting with Windows Server 2008 SP1
- Transferring between UEFI and BIOS is not supported if the backup is stored on a tape device.

When transferring a system between UEFI and BIOS is not supported, the agent behaves as if the **As on the backed-up machine** setting is chosen. If the target machine supports both UEFI and BIOS, you need to manually enable the boot mode corresponding to the original machine. Otherwise, the system will not boot.

15.14.9.3 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**.

15.14.9.4 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.

Save system information if a recovery with reboot fails

This option is effective for a disk or volume recovery to a physical machine running Windows or Linux.

The preset is: **Disabled**.

When this option is enabled, you can specify a folder on the local disk (including flash or HDD drives attached to the target machine) or on a network share where the log, system information, and crash dump files will be saved. This file will help the technical support personnel to identify the problem.

15.14.9.5 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.

15.14.9.6 File-level security

This option is effective when recovering files from disk- and file-level backups of NTFS-formatted volumes.

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

The preset is: **Enabled**.

You can choose whether to recover the permissions or let the files inherit their NTFS permissions from the folder to which they are recovered.

15.14.9.7 Flashback

This option is effective when recovering disks and volumes on physical and virtual machines, except for Mac.

This option works only if the volume layout of the disk being recovered exactly matches that of the target disk.

If the option is enabled, only the differences between the data in the backup and the target disk data are recovered. This accelerates recovery of physical and virtual machines. The data is compared at the block level.

When recovering a physical machine, the preset is: **Disabled**.

When recovering a virtual machine, the preset is: **Enabled**.

15.14.9.8 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**.

15.14.9.9 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. 163) 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.

15.14.9.10 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.

15.14.9.11 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.

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 command execution fails*	Selected	Cleared	Selected	Cleared
Do not recover until the command execution is complete	Selected	Selected	Cleared	Cleared
Result				
	Preset Perform the recovery only after the command is successfully executed. Fail the 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.

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.

15.14.9.12SID 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.

15.14.9.13VM power management

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

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.

15.14.9.14Windows 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**.

15.15 Operations with backups

15.15.1 The Backup storage tab

The **Backup storage** tab provides access to all backups, including backups of offline machines and machines that are no longer registered in the Cyber Protection service.

Backups that are stored in a shared location (such as an SMB or NFS share) are visible to all users that have the read permission for the location.

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

Backups created by the *cloud* Agent for Office 365 and backups of G Suite data are shown not in the **Cloud storage** location, but in a separate section named **Cloud applications backups**.

Backup locations that are used in protection plans are automatically added to the **Backup storage** 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.

If you added or removed some backups by using a file manager, click the gear icon next to the location name, and then click **Refresh**.

A backup location (except for the cloud storage) disappears from the **Backup storage** tab if all machines that had ever backed up to the location were deleted from the Cyber Protection service. This ensures that you do not have to pay for the backups stored in this location. As soon as a backup to this location occurs, the location is re-added along with all backups that are stored in it.

On the **Backup storage** tab, you can filter backups in the list by using the following criteria:

- **Only with forensic data** – only backups having forensic data (p. 154) will be shown.
- **Only pre-update backups created by Patch management** – only backups that were created during patch management run before patch installation (p. 356) will be shown.

To select a recovery point by using the Backup storage tab

1. On the **Backup storage** 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> - <protection 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.

15.15.2 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> - <protection 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.

15.15.3 Deleting backups

Warning When a backup is deleted, all of its data is permanently erased. Deleted data cannot be recovered.

To delete backups of a machine that is online and present in the service 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 X sign.
 - To delete all backups in the selected location, click **Delete all**.
5. Confirm your decision.

To delete backups of any machine

1. On the **Backup storage** 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> - <protection 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 X sign.

- To delete the selected group, click **Delete**.
4. Confirm your decision.

To delete backups directly from the cloud storage

1. Log in to the cloud storage, as described in "Downloading files from the cloud storage" (p. 188).
2. Click the name of the machine whose backups you want to delete.
The software displays one or more backup groups.
3. Click the gear icon corresponding to the backup group that you want to delete.
4. Click **Remove**.
5. Confirm the operation.

What to do if you deleted local backups by using a file manager

We recommend that you delete backups by using the service console, whenever possible. If you deleted local backups by using a file manager, do the following:

1. On the **Backup storage** tab, click the gear icon next to the location name.
2. Click **Refresh**.

This way you will inform the Cyber Protection service that the local storage usage is decreased.

15.16 Protecting Microsoft 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 protection plan can be used for both disaster recovery and data protection purposes.

For Microsoft Exchange Server, you can opt for **Mailbox backup**. This is a backup of individual mailboxes via the Exchange Web Services protocol. The mailboxes or mailbox items can be recovered to a live Exchange Server or to Microsoft Office 365. Mailbox backup is supported for Microsoft Exchange Server 2010 Service Pack 1 (SP1) and later.

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. 212) Databases as files (p. 212)	Entire machine (p. 178) Databases to a live SQL Server instance (p. 212) Databases as files (p. 212)	Entire machine (p. 178)
Microsoft Exchange Server	Databases to a live Exchange (p. 215) Databases as files (p. 215) Granular recovery to a live Exchange or to Office 365* (p. 217)	Entire machine (p. 178) Databases to a live Exchange (p. 215) Databases as files (p. 215) Granular recovery to a live Exchange or to Office 365* (p. 217)	Entire machine (p. 178)
Microsoft SharePoint database servers	Databases to a live SQL Server instance (p. 212) Databases as files (p. 212) Granular recovery by using SharePoint Explorer	Entire machine (p. 178) Databases to a live SQL Server instance (p. 212) Databases as files (p. 212) Granular recovery by using SharePoint Explorer	Entire machine (p. 178)
Microsoft SharePoint front-end web servers	-	-	Entire machine (p. 178)
Active Directory Domain Services	-	Entire machine (p. 178)	-

* Granular recovery is also available from a mailbox backup. Recovery of Exchange data items to Office 365, and vice versa, is supported on the condition that Agent for Office 365 is installed locally.

15.16.1 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.
- 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 protection plan, ensure that:

- For physical machines and machines with the agent installed inside, the Volume Shadow Copy Service (VSS) (p. 172) backup option is enabled.
- For virtual machines, the Volume Shadow Copy Service (VSS) for virtual machines (p. 173) backup option is enabled.

Additional requirements for application-aware backups

When creating a protection plan, ensure that **Entire machine** is selected for backup. The **Sector-by-sector** backup option must be disabled in a protection plan, otherwise it will be impossible to perform a recovery of application data from such backups. If the plan is executed in the **Sector-by-sector** mode due to an automatic switch to this mode, then recovery of application data will also be impossible.

Requirements for ESXi virtual machines

If the application runs on a virtual machine that is backed up by Agent for VMware, ensure that:

- The virtual machine being backed up meets the requirements for application-consistent backup and restore listed in the article "Windows Backup Implementations" in the VMware documentation:
<https://code.vmware.com/docs/1674/virtual-disk-programming-guide/doc/vddkBkupVadp.9.6.html>.
- VMware Tools is installed and up-to-date on the machine.
- User Account Control (UAC) is disabled on the machine. 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.

Requirements for Hyper-V virtual machines

If the application runs on a virtual machine that is backed up by Agent for Hyper-V, ensure that:

- The guest operating system is Windows Server 2008 or later.
- For Hyper-V 2008 R2: the guest operating system is Windows Server 2008/2008 R2/2012.
- The virtual machine has no dynamic disks.
- The network connection exists between the Hyper-V host and the guest operating system. This is required to execute remote WMI queries inside the virtual machine.
- User Account Control (UAC) is disabled on the machine. 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.
- The virtual machine configuration matches the following criteria:

- Hyper-V Integration Services is installed and up-to-date. The critical update is <https://support.microsoft.com/en-us/help/3063109/hyper-v-integration-components-update-for-windows-virtual-machines>
- In the virtual machine settings, the **Management > Integration Services > Backup (volume checkpoint)** option is enabled.
- For Hyper-V 2012 and later: the virtual machine has no checkpoints.
- For Hyper-V 2012 R2 and later: the virtual machine has a SCSI controller (check **Settings > Hardware**).

15.16.2 Database backup

Before backing up databases, ensure that the requirements listed in "Prerequisites" (p. 203) are met.

Select the databases as described below, and then specify other settings of the protection plan as appropriate (p. 111).

15.16.2.1 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 protection plan options (p. 162).

To select SQL databases

1. Click **Devices > Microsoft SQL**.

The software shows the tree of SQL Server Always On Availability Groups (AAG), machines running Microsoft SQL Server, SQL Server instances, and databases.

2. Browse to the data that you want to back up.

Expand the tree nodes or double-click items in the list to the right of the tree.

3. Select the data that you want to back up. You can select AAGs, machines running SQL Server, SQL Server instances, or individual databases.

- If you select an AAG, all databases that are included into the selected AAG will be backed up. For more information about backing up AAGs or individual AAG databases, refer to "Protecting Always On Availability Groups (AAG)" (p. 206).
- If you select a machine running an SQL Server, all databases that are attached to all SQL Server instances running on the selected machine will be backed up.
- If you select a SQL Server instance, all databases that are attached to the selected instance will be backed up.
- If you select databases directly, only the selected databases will be backed up.

4. Click **Protect**. 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.

15.16.2.2 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
2010/2013/2016/2019	Databases, Database Availability Groups (DAG)	Membership in the Server 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 **Devices > Microsoft Exchange**.

The software shows the tree of Exchange Server Database Availability Groups (DAG), machines running Microsoft Exchange Server, and Exchange Server databases. If you configured Agent for Exchange as described in "Mailbox backup" (p. 211), mailboxes are also shown in this tree.

2. Browse to the data that you want to back up.

Expand the tree nodes or double-click items in the list to the right of the tree.

3. Select the data that you want to back up.

- If you select a DAG, one copy of each clustered database will be backed up. For more information about backing up DAGs, refer to "Protecting Database Availability Groups (DAG)".
- If you select a machine running Microsoft Exchange Server, all databases that are mounted to the Exchange Server running on the selected machine will be backed up.
- If you select databases directly, only the selected databases will be backed up.
- If you configured Agent for Exchange as described in "Mailbox backup" (p. 211), you can select mailboxes for backup.

4. If prompted, provide the credentials to access the data.

5. Click **Protect**.

15.16.2.3 Protecting Always On Availability Groups (AAG)

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

SQL Server high-availability solutions overview

The Windows Server Failover Clustering (WSFC) functionality enables you to configure a highly available SQL Server through redundancy at the instance level (Failover Cluster Instance, FCI) or at the database level (AlwaysOn Availability Group, AAG). You can also combine both methods.

In a Failover Cluster Instance, SQL databases are located on a shared storage. This storage can only be accessed from the active cluster node. If the active node fails, a failover occurs and a different node becomes active.

In an availability group, each database replica resides on a different node. If the primary replica becomes not available, a secondary replica residing on a different node is assigned the primary role.

Thus, the clusters are already serving as a disaster recovery solution themselves. However, there might be cases when the clusters cannot provide data protection: for example, in case of a database logical corruption, or when the entire cluster is down. Also cluster solutions do not protect from harmful content changes, as they usually immediately replicate to all cluster nodes.

Supported cluster configurations

This backup software supports *only* the Always On Availability Group (AAG) for SQL Server 2012 or later. Other cluster configurations, such as Failover Cluster Instances, database mirroring, and log shipping are *not* supported.

How many agents are required for cluster data backup and recovery?

For successful data backup and recovery of a cluster Agent for SQL has to be installed on each node of the WSFC cluster.

Backing up databases included in an AAG

1. Install Agent for SQL on each node of the WSFC cluster.

Tip After you install the agent on one of the nodes, the software displays the AAG and its nodes under **Devices > Microsoft SQL > Databases**. To install Agents for SQL on the rest of the nodes, select the AAG, click **Details**, and then click **Install agent** next to each of the nodes.

2. Select the AAG to backup as described in "Selecting SQL databases".

You must select the AAG itself to backup all databases of the AAG. To backup a set of databases, define this set of databases in all nodes of the AAG.

Warning The database set must be exactly the same in all nodes. If even one set is different, or not defined on all nodes, the cluster backup will not work correctly.

3. Configure the "Cluster backup mode" (p. 150) backup option.

Recovery of databases included in an AAG

1. Select the databases that you want to recover, and then select the recovery point from which you want to recover the databases.

When you select a clustered database under **Devices > Microsoft SQL > Databases**, and then click **Recover**, the software shows only the recovery points that correspond to the times when the selected copy of the database was backed up.

The easiest way to view all recovery points of a clustered database is to select the backup of the entire AAG on the Backup storage tab (p. 199). The names of AAG backups are based on the following template: <AAG name> - <protection plan name> and have a special icon.

2. To configure recovery, follow the steps described in "Recovering SQL databases" (p. 212), starting from step 5.

The software automatically defines a cluster node to which the data will be recovered. The node's name is displayed in the **Recover to** field. You can manually change the target node.

Important A database that is included in an Always On Availability Group cannot be overwritten during a recovery because Microsoft SQL Server prohibits this. You need to exclude the target database from the AAG before the recovery. Or, just recover the database as a new non-AAG one. When the recovery is completed, you can reconstruct the original AAG configuration.

15.16.2.4 Protecting Database Availability Groups (DAG)

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

Exchange Server clusters overview

The main idea of Exchange clusters is to provide high database availability with fast failover and no data loss. Usually, it is achieved by having one or more copies of databases or storage groups on the members of the cluster (cluster nodes). If the cluster node hosting the active database copy or the active database copy itself fails, the other node hosting the passive copy automatically takes over the operations of the failed node and provides access to Exchange services with minimal downtime. Thus, the clusters are already serving as a disaster recovery solution themselves.

However, there might be cases when failover cluster solutions cannot provide data protection: for example, in case of a database logical corruption, or when a particular database in a cluster has no copy (replica), or when the entire cluster is down. Also cluster solutions do not protect from harmful content changes, as they usually immediately replicate to all cluster nodes.


Cluster-aware backup

With cluster-aware backup, you back up only one copy of the clustered data. If the data changes its location within the cluster (due to a switchover or a failover), the software will track all relocations of this data and safely back it up.

Supported cluster configurations

Cluster-aware backup is supported *only* for Database Availability Group (DAG) in Exchange Server 2010 or later. Other cluster configurations, such as Single Copy Cluster (SCC) and Cluster Continuous Replication (CCR) for Exchange 2007, are *not* supported.

DAG is a group of up to 16 Exchange Mailbox servers. Any node can host a copy of mailbox database from any other node. Each node can host passive and active database copies. Up to 16 copies of each database can be created.



How many agents are required for cluster-aware backup and recovery?

For successful backup and recovery of clustered databases, Agent for Exchange has to be installed on each node of the Exchange cluster.

Tip After you install the agent on one of the nodes, the service console displays the DAG and its nodes under **Devices > Microsoft Exchange > Databases**. To install Agents for Exchange on the rest of the nodes, select the DAG, click **Details**, and then click **Install agent** next to each of the nodes.

Backing up the Exchange cluster data

1. When creating a protection plan, select the DAG as described in "Selecting Exchange Server data" (p. 206).
2. Configure the "Cluster backup mode" (p. 150) backup option.
3. Specify other settings of the protection plan as appropriate (p. 111).

Important For cluster-aware backup, ensure to select the DAG itself. If you select individual nodes or databases inside the DAG, only the selected items will be backed up and the **Cluster backup mode** option will be ignored.

Recovering the Exchange cluster data

1. Select the recovery point for the database that you want to recover. Selecting an entire cluster for recovery is not possible.

When you select a copy of a clustered database under **Devices > Microsoft Exchange > Databases > <cluster name> > <node name>** and click **Recover**, the software shows only the recovery points that correspond to the times when this copy was backed up.

The easiest way to view all recovery points of a clustered database is to select its backup on the Backup storage tab (p. 199).

2. Follow the steps described in "Recovering Exchange databases", starting from step 5.

The software automatically defines a cluster node to which the data will be recovered. The node's name is displayed in the **Recover to** field. You can manually change the target node.

15.16.3 Application-aware backup

Application-aware disk-level backup is available for physical machines, ESXi virtual machines, and Hyper-V 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 protection plan options (p. 162). The Exchange transaction logs are truncated on virtual machines only. You can enable the VSS full backup option (p. 172) 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) or Agent for Hyper-V.

Agent for VMware (Virtual Appliance) can create application-aware backups, but cannot recover application data from them. To recover application data from backups created by this agent, you need Agent for VMware (Windows), Agent for SQL, or Agent for Exchange on a machine that has access to the location where the backups are stored. When configuring recovery of application data, select the recovery point on the **Backup storage** tab, and then select this machine in **Machine to browse from**.

Other requirements are listed in the "Prerequisites" (p. 203) and "Required user rights" (p. 210) sections.

15.16.3.1 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 **Administrators** group on the machine, and a member of the **Exchange Organization Administrators** role group.

Exchange 2010 and later: The account must be a member of the **Administrators** group on the machine, and a member of the **Organization Management** role group.

- For Active Directory:

The account must be a domain administrator.

Additional requirement for virtual machines

If the application runs on a virtual machine that is backed up by Agent for VMware or Agent for Hyper-V, ensure that User Account Control (UAC) is disabled on the machine. 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.

15.16.4 Mailbox backup

Mailbox backup is supported for Microsoft Exchange Server 2010 Service Pack 1 (SP1) and later.

Mailbox backup is available if at least one Agent for Exchange is registered on the management server. The agent must be installed on a machine that belongs to the same Active Directory forest as Microsoft Exchange Server.

Before backing up mailboxes, you must connect Agent for Exchange to the machine running the **Client Access** server role (CAS) of Microsoft Exchange Server. In Exchange 2016 and later, the CAS role is not available as a separate installation option. It is automatically installed as part of the Mailbox server role. Thus, you can connect the agent to any server running the **Mailbox role**.

To connect Agent for Exchange to CAS

1. Click **Devices > Add**.
2. Click **Microsoft Exchange Server**.
3. Click **Exchange mailboxes**.

If no Agent for Exchange is registered on the management server, the software suggests that you install the agent. After the installation, repeat this procedure from step 1.

4. [Optional] If multiple Agents for Exchange are registered on the management server, click **Agent**, and then change the agent that will perform the backup.
5. In **Client Access server**, specify the fully qualified domain name (FQDN) of the machine where the **Client Access** role of Microsoft Exchange Server is enabled.

In Exchange 2016 and later, the Client Access services are automatically installed as part of the Mailbox server role. Thus, you can specify any server running the **Mailbox role**. We refer to this server as CAS later in this section.

6. In **Authentication type**, select the authentication type that is used by the CAS. You can select **Kerberos** (default) or **Basic**.
7. [Only for basic authentication] Select which protocol will be used. You can select **HTTPS** (default) or **HTTP**.

8. [Only for basic authentication with the HTTPS protocol] If the CAS uses an SSL certificate that was obtained from a certification authority, and you want the software to check the certificate when connecting to the CAS, select the **Check SSL certificate** check box. Otherwise, skip this step.
9. Provide the credentials of an account that will be used to access the CAS. The requirements for this account are listed in "Required user rights" (p. 212).
10. Click **Add**.

As a result, the mailboxes appear under **Devices > Microsoft Exchange > Mailboxes**.

15.16.4.1 Selecting Exchange Server mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select Exchange mailboxes

1. Click **Devices > Microsoft Exchange**.
The software shows the tree of Exchange databases and mailboxes.
2. Click **Mailboxes**, and then select the mailboxes that you want to back up.
3. Click **Protect**.

15.16.4.2 Required user rights

To access mailboxes, Agent for Exchange needs an account with the appropriate rights. You are prompted to specify this account when configuring various operations with mailboxes.

Membership of the account in the **Organization Management** role group enables access to any mailbox, including mailboxes that will be created in the future.

The minimum required user rights are as follows:

- The account must be a member of the **Server Management** and **Recipient Management** role groups.
- The account must have the **ApplicationImpersonation** management role enabled for all users or groups of users whose mailboxes the agent will access.

For information about configuring the **ApplicationImpersonation** management role, refer to the following Microsoft knowledge base article:

<https://msdn.microsoft.com/en-us/library/office/dn722376.aspx>.

15.16.5 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 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. 215).

If you use only Agent for VMware (Windows), recovering databases as files is the only available recovery method. Recovering databases by using Agent for VMware (Virtual Appliance) is not possible.

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. 214).

To recover SQL databases to a SQL Server instance

1. Do one of the following:
 - When recovering from an application-aware backup, under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft SQL**, and then select the databases that you want to recover.
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 one of the following:

 - [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine that has Agent for SQL, and then select a recovery point.
 - Select a recovery point on the Backup storage tab (p. 199).

The machine chosen for browsing in either of the above actions becomes a target machine for the SQL databases recovery.
4. Do one of the following:
 - When recovering from an application-aware backup, click **Recover > SQL databases**, select the databases that you want to recover, and then click **Recover**.
 - When recovering from a database backup, click **Recover > Databases to an instance**.
5. 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 SQL Server instance (running on the same machine) 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.
6. [Optional] [Not available for a database recovered to its original instance as a new database] 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.

7. Click **Start recovery**.

The recovery progress is shown on the Activities tab.

To recover SQL databases as files

1. Do one of the following:

- When recovering from an application-aware backup, under **Devices**, select the machine that originally contained the data that you want to recover.
- When recovering from a database backup, click **Devices > Microsoft SQL**, and then select the databases that you want to recover.

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 one of the following:

- [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine that has Agent for SQL or Agent for VMware, and then select a recovery point.
- Select a recovery point on the Backup storage tab (p. 199).

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

4. Do one of the following:

- When recovering from an application-aware backup, click **Recover > SQL databases**, select the databases that you want to recover, and then click **Recover as files**.
- When recovering from a database backup, click **Recover > Databases as files**.

5. Click **Browse**, and then select a local or a network folder to save the files to.

6. Click **Start recovery**.

The recovery progress is shown on the Activities tab.

15.16.5.1 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.

15.16.5.2 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**.

15.16.6 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.
2010/2013/2016/2019	Databases	Membership in the Server 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. 217).

If you use only Agent for VMware (Windows), recovering databases as files is the only available recovery method. Recovering databases by using Agent for VMware (Virtual Appliance) is not possible.

We will refer to both databases and storage groups as "databases" throughout the below procedures.

To recover Exchange databases to a live Exchange Server

1. Do one of the following:
 - When recovering from an application-aware backup, under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the databases that you want to recover.

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 one of the following:

- [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), click **Select machine**, select an online machine that has Agent for Exchange, and then select a recovery point.
- Select a recovery point on the Backup storage tab (p. 199).

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

4. Do one of the following:
 - When recovering from an application-aware backup, click **Recover > Exchange databases**, select the databases that you want to recover, and then click **Recover**.
 - When recovering from a database backup, click **Recover > Databases to an Exchange server**.
5. 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.

6. Click **Start recovery**.

The recovery progress is shown on the Activities tab.

To recover Exchange databases as files

1. Do one of the following:
 - When recovering from an application-aware backup, under **Devices**, select the machine that originally contained the data that you want to recover.

- When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the databases that you want to recover.
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 one of the following:
 - [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), 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 Backup storage tab (p. 199).
 The machine chosen for browsing in either of the above actions becomes a target machine for the Exchange data recovery.
 4. Do one of the following:
 - When recovering from an application-aware backup, click **Recover > Exchange databases**, select the databases that you want to recover, and then click **Recover as files**.
 - When recovering from a database backup, click **Recover > Databases as files**.
 5. Click **Browse**, and then select a local or a network folder to save the files to.
 6. Click **Start recovery**.

The recovery progress is shown on the Activities tab.

15.16.6.1 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 2010 or later: <http://technet.microsoft.com/en-us/library/aa998871.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)

15.16.7 Recovering Exchange mailboxes and mailbox items

This section describes how to recover Exchange mailboxes and mailbox items from database backups, from application-aware backups, and from mailbox backups. The mailboxes or mailbox items can be recovered to a live Exchange Server or to Microsoft Office 365.

The following items can be recovered:

- Mailboxes (except for archive 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.

Recovery to an Exchange Server

Granular recovery can be performed to Microsoft Exchange Server 2010 Service Pack 1 (SP1) and later. The source backup may contain databases or mailboxes 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.

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. Instead, the full path to a mailbox item is recreated in the target folder.

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.

Recovery to Office 365

Recovery of Exchange data items to Office 365, and vice versa, is supported on the condition that Agent for Office 365 is installed locally.

Recovery can be performed from backups of Microsoft Exchange Server 2010 and later.

When a mailbox is recovered to an existing Office 365 mailbox, the existing items are kept intact, and the recovered items are placed next to them.

When recovering a single mailbox, you need to select the target Office 365 mailbox. When recovering several mailboxes within one recovery operation, the software will try to recover each mailbox to the mailbox of the user with the same name. If the user is not found, the mailbox is

skipped. If some mailboxes are skipped, the recovery will succeed with warnings. If all mailboxes are skipped, the recovery will fail.

For more information about recovery to Office 365, refer to "Protecting Office 365 mailboxes" (p. 230).

15.16.7.1 Recovering mailboxes

To recover mailboxes from an application-aware backup or a database backup

1. [Only when recovering from a database backup to Office 365] If Agent for Office 365 is not installed on the machine running Exchange Server that was backed up, do one of the following:
 - If there is not Agent for Office 365 in your organization, install Agent for Office 365 on the machine that was backed up (or on another machine with the same Microsoft Exchange Server version).
 - If you already have Agent for Office 365 in your organization, copy libraries from the machine that was backed up (or from another machine with the same Microsoft Exchange Server version) to the machine with Agent for Office 365, as described in "Copying Microsoft Exchange libraries" (p. 222).
2. Do one of the following:
 - When recovering from an application-aware backup: under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the database 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:

 - [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), 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 Backup storage tab (p. 199).

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. [Only when recovering to Office 365]:

- a. In **Recover to**, select **Microsoft Office 365**.
- b. [If you selected only one mailbox in step 6] In **Target mailbox**, specify the target mailbox.
- c. Click **Start recovery**.

Further steps of this procedure are not required.

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 a machine where the **Client Access** role (in Microsoft Exchange Server 2010/2013) or **Mailbox role** (in Microsoft Exchange Server 2016 or later) 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 will be used to access the machine. The requirements for this account are listed in "Required user rights" (p. 212).

10. [Optional] Click **Database to re-create any missing mailboxes** to change the automatically selected database.
11. Click **Start recovery**.

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

To recover a mailbox from a mailbox backup

1. Click **Devices > Microsoft Exchange > Mailboxes**.
2. Select the mailbox to recover, and then click **Recovery**.
You can search mailboxes by name. Wildcards are not supported.
If the mailbox was deleted, select it on the Backup storage tab (p. 199), and then click **Show backups**.
3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Mailbox**.
5. Perform steps 8-11 of the above procedure.

15.16.7.2 Recovering mailbox items

To recover mailbox items from an application-aware backup or a database backup

1. [Only when recovering from a database backup to Office 365] If Agent for Office 365 is not installed on the machine running Exchange Server that was backed up, do one of the following:
 - If there is not Agent for Office 365 in your organization, install Agent for Office 365 on the machine that was backed up (or on another machine with the same Microsoft Exchange Server version).
 - If you already have Agent for Office 365 in your organization, copy libraries from the machine that was backed up (or from another machine with the same Microsoft Exchange Server version) to the machine with Agent for Office 365, as described in "Copying Microsoft Exchange libraries" (p. 222).
2. Do one of the following:
 - When recovering from an application-aware backup: under **Devices**, select the machine that originally contained the data that you want to recover.
 - When recovering from a database backup, click **Devices > Microsoft Exchange > Databases**, and then select the database 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:

- [Only when recovering from an application-aware backup] If the backup location is cloud or shared storage (i.e. other agents can access it), 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 Backup storage tab (p. 199).

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. To recover to Office 365, select **Microsoft Office 365** in **Recover to**.
To recover to an Exchange Server, keep the default **Microsoft Exchange** value in **Recover to**.
10. [Only when recovering to an Exchange Server] 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 a machine where the **Client Access** role (in Microsoft Exchange Server 2010/2013) or **Mailbox role** (in Microsoft Exchange Server 2016 or later) 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 will be used to access the machine. The requirements for this account are listed in "Required user rights" (p. 212).
11. 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.

12. [Only when recovering email messages] In **Target folder**, view or change the target folder in the target mailbox. By default, the **Recovered items** folder is selected. Due to Microsoft Exchange limitations, events, tasks, notes, and contacts are restored to their original location regardless of any different **Target folder** specified.
13. Click **Start recovery**.

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

To recover a mailbox item from a mailbox backup

1. Click **Devices > Microsoft Exchange > Mailboxes**.
2. Select the mailbox that originally contained the items that you want to recover, and then click **Recovery**.

You can search mailboxes by name. Wildcards are not supported.

If the mailbox was deleted, select it on the Backup storage tab (p. 199), and then click **Show backups**.

3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Email messages**.
5. 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.

When an email message is selected, you can click **Send as email** to send the message to an email address. The message is sent from your administrator account's email address.

To be able to select folders, click the recover folders icon:



6. Click **Recover**.
7. Perform steps 9-13 of the above procedure.

15.16.7.3 Copying Microsoft Exchange Server libraries

When recovering Exchange mailboxes or mailbox items to Office 365 (p. 217), you may need to copy the following libraries from the machine that was backed up (or from another machine with the same Microsoft Exchange Server version) to the machine with Agent for Office 365.

Copy the following files, according to the Microsoft Exchange Server version that was backed up.

Microsoft Exchange Server version	Libraries	Default location
Microsoft Exchange Server 2010	ese.dll esebcli2.dll store.exe	%ProgramFiles%\Microsoft\Exchange Server\V14\bin

Microsoft Exchange Server 2013	ese.dll	%ProgramFiles%\Microsoft\Exchange Server\V15\bin
	msvcr110.dll	%WINDIR%\system32
Microsoft Exchange Server 2016, 2019	ese.dll	%ProgramFiles%\Microsoft\Exchange Server\V15\bin
	msvcr110.dll	%WINDIR%\system32
	msvcpr110.dll	

The libraries should be placed in the folder **%ProgramData%\Acronis\ese**. If this folder does not exist, create it manually.

15.16.8 Changing the SQL Server or Exchange Server access credentials

You can change access credentials for SQL Server or Exchange Server without re-installing the agent.

To change the SQL Server or Exchange Server access credentials

1. Click **Devices**, and then click **Microsoft SQL** or **Microsoft Exchange**.
2. Select the Always On Availability Group, Database Availability Group, SQL Server instance, or Exchange Server for which you want to change the access credentials.
3. Click **Specify credentials**.
4. Specify the new access credentials, and then click **OK**.

To change the Exchange Server access credentials for mailbox backup

1. Click **Devices > Microsoft Exchange**, and then expand **Mailboxes**.
2. Select the Exchange Server for which you want to change the access credentials.
3. Click **Settings**.
4. Under **Exchange administrator account**, specify the new access credentials, and then click **Save**.

15.17 Protecting mobile devices

The backup app allows you to back up your mobile data to the Cloud storage and then recover it in case of loss or corruption. Note that backup to the cloud storage requires an account and the Cloud subscription.

Supported mobile devices

You can install the backup app on a mobile device that runs one of the following operating systems:

- iOS 10.3 and later (iPhone, iPod, and iPads)
- Android 5.0 and later

What you can back up

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

What you need to know

- You can back up the data only to the cloud storage.
- Any time you open the app, you will see the summary of data changes and can start a backup manually.
- The **Continuous backup** functionality is enabled by default. If this setting is turned on:
 - For Android 7.0 or higher, the backup app automatically detects new data on-the-fly and uploads it to the Cloud,
 - For Android 5 and 6, it checks for changes every three hours. You can turn off continuous backup in the app settings.
- The **Use Wi-Fi only** option is enabled by default in the app settings. If this setting is turned on, the backup app will back up your data only when a Wi-Fi connection is available. If the Wi-Fi connection is lost, a backup process does not start. For the app to use cellular connection as well, turn this option off.
- You have two ways to save energy:
 - The **Back up while charging** functionality which is disabled by default. If this setting is turned on, the backup app will back up your data only when your device is connected to a power source. When the device is disconnected from a power source during a continuous backup process, the backup is paused.
 - The **Save power mode** which is enabled by default. If this setting is turned on, the backup app will back up your data only when your device battery is not low. When the device battery gets low, the continuous backup is paused. This option is available for Android 8 or higher.
- 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 any device by using the service console.
- The data backed up from mobile devices registered under your account is available only under this account. Nobody else can view or recover your data.
- In the backup app, you can recover only the latest data versions. If you need to recover from a specific backup version, use the service console on either a tablet or a computer.
- Retention rules are not applied to backups of mobile devices.
- [Only for Android devices] If an SD card is present during a backup, the data stored on this card is also backed up. The data will be recovered to an SD card, to the folder **Recovered by Backup** if it is present during recovery, or the app will ask for a different location to recover the data to.

Where to get the backup app

Depending on your mobile device, install the app from the App Store or Google Play.

How to start backing up your data

1. Open the app.
2. Sign in with your account.
3. Tap **Set up** to create your backup. Note that this button occurs only when you have no backup of your mobile device.
4. Select the data categories that you want to back up. By default, all categories are selected.
5. [optional step] Enable **Encrypt Backup** to protect your backup by encryption. In this case, you will need to also:
 1. Enter an encryption password twice.

Make sure you remember the password, because a forgotten password can never be restored or changed.

2. Tap **Encrypt**.
6. Tap **Back up**.
7. Allow the app access to your personal data. If you deny access to some data categories, they will not be backed up.

The backup starts.


How to recover data to a mobile device

1. Open the backup app.
2. Tap **Browse**.
3. Tap the device name.
4. Do one of the following:
 - To recover all of the backed-up data, tap **Recover all**. No more actions are required.
 - To recover one or more data categories, tap **Select**, and then tap the check boxes for the required data categories. Tap **Recover**. No more actions are required.
 - To recover one or more data items belonging to the same data category, tap the data category. Proceed to further steps.
5. Do one of the following:
 - To recover a single data item, tap it.
 - To recover several data items, tap **Select**, and then tap the check boxes for the required data items.
6. Tap **Recover**.


How to review data via the service console

1. On a computer, open a browser and type the service console URL.
2. Sign in with your account.
3. In **All devices**, click **Recover** under your mobile device name.
4. Do any of the following:

- To download all photos, videos, contacts, calendars, or reminders, select the respective data category. Click **Download**.



- To download individual photos, videos, contacts, calendars, or reminders, click the respective data category name, and then select the check boxes for the required data items. Click **Download**.



- To preview a photo, or a contact, click the respective data category name, and then click the required data item.

15.18 Protecting Hosted Exchange data

What items can be backed up?

You can back up user mailboxes, shared mailboxes, and group mailboxes. Optionally, you can choose to back up the archive mailboxes (**In-Place Archive**) of the selected mailboxes.

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders
- Email messages
- Calendar events
- Tasks
- Contacts
- Journal entries
- Notes

You can use search to locate the items.

When recovering mailboxes, mailbox items, public folders, and public folder items, you can select whether to overwrite the items in the target location.

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. Instead, the full path to a mailbox item is recreated in the target folder.

15.18.1 Selecting mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select Exchange Online mailboxes

1. Click **Devices > Hosted Exchange**.
2. If multiple Hosted Exchange organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the mailboxes of all users and all shared mailboxes (including mailboxes that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up individual user or shared mailboxes, expand the **Users** node, select **All users**, select the users whose mailboxes you want to back up, and then click **Backup**.
 - To back up all group mailboxes (including mailboxes of groups that will be created in the future), expand the **Groups** node, select **All groups**, and then click **Group backup**.

- To back up individual group mailboxes, expand the **Groups** node, select **All groups**, select the groups whose mailboxes you want to back up, and then click **Backup**.

15.18.2 Recovering mailboxes and mailbox items

15.18.2.1 Recovering mailboxes

1. Click **Devices > Hosted Exchange**.
2. If multiple Hosted Exchange organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox you want to recover, and then click **Recovery**.
 - To recover a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that you want to recover, and then click **Recovery**.
 - To recover a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Entire mailbox**.
6. If multiple Hosted Exchange organizations are added to the Cyber Protection service, click **Hosted Exchange organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

7. In **Recover to 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 organization is selected, you must specify the target mailbox.
8. Click **Start recovery**.
9. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
10. Click **Proceed** to confirm your decision.

15.18.2.2 Recovering mailbox items

1. Click **Devices > Hosted Exchange**.
2. If multiple Hosted Exchange organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:

- To recover items from a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
- To recover items from a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that originally contained the items that you want to recover, and then click **Recovery**.
- To recover items from a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
- If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Email messages**.
6. Browse to the required folder or use search to obtain the list of the required items.

The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, attachment name, 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.

7. Select the items that you want to recover. To be able to select folders, click the "recover folders" icon:



Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
- When an email message or a calendar item is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
- Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. Click **Recover**.
9. If multiple Hosted Exchange organizations were added to the Cyber Protection service, click **Hosted Exchange organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

10. In **Recover to 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 organization is selected, you must specify the target mailbox.

11. [Only when recovering to a user or a shared mailbox] In **Path**, view or change the target folder in the target mailbox. By default, the **Recovered items** folder is selected.

Group mailbox items are always recovered to the **Inbox** folder.

12. Click **Start recovery**.
13. Select one of the overwriting options:

- **Overwrite existing items**
- **Do not overwrite existing items**

14. Click **Proceed** to confirm your decision.

15.19 Protecting Office 365 data

Why back up Office 365 data?

Even though Microsoft Office 365 is a set of cloud services, regular backups provide an additional layer of protection from user errors and intentional malicious actions. You can recover deleted items from a backup even after the Office 365 retention period has expired. Also, you can keep a local copy of the Exchange Online mailboxes if it is required for regulatory compliance.

Agent for Office 365

Depending on the desired functionality, you can choose to install Agent for Office 365 locally, use the agent installed in the cloud, or both. The following table summarizes the functionality of the local and the cloud agent.

	Local Agent for Office 365	Cloud Agent for Office 365
Data items that can be backed up	Exchange Online: user and shared mailboxes	▪ Exchange Online: user, shared, and group mailboxes; public folders ▪ OneDrive: user files and folders ▪ SharePoint Online: classic site collections, group (team) sites, communication sites, individual data items
Backup of archive mailboxes (In-Place Archive)	No	Yes
Backup schedule	User-defined (p. 127)	Cannot be changed. Each protection plan runs daily at the same time of day.*
Backup locations	Cloud storage, local folder, network folder	Cloud storage only
Automatic protection of new Office 365 users, groups, sites	No	Yes, by applying a protection plan to the All users, All groups, All sites groups
Protecting more than one Office 365 organization	No	Yes
Granular recovery	Yes	Yes
Recovery to another user within one organization	Yes	Yes
Recovery to another organization	No	Yes

	Local Agent for Office 365	Cloud Agent for Office 365
Recovery to an on-premises Microsoft Exchange Server	No	No
Maximum number of items that can be backed up without performance degradation	When backing up to the cloud storage: 5000 mailboxes per company When backing up to other destinations: 2000 mailboxes per protection plan (no limitation for number of mailboxes per company)	10 000 protected items (mailboxes, OneDrives, or sites) per company**
Maximum number of manual backup runs	No	10 manual runs during an hour (p. 378)
Maximum number of simultaneous recovery operations	No	10 operations, including GSuite recovery operations

* Because a cloud agent serves multiple customers, it determines the start time for each protection plan on its own, to ensure even load during a day and the equal quality of service for all customers.

Note The protection schedule might be affected by the operation of third-party services, for example, the accessibility of Microsoft Office 365 servers, throttling settings on the Microsoft servers, and others. See also <https://docs.microsoft.com/en-us/graph/throttling> <https://docs.microsoft.com/en-us/graph/throttling>.

** It is recommended that you back up your protected items gradually and in this order:

1. Mailboxes.
2. After all mailboxes are backed up, proceed with OneDrives.
3. After OneDrive backup is completed, proceed with the SharePoint Online sites.

The first full backup may take several days, depending on the number of protected items and their size.

Limitations

- A mailbox backup includes only folders visible to users. The **Recoverable items** folder and its subfolders (**Deletions**, **Versions**, **Purges**, **Audits**, **DiscoveryHold**, **Calendar Logging**) are not included in a mailbox backup.
- Automatic creation of users, public folders, groups, or sites during a recovery is not possible. For example, if you want to recover a deleted SharePoint Online site, first create a new site manually, and then specify it as the target site during a recovery.

Required user rights

In the Cyber Protection service

Any Agent for Office 365, either local or cloud, must be registered under a company administrator account and used on a customer tenant level. Company administrators acting on a unit level, unit administrators, and users cannot back up or recover Office 365 data.

In Microsoft Office 365

Your account must be assigned the global administrator role in Microsoft Office 365.

To back up and recover Office 365 public folders, at least one of your Office 365 administrator accounts must have a mailbox and read/write rights to the public folders that you want to back up.

- The local agent will log in to Office 365 by using this account. To enable the agent to access the contents of all mailboxes, this account will be assigned the **ApplicationImpersonation** management role. If you change this account password, update the password in the service console, as described in "Changing the Office 365 access credentials" (p. 234).
- The cloud agent does not log in to Office 365. The agent is given the necessary permissions directly by Microsoft Office 365. You only need to confirm granting these permissions once, being signed in as a global administrator. The agent does not store your account credentials and does not use them to perform backup and recovery. Changing this account password or disabling this account or deleting this account in Office 365 does not affect agent operation.

15.19.1 Using the locally installed Agent for Office 365

15.19.1.1 Adding a Microsoft Office 365 organization

To add a Microsoft Office 365 organization

1. Sign in to the service console as a company administrator.
2. Click the account icon in the top-right corner, and then click **Downloads > Agent for Office 365**.
3. Download the agent and install it on a Windows machine that is connected to the Internet.
4. After the installation is complete, click **Devices > Microsoft Office 365**, and then enter the Office 365 global administrator credentials.

Important *There must be only one locally installed Agent for Office 365 in an organization (company group).*

As a result, your organization data items appear in the service console on the **Microsoft Office 365** page.

15.19.1.2 Protecting Exchange Online mailboxes

What items can be backed up?

You can back up user mailboxes and shared mailboxes. Group mailboxes and archive mailboxes (**In-Place Archive**) cannot be backed up.

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- 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. Instead, the full path to a mailbox item is recreated in the target folder.

Selecting mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select mailboxes

1. Click **Microsoft Office 365**.
2. If prompted, sign in as a global administrator to Microsoft Office 365.
3. Select the mailboxes that you want to back up.
4. Click **Backup**.

Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **Microsoft Office 365**.
2. Select the mailbox to recover, and then click **Recovery**.
You can search mailboxes by name. Wildcards are not supported.
If the mailbox was deleted, select it on the Backup storage tab (p. 199), and then click **Show backups**.
3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Mailbox**.
5. In **Target mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist, you must specify the target mailbox.
6. Click **Start recovery**.

Recovering mailbox items

1. Click **Microsoft Office 365**.
2. Select the mailbox that originally contained the items that you want to recover, and then click **Recovery**.
You can search mailboxes by name. Wildcards are not supported.
If the mailbox was deleted, select it on the Backup storage tab (p. 199), and then click **Show backups**.
3. Select a recovery point. Note that recovery points are filtered by location.
4. Click **Recover > Email messages**.
5. 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, attachment name, 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.

When an email message is selected, you can click **Send as email** to send the message to an email address. The message is sent from your administrator account's email address.

To be able to select folders, click the "recover folders" icon:



6. Click **Recover**.
7. In **Target mailbox**, view, change, or specify the target mailbox.
By default, the original mailbox is selected. If this mailbox does not exist, you must specify the target mailbox.
8. Click **Start recovery**.
9. Confirm your decision.

The mailbox items are always recovered to the **Recovered items** folder of the target mailbox.

Changing the Office 365 access credentials

You can change access credentials for Office 365 without re-installing the agent.

To change the Office 365 access credentials

1. Click **Devices > Microsoft Office 365**.
2. Click **Specify credentials**.
3. Enter the Office 365 global administrator credentials, and then click **OK**.
The agent will log in to Office 365 by using this account. To enable the agent to access the contents of all mailboxes, this account will be assigned the **ApplicationImpersonation** management role.

15.19.2 Using the cloud Agent for Office 365

15.19.2.1 Adding a Microsoft Office 365 organization

To add a Microsoft Office 365 organization

1. Sign in to the service console as a company administrator.
2. Click **Devices > Add > Microsoft Office 365 for Business**.
3. Select the Microsoft data center used by your organization.
The software redirects you to the Microsoft Office 365 login page.
4. Sign in with the Office 365 global administrator credentials.
Microsoft Office 365 displays a list of permissions that are necessary to back up and recover your organization's data.
5. Confirm that you grant the Cyber Protection service these permissions.

As a result, your organization's data items appear in the service console on the **Microsoft Office 365** page.

Tips for further usage

- The cloud agent synchronizes with Office 365 every 24 hours, starting from the moment when the organization is added to the Cyber Protection service. If you add or remove a user, group, or site, you will not see this change in the service console immediately. To forcibly synchronize the

cloud agent with Office 365, select the organization on the **Microsoft Office 365** page, and then click **Refresh**.

- If you applied a protection plan to the **All users**, **All groups**, or **All sites** group, the newly added items will be included in the backup only after the synchronization.
- According to Microsoft policy, after a user, group, or site is removed from Office 365 GUI, it remains available for a few days via the API. During these days, the removed item is inactive (grayed out) in the service console and is not backed up. When the removed item becomes unavailable via the API, it disappears from the service console. Its backups (if any) can be found at **Backups > Cloud applications backups**.

15.19.2.2 Deleting a Microsoft Office 365 organization

To delete a Microsoft Office 365 organization

1. Sign in to the service console as a company administrator.
2. Go to **Devices > Microsoft Office 365**.
3. Select the organization and click **Delete group**.

As a result, the backup plans applied to this group will be revoked.

However, you should additionally revoke access rights of the Backup Service application to Office 365 organization data manually.

To revoke access rights

1. Log in to Office 365 under a global administrator.
2. Go to **Admin Center > Azure Active Directory > Enterprise applications > All applications**.
3. Select the **Backup Service** application and drill down to it.
4. Go to the **Properties** tab and on the action panel click **Delete**.
5. Confirm the deletion operation.

As a result, access rights to the Office 365 organization data will be revoked from the Backup Service application.

15.19.2.3 Protecting Exchange Online data

What items can be backed up?

You can back up user mailboxes, shared mailboxes, and group mailboxes. Optionally, you can choose to back up the archive mailboxes (**In-Place Archive**) of the selected mailboxes.

Starting from version 8.0 of the Cyber Protection service, you can back up public folders. If your organization was added to the Cyber Protection service before the version 8.0 release, you need to re-add the organization to obtain this functionality. Do not delete the organization, simply repeat the steps described in "Adding a Microsoft Office 365 organization" (p. 234). As a result, the Cyber Protection service obtains the permission to use the corresponding API.

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders
- Email messages
- Calendar events

- Tasks
- Contacts
- Journal entries
- Notes

The following items can be recovered from a public folder backup:

- Subfolders
- Posts
- Email messages

You can use search to locate the items.

When recovering mailboxes, mailbox items, public folders, and public folder items, you can select whether to overwrite the items in the target location.

Selecting mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select Exchange Online mailboxes

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the mailboxes of all users and all shared mailboxes (including mailboxes that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up individual user or shared mailboxes, expand the **Users** node, select **All users**, select the users whose mailboxes you want to back up, and then click **Backup**.
 - To back up all group mailboxes (including mailboxes of groups that will be created in the future), expand the **Groups** node, select **All groups**, and then click **Group backup**.
 - To back up individual group mailboxes, expand the **Groups** node, select **All groups**, select the groups whose mailboxes you want to back up, and then click **Backup**.

Note The cloud Agent for Office 365 uses an account with the appropriate rights to access a group mailbox. Thus, to back up a group mailbox, at least one of the group owners must be licensed Office 365 user with a mailbox. If the group is private or with hidden membership, the owner must also be a member of the group.

4. On the protection plan panel:
 - Ensure that the **Office 365 mailboxes** item is selected in **What to back up**.
 - If you do not want to backup the archive mailboxes, disable the **Archive mailbox** switch.

Selecting public folders

Select the public folders as described below, and then specify other settings of the protection plan as appropriate (p. 111).

Note Public folders consume licenses from your backup quota for Office 365 seats.

To select Exchange Online public folders

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, expand the organization whose data you want to back up. Otherwise, skip this step.
3. Expand the **Public folders** node, and then select **All public folders**.
4. Do one of the following:
 - To back up all public folders (including public folders that will be created in the future), click **Group backup**.
 - To back up individual public folders, select the public folders that you want to back up, and then click **Backup**.
5. On the protection plan panel, ensure that the **Office 365 mailboxes** item is selected in **What to back up**.

Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox you want to recover, and then click **Recovery**.
 - To recover a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that you want to recover, and then click **Recovery**.
 - To recover a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.

Tip To see only the recovery points that contain mailboxes, select **Mailboxes** in **Filter by content**.

5. Click **Recover > Entire mailbox**.
6. If multiple Office 365 organizations are added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to 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 organization is selected, you must specify the target mailbox.
8. Click **Start recovery**.
9. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
10. Click **Proceed** to confirm your decision.

Recovering mailbox items

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover items from a user mailbox, expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a shared mailbox, expand the **Users** node, select **All users**, select the shared mailbox that originally contained the items that you want to recover, and then click **Recovery**.
 - To recover items from a group mailbox, expand the **Groups** node, select **All groups**, select the group whose mailbox originally contained the items that you want to recover, and then click **Recovery**.
 - If the user, group, or the shared mailbox was deleted, select the item in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.

Tip. To see only the recovery points that contain mailboxes, select **Mailboxes** in **Filter by content**.

5. Click **Recover > Email messages**.
6. Browse to the required folder or use search to obtain the list of the required items.

The following search options are available. Wildcards are not supported.

 - For email messages: search by subject, sender, recipient, attachment name, 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.
7. Select the items that you want to recover. To be able to select folders, click the "recover folders" icon:



icon:

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
 - When an email message or a calendar item is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
 - Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. Click **Recover**.
 9. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
 10. In **Recover to 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 organization is selected, you must specify the target mailbox.

11. [Only when recovering to a user or a shared mailbox] In **Path**, view or change the target folder in the target mailbox. By default, the **Recovered items** folder is selected.

Group mailbox items are always recovered to the **Inbox** folder.

12. Click **Start recovery**.
13. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
14. Click **Proceed** to confirm your decision.

Recovering public folders and folder items

In order to recover a public folder or public folder items, at least one administrator of the target Office 365 organization must have the **Owner's** rights for the target public folder. If the recovery fails with an error about denied access, assign these rights in the target folder properties, select the target organization in the service console, click **Refresh**, and then repeat the recovery.

To recover a public folder or folder items

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations are added to the Cyber Protection service, expand the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - Expand the **Public folders** node, select **All public folders**, select the public folder that you want to recover or that originally contained the items that you want to recover, and then click **Recovery**.
 - If the public folder was deleted, select it in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search public folders by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover data**.
6. Browse to the required folder or use search to obtain the list of the required items.

You can search email messages and posts by subject, sender, recipient, and date. Wildcards are not supported.
7. Select the items that you want to recover. To be able to select folders, click the "recover folders"

icon: 

Additionally, you can do any of the following:

- When an email message or a post is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
 - When an email message or a post is selected, click **Send as email** to send the item to specified email addresses. You can select the sender and write a text to be added to the forwarded item.
 - Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. Click **Recover**.

9. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
10. In **Recover to public folder**, view, change, or specify the target public folder.
By default, the original folder is selected. If this folder does not exist or a non-original organization is selected, you must specify the target folder.
11. In **Path**, view or change the target subfolder in the target public folder. By default, the original path will be recreated.
12. Click **Start recovery**.
13. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
14. Click **Proceed** to confirm your decision.

15.19.2.4 Protecting OneDrive files

What items can be backed up?

You can back up an entire OneDrive, or individual files and folders.

Files are backed up together with their sharing permissions. Advanced permission levels (**Design**, **Full**, **Contribute**) are not backed up.

What items can be recovered?

You can recover an entire OneDrive or any file or folder that was backed up.

You can use search to locate the items.

You can choose whether to recover the sharing permissions or let the files inherit the permissions from the folder to which they are recovered.

Sharing links for files and folders are not recovered.

Selecting OneDrive files

Select the files as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select OneDrive files

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the files of all users (including users that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up the files of individual users, expand the **Users** node, select **All users**, select the users whose files you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **OneDrive** item is selected in **What to back up**.

- In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all files).
 - Specify the files and folders to back up by adding their names or paths.
You can use wildcard characters (*, **, and ?). For more details about specifying paths and using wildcards, refer to "File filters" (p. 153).
 - Specify the files and folders to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single user.
- [Optional] In **Items to back up**, click **Show exclusions** to specify the files and folders to skip during the backup.
File exclusions override the file selection; i.e. if you specify the same file in both fields, this file will be skipped during a backup.

Recovering OneDrive and OneDrive files

Recovering an entire OneDrive

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose OneDrive you want to recover, and then click **Recovery**.
If the user was deleted, select the user in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.
You can search users by name. Wildcards are not supported.
4. Select a recovery point.

***Tip.** To see only the recovery points that contain OneDrive files, select **OneDrive** in **Filter by content**.*

5. Click **Recover > Entire OneDrive**.
6. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to drive**, view, change, or specify the target user.
By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user.
8. Select whether to recover the sharing permissions for the files.
9. Click **Start recovery**.
10. Select one of the overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**
11. Click **Proceed** to confirm your decision.

Recovering OneDrive files

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.

3. Expand the **Users** node, select **All users**, select the user whose OneDrive files you want to recover, and then click **Recovery**.
If the user was deleted, select the user in the **Cloud Applications Backups** section of the Backup storage tab (p. 199), and then click **Show backups**.
You can search users by name. Wildcards are not supported.
4. Select a recovery point.

***Tip.** To see only the recovery points that contain OneDrive files, select **OneDrive** in **Filter by content**.*

5. Click **Recover > Files/folders**.
6. Browse to the required folder or use search to obtain the list of the required files and folders.
The search is not available if the backup is encrypted.
7. Select the files that you want to recover.
If the backup is not encrypted and you selected a single file, you can click **Show versions** to select the file version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. If you want to download a file, select the file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
9. Click **Recover**.
10. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
11. In **Recover to drive**, view, change, or specify the target user.
By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user.
12. In **Path**, view or change the target folder in the target user's OneDrive. By default, the original location is selected.
13. Select whether to recover the sharing permissions for the files.
14. Click **Start recovery**.
15. Select one of the file overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**
16. Click **Proceed** to confirm your decision.

15.19.2.5 Protecting SharePoint Online sites

What items can be backed up?

You can back up SharePoint classic site collections, group (modern team) sites, and communication sites. Also, you can select individual subsites, lists, and libraries for backup.

The following items are *skipped* during a backup:

- The **Look and Feel** site settings (except for **Title, description, and logo**).
- Site page comments and page comments settings (comments **On/Off**).
- The **Site features** site settings.

- Web part pages and web parts embedded in the wiki pages (due to SharePoint Online API limitations).
- Checked out files—files that are manually checked out for editing and all files that are created or uploaded in libraries, for which the option **Require Check Out** was enabled. To backup these files, first check them in.
- OneNote files (due to SharePoint Online API limitations).
- External data and Managed Metadata types of columns.
- The default site collection "domain-my.sharepoint.com". This is a collection where all of the organization users' OneDrive files reside.
- The contents of the recycle bin.

Limitations

- Titles and descriptions of sites/subsites/lists/columns are truncated during a backup if the title/description size is greater than 10000 bytes.
- You cannot back up previous versions of files created in SharePoint Online. Only the latest versions of the files are protected.
- You cannot back up sites created in the Business Productivity Online Suite (BPOS), the predecessor of Microsoft Office 365.
- You cannot back up the settings for sites that use the managed path **/portals** (for example, <https://<tenant>.sharepoint.com/portals/...>).

What items can be recovered?

The following items can be recovered from a site backup:

- Entire site
- Subsites
- Lists
- List items
- Document libraries
- Documents
- List item attachments
- Site pages and wiki pages

You can use search to locate the items.

Items can be recovered to the original or a non-original site. The path to a recovered item is the same as the original one. If the path does not exist, it is created.

You can choose whether to recover the sharing permissions or let the items inherit the permissions from the parent object after the recovery.

What items cannot be recovered?

- Subsites based on the **Visio Process Repository** template.
- Lists of the following types: **Survey list**, **Task list**, **Picture library**, **Links**, **Calendar**, **Discussion Board**, **External**, and **Import Spreadsheet**.
- Lists for which multiple content types are enabled.

Selecting SharePoint Online data

Select the data as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select SharePoint Online data

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up all classic SharePoint sites in the organization, including sites that will be created in the future, expand the **Site collections** node, select **All site collections**, and then click **Group backup**.
 - To back up individual classic sites, expand the **Site collections** node, select **All site collections**, select the sites that you want to back up, and then click **Backup**.
 - To back up all group (modern team) sites, including sites that will be created in the future, expand the **Groups** node, select **All groups**, and then click **Group backup**.
 - To back up individual group (modern team) sites, expand the **Groups** node, select **All groups**, select the groups whose sites you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **SharePoint sites** item is selected in **What to back up**.
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all items of the selected sites).
 - Specify the subsites, lists, and libraries to back up by adding their names or paths.
To back up a subsite or a top-level site list/library, specify its display name in the following format: `/display name/**`
To back up a subsite list/library, specify its display name in the following format: `/subsite display name/list display name/**`
The display names of subsites, lists, and libraries are shown on the **Site contents** page of a SharePoint site or subsite.
 - Specify the subsites to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single site.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the subsites, lists, and libraries to skip during the backup.
Item exclusions override the item selection; i.e. if you specify the same subsite in both fields, this subsite will be skipped during a backup.

Recovering SharePoint Online data

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Do one of the following:
 - To recover data from a group (modern team) site, expand the **Groups** node, select **All groups**, select the group whose site originally contained the items that you want to recover, and then click **Recovery**.

- To recover data from a classic site, expand the **Site Collections** node, select **All site collections**, select the site that originally contained the items that you want to recover, and then click **Recovery**.
- If the site was deleted, select it in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search groups and sites by name. Wildcards are not supported.

4. Select a recovery point.

Tip. To see only the recovery points that contain SharePoint sites, select **SharePoint sites** in **Filter by content**.

5. Click **Recover SharePoint files**.

6. Browse to the required folder or use search to obtain the list of the required data items.

The search is not available if the backup is encrypted.

7. Select the items that you want to recover.

If the backup is not encrypted, you used search, and selected a single item in the search results, you can click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. If you want to download an item, select the item, click **Download**, select the location to save the item to, and click **Save**. Otherwise, skip this step.

9. Click **Recover**.

10. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

11. In **Recover to site**, view, change, or specify the target site.

By default, the original site is selected. If this site does not exist or a non-original organization is selected, you must specify the target site.

12. Select whether to recover the sharing permissions of the recovered items.

13. Click **Start recovery**.

14. Select one of the overwriting options:

- **Overwrite existing files**
- **Overwrite an existing file if it is older**
- **Do not overwrite existing files**

15. Click **Proceed** to confirm your decision.

15.19.2.6 Protecting Office 365 Teams

What items can be backed up?

You can back up entire teams. This includes team name, team members list, team channels and their content, team mailbox and meetings, and team site.

What items can be recovered?

- Entire team
- Team channels
- Channel files
- Team mailbox

- Email folders in the team mailbox
- Email messages in the team mailbox
- Meetings
- Team site

You cannot recover conversations in team channels, but you can download them as a single html file.

Limitations

The following items are not backed up:

- The settings of the general channel (moderation preferences) – due to a Microsoft Teams beta API limitation.
- The settings of the custom channels (moderation preferences) – due to a Microsoft Teams beta API limitation.
- Meeting notes, chats.
- Stickers and praises.

Backup and recovery are supported for the following channel tabs:

- Word
- Excel
- PowerPoint
- PDF
- Document Library

Files that are shared in private channels are backed up, but not restored due to an API limitation.

Note *These files are stored in specific locations, separately from the files that are shared in public channels.*

Selecting teams

Select teams as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select teams

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose teams you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up all the teams in the organization (including teams that will be created in the future), expand the **Teams** node, select **All teams**, and then click **Group backup**.
 - To back up individual teams, expand the **Teams** node, select **All teams**, select the teams that you want to back up, and then click **Backup**.

You can search teams by name. Wildcards are not supported.

4. On the protection plan panel:
 - Ensure that the **Microsoft Teams** item is selected in **What to back up**.
 - [Optional] In **How long to keep**, set the cleanup options.

- [Optional] If you want to encrypt your backup, enable the **Encryption** switch, and then set your password and select the encryption algorithm.

Recovering an entire team

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team that you want to recover, and then click **Recovery**.
You can search teams by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Entire Team**.
If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
6. In **Recover to team**, view, change, or specify the target team.
By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target team.
7. Click **Start recovery**.
8. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**
9. Click **Proceed** to confirm your decision.

When you delete a channel in Microsoft Teams' graphic interface, it is not immediately removed from the system. Thus, when you recover the whole team, this channel's name cannot be used and a postfix will be added to it.

Conversations are recovered as a single html file in the **Files** tab of the channel. You can find this file in a folder named according to the following pattern: <Team name>_<Channel name>_conversations_backup_<date of recovery>T<time of recovery>Z.

Note After recovering a team or team channels, go to Microsoft Teams, select the channels that were recovered, and then click their **Files** tab. Otherwise, the subsequent backups of these channels will not include this tab's content – due to a Microsoft Teams beta API limitation.

Recovering team channels or files in team channels

To recover team channels

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose channels you want to recover, and then click **Recovery**.
4. Select a recovery point.
5. Click **Recover > Channels**.

6. Select the channels that you want to recover, and then click **Recover**. To select a channel in the main pane, select the check box in front of its name.
The following search options are available:
 - For **Conversations**: sender, subject, content, language, attachment name, date or date range.
 - For **Files**: file name or folder name, file type, size, date or date range of the last change.
7. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
8. In **Recover to team**, view, change, or specify the target team.
By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target team.
9. In **Recover to channel**, view, change, or specify the target channel.
10. Click **Start recovery**.
11. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**
12. Click **Proceed** to confirm your decision.

Conversations are recovered as a single html file in the **Files** tab of the channel. You can find this file in a folder named according to the following pattern: <Team name>_<Channel name>_conversations_backup_<date of recovery>T<time of recovery>Z.

Note After recovering a team or team channels, go to Microsoft Teams, select the channels that were recovered, and then click their **Files** tab. Otherwise, the subsequent backups of these channels will not include this tab's content – due to a Microsoft Teams beta API limitation.

To recover files in a team channel

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose channels you want to recover, and then click **Recovery**.
4. Select a recovery point.
5. Click **Recover > Channels**.
6. Select the desired channel, and then open the **Files** folder.
Browse to the required items or use search to obtain the list of the required items. The following search options are available: file name or folder name, file type, size, date or date range of the last change.
7. Select the items that you want to recover, and then click **Recover**.
8. If multiple Office 365 organizations were added to the Cyber Protection service, click Office 365 organization to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
9. In **Recover to team**, view, change, or specify the target team.

By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target team.

10. In **Recover to channel**, view, change, or specify the target channel.
11. Select whether to recover the sharing permissions of the recovered items.
12. Click **Start recovery**.
13. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**
14. Click **Proceed** to confirm your decision.

You cannot recover individual conversations. In the main pane, you can only browse the **Conversation** folder or download its content as a single html file. To do so, click the "recover folders" icon , select the desired **Conversations** folder, and then click **Download**.

You can search the messages in the **Conversation** folder by:

- Sender
- Content
- Attachment name
- Date

Recovering team mailboxes

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose mailbox you want to recover, and then click **Recovery**.

You can search teams by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Email messages**.
6. Click the "recover folders" icon , select the root mailbox folder, and then click **Recover**.

Note You can also recover individual folders from the selected mailbox.

7. Click **Recover**.
8. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

9. In **Recover to 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 organization is selected, you must specify the target mailbox.

10. Click **Start recovery**.
11. Select one of the overwriting options:

- **Overwrite existing items**
- **Do not overwrite existing items**

12. Click **Proceed** to confirm your decision.

Recovering email messages and meetings

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose email messages or meetings you want to recover, and then click **Recovery**.

You can search teams by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Email messages**.
6. Browse to the required item or use search to obtain the list of the required items.

The following search options are available:

- For email messages: search by subject, sender, recipient, and date.
- For meetings: search by event name and date.

7. Select the items that you want to recover, and then click **Recover**.

Note You can find the meetings in the **Calendar** folder.

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
 - When an email message or a meeting is selected, click **Send as email** to send the item to the specified email addresses. You can select the sender and write a text to be added to the forwarded item.
8. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
 9. In **Recover to 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 organization is selected, you must specify the target mailbox.
 10. Click **Start recovery**.
 11. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
 12. Click **Proceed** to confirm your decision.

Recovering a team site or specific items of a site

1. Click **Microsoft Office 365**.
2. If multiple Office 365 organizations were added to the Cyber Protection service, select the organization whose backed-up teams you want to recover. Otherwise, skip this step.
3. Expand the **Teams** node, select **All teams**, select the team whose site you want to recover, and then click **Recovery**.

You can search teams by name. Wildcards are not supported.

4. Select a recovery point.
5. Click **Recover > Team site**.
6. Browse to the required item or use search to obtain the list of the required items.
The search is not available if the backup is encrypted.
7. Select the items that you want to recover, and then click **Recover**.
8. If multiple Office 365 organizations were added to the Cyber Protection service, click **Office 365 organization** to view, change, or specify the target organization.
By default, the original organization and team are selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
9. In **Recover to team**, view, change, or specify the target team.
By default, the original team is selected. If this team does not exist or a non-original organization is selected, you must specify the target site.
10. Select whether to recover the sharing permissions of the recovered items.
11. Click **Start recovery**.
12. Select one of the overwriting options:
 - **Overwrite existing content if it is older**
 - **Overwrite existing content**
 - **Do not overwrite existing content**
13. Click **Proceed** to confirm your decision.

15.19.2.7 Upgrading the cloud agent

This section describes how to upgrade to the current version of the backup solution for Microsoft Office 365. This version supports OneDrive and SharePoint Online backup, and provides improved backup and recovery performance. Starting from version 8.0 of the Cyber Protection service, the following functionality is no longer supported by the old solution: editing, deleting, applying, and revoking a protection plan.

The upgrade availability depends on the data center readiness and the settings made by your service provider. If the upgrade is available, the service console shows a notification at the top of the **Microsoft Office 365 (v1)** tab.

The upgrade process

During the upgrade, your Office 365 organization users are added to the new backup solution. The protection plans are migrated and applied to the appropriate users.

The earlier created backups are copied from one location in the cloud to another. On the **Backup storage** tab, the copied backups are shown in a separate section named **Cloud applications backups**, while the original backups remain in the **Cloud storage** location. When the upgrade process is complete, the original backups are deleted from the **Cloud storage** location.

The upgrade may take several hours, or even days, depending on the number of users in the organization, the number of backups, and the Office 365 access speed. During the upgrade, recovery from the earlier created backups is possible. However, backups and protection plans created during the upgrade will be lost.

In the unlikely case of an upgrade failure, the backup solution remains fully operational and the upgrade can be restarted from the point of failure.

To start the upgrade process

1. Click **Microsoft Office 365 (v1)**.
2. Click **Upgrade** in the notification at the top of the screen.
3. Confirm that you want to start the upgrade process.
4. Select the Microsoft data center used by your organization.
The software redirects you to the Microsoft Office 365 login page.
5. Sign in with the Office 365 global administrator credentials.
Microsoft Office 365 displays a list of permissions that are necessary to back up and recover your organization's data.
6. Confirm that you grant the Cyber Protection service these permissions.
You are redirected to the service console and the upgrade process begins. The upgrade progress is shown on the **Microsoft Office 365 > Activities** panel.

15.20 Protecting G Suite data

What does G Suite protection mean?

- Cloud-to-cloud backup and recovery of G Suite user data (Gmail mailboxes, Calendars, Contacts, Google Drives) and G Suite Shared drives.
- Granular recovery of emails, files, contacts, and other items.
- Support for several G Suite organizations and cross-organization recovery.
- Optional notarization of the backed-up files by means of the Ethereum blockchain database. When enabled, you can prove that a file is authentic and unchanged since it was backed up.
- Optional full-text search. When enabled, you can search emails by their content.
- Up to 5000 items (mailboxes, Google Drives, and Shared drives) per company can be protected without performance degradation.

Supported G Suite editions

- G Suite Basic. Only Gmail, Drive, Calendar, and Contacts services.
- G Suite Business. Only Gmail, Drive (including Shared drives), Calendar, and Contacts services.
- G Suite Enterprise. Only Gmail, Drive (including Shared drives), Calendar, and Contacts services.
- G Suite for Education. Only Gmail, Drive (including Shared drives), Calendar, and Contacts services. Classroom service is not supported.

Required user rights

In the Cyber Protection service

In the Cyber Protection service, you need to be a company administrator acting on a customer tenant level. Company administrators acting on a unit level, unit administrators, and users cannot back up or recover G Suite data.

In G Suite

To add your G Suite organization to the Cyber Protection service, you must be signed in as a Super Admin with enabled API access (**Security > API reference > Enable API access** in the Google Admin console).

The Super Admin password is not stored anywhere and is not used to perform backup and recovery. Changing this password in G Suite does not affect Cyber Protection service operation.

If the Super Admin who added the G Suite organization is deleted from G Suite or assigned a role with less privileges, the backups will fail with an error like "access denied". In this case, repeat the "Adding a G Suite organization" (p. 253) procedure and specify valid Super Admin credentials. To avoid this situation, we recommend creating a dedicated Super Admin user for backup and recovery purposes.

About the backup schedule

Because the cloud agent serves multiple customers, it determines the start time for each protection plan on its own, to ensure an even load during a day and an equal quality of service for all of the customers.

Each protection plan runs daily at the same time of day.

Limitations

- Search in encrypted backups is not supported.
- No more than 10 manual backup runs during an hour (p. 378).
- No more than 10 simultaneous recovery operations (this number includes both Office 365 and GSuite recovery).

15.20.1 Adding a G Suite organization

To add a G Suite organization

1. Sign in to the service console as a company administrator.
2. Click **Devices > Add > G Suite**.
3. Follow the instructions displayed by the software:
 - a. Click **Open marketplace**.
 - b. Sign in with the Super Admin credentials.
 - c. Click **Domain install**.
 - d. Confirm the domain-wide installation.

G Suite displays a list of permissions that are necessary to back up and recover your organization's data.
 - e. Confirm that you grant the Cyber Protection service these permissions.
 - f. Complete the installation wizard.
 - g. Go to the Apps launcher icon, find the Cyber Protection Service application in the list, and then click on it.

You are redirected back to the service console. Your organization's data items appear in the service console on the **G Suite** page.

Tips for further usage

- After adding a G Suite organization, the user data and Shared drives in both the primary domain and all the secondary domains, if there are any, will be backed up. The backed-up resources will be displayed in one list, and will not be grouped by their domain.
- The cloud agent synchronizes with G Suite every 24 hours, starting from the moment when the organization is added to the Cyber Protection service. If you add or remove a user or Shared

drive, you will not see this change in the service console immediately. To forcibly synchronize the cloud agent with G Suite, select the organization on the **G Suite** page, and then click **Refresh**.

- If you applied a protection plan to the **All users** or **All Shared drives** group, the newly added items will be included in the backup only after the synchronization.
- According to Google policy, after a user or Shared drive is removed from the G Suite GUI, it remains available for a few days via the API. During these days, the removed item is inactive (grayed out) in the service console and is not backed up. When the removed item becomes unavailable via the API, it disappears from the service console. Its backups (if any) can be found at **Backups > Cloud Applications Backups**.

15.20.2 Protecting Gmail data

What items can be backed up?

You can back up Gmail users' mailboxes. A mailbox backup also includes the Calendar and Contacts data. Optionally, you can choose to back up the shared calendars.

The following items are *skipped* during a backup:

- The **Birthdays, Reminders, Tasks** calendars
- Folders attached to calendar events
- The **Directory** folder in Contacts

The following Calendar items are *skipped*, due to Google Calendar API limitations:

- Appointment slots
- The conferencing field of an event
- The calendar setting **All-day event notifications**
- The calendar setting **Auto-accept invitations** (in calendars for rooms or shared spaces)

The following Contacts items are *skipped*, due to Google People API limitations:

- The **Other contacts** folder
- The external profiles of a contact (**Directory profile, Google profile**)
- The contact field **File as**

What items can be recovered?

The following items can be recovered from a mailbox backup:

- Mailboxes
- Email folders (According to Google terminology, "labels". **Labels** are presented in the backup software as folders, for consistency with other data presentation.)
- Email messages
- Calendar events
- Contacts

You can use search to locate items in a backup, unless the backup is encrypted. Search in encrypted backups is not supported.

When recovering mailboxes and mailbox items, you can select whether to overwrite the items in the target location.

Limitations

- Contact photos cannot be recovered
- The **Out of office** calendar item is recovered as a regular calendar event, due to Google Calendar API limitations

15.20.2.1 Selecting mailboxes

Select the mailboxes as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select Gmail mailboxes

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the mailboxes of all users (including mailboxes that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up individual user mailboxes, expand the **Users** node, select **All users**, select the users whose mailboxes you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **Gmail** item is selected in **What to back up**.
 - If you want to back up calendars that are shared with the selected users, enable the **Include shared calendars** switch.
 - Decide whether you need full-text search (p. 255) through the backed-up email messages. To access this option, click the gear icon > **Backup options** > **Full-text search**.

Full-text search

This option defines whether the email messages content is indexed by the cloud agent.

The preset is: **Enabled**.

If this option is enabled, the messages content is indexed and you can search messages by their content. Otherwise, only searching by subject, sender, recipient, or date is available.

Note Search in encrypted backups is not supported.

The indexing process does not affect the backup performance because it is performed by a different software component. Indexing of the first (full) backup may take some time, therefore, there may be a delay between the backup completion and the content appearing in the search results.

The index occupies 10-30 percent of storage space occupied by the mailbox backups. To learn the exact value, click **Backup storage** > **Cloud applications backups** and view the **Index size** column. You may want to disable full-text search in order to save this space. The value in the **Index size** column will decrease to a few megabytes after the next backup. This minimal amount of metadata is necessary to perform a search by subject, sender, recipient, or date.

When you re-enable full-text search, the software indexes all of the backups previously created by the protection plan. This also takes some time.

15.20.2.2 Recovering mailboxes and mailbox items

Recovering mailboxes

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose mailbox you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.

Tip. To see only the recovery points that contain mailboxes, select **Gmail** in **Filter by content**.

5. Click **Recover > Entire mailbox**.
6. If multiple G Suite organizations are added to the Cyber Protection service, click **G Suite organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to 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 organization is selected, you must specify the target mailbox.
8. Click **Start recovery**.
9. Select one of the overwriting options:
 - **Overwrite existing items**
 - **Do not overwrite existing items**
10. Click **Proceed** to confirm your decision.

Recovering mailbox items

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose mailbox originally contained the items that you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users and groups by name. Wildcards are not supported.

4. Select a recovery point.

Tip. To see only the recovery points that contain mailboxes, select **Gmail** in **Filter by content**.

5. Click **Recover > Email messages**.
6. Browse to the required folder. If the backup is not encrypted, you can use search to obtain the list of the required items.

The following search options are available. Wildcards are not supported.

- For email messages: search by subject, sender, recipient, date, attachment name, and message content. The last two options yield results only if the **Full-text search** option was

enabled during backup. The language of the message fragment being searched can be specified as an additional parameter.

- For events: search by title and date.
- For contacts: search by name, email address, and phone number.

7. Select the items that you want to recover. To be able to select folders, click the "recover folders"

icon: 

Additionally, you can do any of the following:

- When an item is selected, click **Show content** to view its contents, including attachments. Click the name of an attached file to download it.
- Only if the backup is not encrypted, you used search, and selected a single item in the search results: click **Show versions** to select the item version to recover. You can select any backed-up version, earlier or later than the selected recovery point.

8. Click **Recover**.

9. If multiple G Suite organizations were added to the Cyber Protection service, click **G suite organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

10. In **Recover to 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 organization is selected, you must specify the target mailbox.

11. In **Path**, view or change the target folder in the target mailbox. By default, the original folder is selected.

12. Click **Start recovery**.

13. Select one of the overwriting options:

- **Overwrite existing items**
- **Do not overwrite existing items**

14. Click **Proceed** to confirm your decision.

15.20.3 Protecting Google Drive files

What items can be backed up?

You can back up an entire Google Drive, or individual files and folders. Optionally, you can choose to back up files that are shared with the Google Drive user.

Files are backed up together with their sharing permissions.

The following items are *skipped* during a backup:

- A shared file, if the user has a commenter or viewer access to the file and the file owner disabled the options to download, print, and copy for commenters and viewers.
- The **Computers** folder (created by the Backup and Sync client)

Limitations

- Out of Google-specific file formats, only Google docs, Google sheets, Google slides, and Google Drawings are backed up.

What items can be recovered?

You can recover an entire Google Drive, or any file or folder that was backed up.

You can use search to locate items in a backup, unless the backup is encrypted. Search in encrypted backups is not supported.

You can choose whether to recover the sharing permissions or let the files inherit the permissions from the folder to which they are recovered.

Limitations

- Comments in files are not recovered.
- Sharing links for files and folders are not recovered.
- The read-only **Owner settings** for shared files (**Prevent editors from changing access and adding new people** and **Disable options to download, print and copy for commenters and viewers**) cannot be changed during a recovery.
- Ownership of a shared folder cannot be changed during a recovery if the **Prevent editors from changing access and adding new people** option is enabled for this folder. This setting prevents the Google Drive API from listing the folder permissions. Ownership of the files in the folder is recovered correctly.

15.20.3.1 Selecting Google Drive files

Select the files as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select Google Drive files

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the files of all users (including users that will be created in the future), expand the **Users** node, select **All users**, and then click **Group backup**.
 - To back up the files of individual users, expand the **Users** node, select **All users**, select the users whose files you want to back up, and then click **Backup**.
4. On the protection plan panel:
 - Ensure that the **Google Drive** item is selected in **What to back up**.
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all files).
 - Specify the files and folders to back up by adding their names or paths.
You can use wildcard characters (*, **, and ?). For more details about specifying paths and using wildcards, refer to "File filters" (p. 153).
 - Specify the files and folders to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single user.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the files and folders to skip during the backup.
File exclusions override the file selection; i.e. if you specify the same file in both fields, this file will be skipped during a backup.
 - If you want to back up the files that are shared with the selected users, enable the **Include shared files** switch.
 - If you want to enable notarization of all files selected for backup, enable the **Notarization** switch. For more information about notarization, refer to "Notarization" (p. 263).

15.20.3.2 Recovering Google Drive and Google Drive files

Recovering an entire Google Drive

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose Google Drive you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users by name. Wildcards are not supported.

4. Select a recovery point.

Tip. To see only the recovery points that contain Google Drive files, select **Google Drive** in **Filter by content**.

5. Click **Recover > Entire Drive**.
6. If multiple G Suite organizations were added to the Cyber Protection service, click **G Suite organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to drive**, view, change, or specify the target user or the target Shared drive.

By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user or the target Shared drive.

If the backup contains shared files, the files will be recovered to the root folder of the target drive.
8. Select whether to recover the sharing permissions for the files.
9. Click **Start recovery**.
10. Select one of the overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**
11. Click **Proceed** to confirm your decision.

Recovering Google Drive files

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Users** node, select **All users**, select the user whose Google Drive files you want to recover, and then click **Recovery**.

If the user was deleted, select the user in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.

You can search users by name. Wildcards are not supported.

4. Select a recovery point.

Tip. To see only the recovery points that contain Google Drive files, select **Google Drive** in **Filter by content**.

5. Click **Recover > Files/folders**.

6. Browse to the required folder or use search to obtain the list of the required files and folders.
The search is not available if the backup is encrypted.
7. Select the files that you want to recover.
If the backup is not encrypted and you selected a single file, you can click **Show versions** to select the file version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. If you want to download a file, select the file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
9. Click **Recover**.
10. If multiple G Suite organizations were added to the Cyber Protection service, click **G Suite organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
11. In **Recover to drive**, view, change, or specify the target user or the target Shared drive.
By default, the original user is selected. If this user does not exist or a non-original organization is selected, you must specify the target user or the target Shared drive.
12. In **Path**, view or change the target folder in the target user's Google Drive or in the target Shared drive. By default, the original location is selected.
13. Select whether to recover the sharing permissions for the files.
14. Click **Start recovery**.
15. Select one of the file overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**
16. Click **Proceed** to confirm your decision.

15.20.4 Protecting Shared drive files

What items can be backed up?

You can back up an entire Shared drive, or individual files and folders.

Files are backed up together with their sharing permissions.

Limitations

- A Shared drive without members cannot be backed up, due to Google Drive API limitations.
- Out of Google-specific file formats only Google docs, Google sheets, Google slides, and Google Drawings are backed up.

What items can be recovered?

You can recover an entire Shared drive, or any file or folder that was backed up.

You can use search to locate items in a backup, unless the backup is encrypted. Search in encrypted backups is not supported.

You can choose whether to recover the sharing permissions or let the files inherit the permissions from the folder to which they are recovered.

The following items are not recovered:

- Sharing permissions for a file that was shared with a user outside the organization are not recovered if sharing outside the organization is disabled in the target Shared drive.
- Sharing permissions for a file that was shared with a user who is not a member of the target Shared drive are not recovered if **Sharing with non-members** is disabled in the target Shared drive.

Limitations

- Comments in files are not recovered.
- Sharing links for files and folders are not recovered.

15.20.4.1 Selecting Shared drive files

Select the files as described below, and then specify other settings of the protection plan as appropriate (p. 111).

To select Shared drive files

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose users' data you want to back up. Otherwise, skip this step.
3. Do one of the following:
 - To back up the files of all Shared drive (including Shared drive that will be created in the future), expand the **Shared drives** node, select **All Shared drives**, and then click **Group backup**.
 - To back up the files of individual Shared drives, expand the **Shared drives** node, select **All Shared drives**, select the Shared drives to back up, and then click **Backup**.
4. On the protection plan panel:
 - In **Items to back up**, do one of the following:
 - Keep the default setting **[All]** (all files).
 - Specify the files and folders to back up by adding their names or paths.
You can use wildcard characters (*, **, and ?). For more details about specifying paths and using wildcards, refer to "File filters" (p. 153).
 - Specify the files and folders to back up by browsing.
The **Browse** link is available only when creating a protection plan for a single Shared drive.
 - [Optional] In **Items to back up**, click **Show exclusions** to specify the files and folders to skip during the backup.
File exclusions override the file selection; i.e. if you specify the same file in both fields, this file will be skipped during a backup.
 - If you want to enable notarization of all files selected for backup, enable the **Notarization** switch. For more information about notarization, refer to "Notarization" (p. 263).

15.20.4.2 Recovering Shared drive and Shared drive files

Recovering an entire Shared drive

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.

3. Expand the **Shared drives** node, select **All Shared drives**, select the Shared drive that you want to recover, and then click **Recovery**.
If the Shared drive was deleted, select it in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.
You can search Shared drives by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Entire Shared drive**.
6. If multiple G Suite organizations were added to the Cyber Protection service, click **G Suite organization** to view, change, or specify the target organization.
By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.
7. In **Recover to drive**, view, change, or specify the target Shared drive or the target user. If you specify a user, the data will be recovered to this user's Google Drive.
By default, the original Shared drive is selected. If this Shared drive does not exist or a non-original organization is selected, you must specify the target Shared drive or the target user.
8. Select whether to recover the sharing permissions for the files.
9. Click **Start recovery**.
10. Select one of the overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**
11. Click **Proceed** to confirm your decision.

Recovering Shared drive files

1. Click **G Suite**.
2. If multiple G Suite organizations were added to the Cyber Protection service, select the organization whose backed-up data you want to recover. Otherwise, skip this step.
3. Expand the **Shared drives** node, select **All Shared drives**, select the Shared drive that originally contained the files you want to recover, and then click **Recovery**.
If the Shared drive was deleted, select it in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.
You can search Shared drives by name. Wildcards are not supported.
4. Select a recovery point.
5. Click **Recover > Files/folders**.
6. Browse to the required folder or use search to obtain the list of the required files and folders.
The search is not available if the backup is encrypted.
7. Select the files that you want to recover.
If the backup is not encrypted and you selected a single file, you can click **Show versions** to select the file version to recover. You can select any backed-up version, earlier or later than the selected recovery point.
8. If you want to download a file, select the file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
9. Click **Recover**.
10. If multiple G Suite organizations were added to the Cyber Protection service, click **G Suite organization** to view, change, or specify the target organization.

By default, the original organization is selected. If this organization is no longer registered in the Cyber Protection service, you must specify the target organization.

11. In **Recover to drive**, view, change, or specify the target Shared drive or the target user. If you specify a user, the data will be recovered to this user's Google Drive.

By default, the original Shared drive is selected. If this Shared drive does not exist or a non-original organization is selected, you must specify the target Shared drive or the target user.

12. In **Path**, view or change the target folder in the target Shared drive or the target user's Google Drive. By default, the original location is selected.
13. Select whether to recover the sharing permissions for the files.
14. Click **Start recovery**.
15. Select one of the file overwriting options:
 - **Overwrite existing files**
 - **Overwrite an existing file if it is older**
 - **Do not overwrite existing files**
16. Click **Proceed** to confirm your decision.

15.20.5 Notarization

Notarization enables you to prove that a file is authentic and unchanged since it was backed up. We recommend that you enable notarization when backing up your legal document files or other files that require proved authenticity.

Notarization is available only for backups of Google Drive files and G Suite Shared drive files.

How to use notarization

To enable notarization of all files selected for backup, enable the **Notarization** switch when creating a protection plan.

When configuring recovery, the notarized files will be marked with a special icon, and you can verify the file authenticity (p. 189).

How it works

During a backup, the agent calculates the hash codes of the backed-up files, builds a hash tree (based on the folder structure), saves the tree in the backup, and then sends the hash tree root to the notary service. The notary service saves the hash tree root in the Ethereum blockchain database to ensure that this value does not change.

When verifying the file authenticity, the agent calculates the hash of the file, and then compares it with the hash that is stored in the hash tree inside the backup. If these hashes do not match, the file is considered not authentic. Otherwise, the file authenticity is guaranteed by the hash tree.

To verify that the hash tree itself was not compromised, the agent sends the hash tree root to the notary service. The notary service compares it with the one stored in the blockchain database. If the hashes match, the selected file is guaranteed to be authentic. Otherwise, the software displays a message that the file is not authentic.

15.20.5.1 Verifying file authenticity with Notary Service

If notarization was enabled during backup, you can verify the authenticity of a backed-up file.

To verify the file authenticity

1. Do one of the following:
 - To verify the authenticity of a Google Drive file, select the file as described in steps 1-7 of the "Recovering Google Drive files" (p. 259) section.
 - To verify the authenticity of a G Suite Shared drive file, select the file as described in steps 1-7 of the "Recovering Shared drive files" (p. 262) section.
2. Ensure that the selected file is marked with the following icon: . This means that the file is notarized.
3. Do one of the following:
 - Click **Verify**.
The software checks the file authenticity and displays the result.
 - Click **Get certificate**.
A certificate that confirms the file notarization is opened in a web browser window. The window also contains instructions that allow you to verify the file authenticity manually.

15.21 Protecting Oracle Database

Protection of Oracle Database is described in a separate document available at https://dl.managed-protection.com/u/pdf/OracleBackup_whitepaper.pdf

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

15.22 Protecting SAP HANA

Protection of SAP HANA is described in a separate document available at https://dl.managed-protection.com/u/pdf/SAP%20HANA_backup_whitepaper.pdf

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

15.23 Protecting websites and hosting servers

15.23.1 Protecting websites

A website can be corrupted as a result of unauthorized access or a malware attack. Back up your website if you want to easily revert it to a healthy state, in case of corruption.

What do I need to back up a website?

The website must be accessible via the SFTP or SSH protocol. You do not need to install an agent, just add a website as described later in this section.

What items can be backed up?

You can back up the following items:

- **Website content files**
All files accessible to the account you specify for the SFTP or SSH connection.
- **Linked databases (if any) hosted on MySQL servers.**
All databases accessible to the MySQL account you specify.

If your website employs databases, we recommend that you back up both the files and the databases, to be able to recover them to a consistent state.

Limitations

- The only backup location available for website backup is the cloud storage.
- It is possible to apply several protection plans to a website, but only one of them can run on a schedule. Other plans need to be started manually.
- The only available backup option is "Backup file name" (p. 145).
- The website protection plans are not shown on the **Plans > Protection** tab.

15.23.1.1 Backing up a website

To add a website

1. Click **Devices > Add**.
2. Click **Website**.
3. Configure the following access settings for the website:
 - In **Website name**, create and type a name for your website. This name will be displayed in the service console.
 - In **Host**, specify the host name or IP address that will be used to access the website via SFTP or SSH. For example, `my.server.com` or `10.250.100.100`.
 - In **Port**, specify the port number.
 - In **User name** and **Password**, specify the credentials of the account that can be used to access the website via SFTP or SSH.

Important Only the files that are accessible to the specified account will be backed up.

Instead of a password, you can specify your private SSH key. To do this, select the **Use SSH private key instead of password** check box, and then specify the key.

4. Click **Next**.
5. If your website uses MySQL databases, configure the access settings for the databases. Otherwise, click **Skip**.
 - a. In **Connection type**, select how to access the databases from the cloud:
 - **Via SSH from host**—The databases will be accessed via the host specified in step 3.
 - **Direct connection**—The databases will be accessed directly. Choose this setting only if the databases are accessible from the Internet.
 - b. In **Host**, specify the name or IP address of the host where the MySQL server is running.
 - c. In **Port**, specify the port number for the TCP/IP connection to the server. The default port number is 3306.
 - d. In **User name** and **Password**, specify the MySQL account credentials.

Important Only the databases that are accessible to the specified account will be backed up.

- e. Click **Create**.

The website appears in the service console under **Devices > Websites**.

To change the connection settings

1. Select the website under **Devices > Websites**.
2. Click **Details**.
3. Click the pencil icon next to the website or the database connection settings.
4. Do the necessary changes, and then click **Save**.

To create a protection plan for websites

1. Select a website or several websites under **Devices > Websites**.
2. Click **Protect**.
3. [Optional] Enable backup of databases.
If several websites are selected, backup of databases is disabled by default.
4. [Optional] Change the retention rules (p. 136).
5. [Optional] Enable encryption of backups (p. 138).
6. [Optional] Click the gear icon to edit the **Backup file name** (p. 145) option. This makes sense in two cases:
 - If you backed up this website earlier and want to continue the existing sequence of backups
 - If you want to see the custom name on the **Backup storage** tab
7. Click **Apply**.

You can edit, revoke, and delete protection plans for websites in the same way as for machines. These operations are described in "Operations with protection plans".

15.23.1.2 Recovering a website

To recover a website

1. Do one of the following:
 - Under **Devices > Websites**, select the website that you want to recover, and then click **Recovery**.
You can search websites by name. Wildcards are not supported.
 - If the website was deleted, select it in the **Cloud applications backups** section of the Backup storage tab (p. 199), and then click **Show backups**.
To recover a deleted website, you need to add the target site as a device.
2. Select the recovery point.
3. Click **Recover**, and then select what you want to recover: **Entire website**, **Databases** (if any), or **Files/folders**.
To ensure that your website is in a consistent state, we recommend recovering both files and databases, in any order.
4. Depending on your choice, follow one of the procedures described below.

To recover the entire website

1. In **Recover to website**, view or change the target website.
By default, the original website is selected. If it does not exist, you must select the target website.
2. Select whether to recover the sharing permissions of the recovered items.
3. Click **Start recovery**, and then confirm the action.

To recover the databases

1. Select the databases that you want to recover.
2. If you want to download a database as a file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
3. Click **Recover**.
4. In **Recover to website**, view or change the target website.

By default, the original website is selected. If it does not exist, you must select the target website.

5. Click **Start recovery**, and then confirm the action.

To recover the website files/folders

1. Select the files/folders that you want to recover.
2. If you want to save a file, click **Download**, select the location to save the file to, and then click **Save**. Otherwise, skip this step.
3. Click **Recover**.
4. In **Recover to website**, view or change the target website.

By default, the original website is selected. If it does not exist, you must select the target website.

5. Select whether to recover the sharing permissions of the recovered items.
6. Click **Start recovery**, and then confirm the action.

15.23.2 Protecting web hosting servers

Web hosting administrators that use the Plesk or cPanel platforms can integrate these platforms with the Cyber Protection service.

The integration enables an administrator to do the following:

- Back up an entire Plesk or cPanel server to the cloud storage, with disk-level backup
- Recover the entire server, including all of the websites
- For Plesk: perform granular recovery of websites, individual files, mailboxes, or databases
- For cPanel: perform granular recovery of websites, individual files, mailboxes, mail filters, mail forwarders, databases, and accounts
- Enable self-service recovery for Plesk and cPanel customers

The integration is performed by using the Cyber Protection service extension. If you need the extension for Plesk or cPanel, contact the provider of the Cyber Protection service.

Supported Plesk and cPanel versions

- Plesk for Linux 17.0 and later
- Any cPanel version with PHP 5.6 and later

Quotas

Each backed-up Plesk or cPanel server consumes the **Web hosting servers** quota. If this quota is disabled or the overage for this quota is exceeded, the following will happen:

- If the server is physical, the **Servers** quota will be used. If this quota is disabled or the overage for this quota is exceeded, the backup will fail.
- If the server is virtual, the **Virtual machines** quota will be used. If this quota is disabled or the overage for this quota is exceeded, the backup will fail.

15.24 Special operations with virtual machines

15.24.1 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 restore, 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 Cyber Protection 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.
- Backups that contain Linux logical volumes (LVM) must be created by Agent for VMware or Agent for Hyper-V. The virtual machine must be of the same type as the original machine (ESXi or Hyper-V).

15.24.1.1 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 Backup storage tab (p. 199).
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. If you are planning to preserve these changes by making the virtual machine permanent (p. 270), select a datastore that is suitable for running the machine in production.
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:



or

. Such virtual machines cannot be selected for backup.

15.24.1.2 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.

15.24.1.3 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 protection agent.

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.

The location of the final virtual disks is defined in the parameters of the **Run as VM** operation (**Datastore** for ESXi or **Path** for Hyper-V). Prior to starting the finalization, ensure that free space, sharing capabilities, and performance of this datastore are suitable for running the machine in production.

Note: Finalization is not supported for Hyper-V running in Windows Server 2008/2008 R2 and Microsoft Hyper-V Server 2008/2008 R2 because the necessary API is missing in these Hyper-V versions.

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.

What you need to know about finalization

Finalization vs. regular recovery

The finalization process is slower than a regular recovery for the following reasons:

- During a finalization, the agent performs random access to different parts of the backup. When an entire machine is being recovered, the agent reads data from the backup sequentially.

- If the virtual machine is running during the finalization, the agent reads data from the backup more often, to maintain both processes simultaneously. During a regular recovery, the virtual machine is stopped.

Finalization of machines running from cloud backups

Because of intensive access to the backed-up data, the finalization speed highly depends on the connection bandwidth between the backup location and the agent. The finalization will be slower for backups located in the cloud as compared to local backups. If the Internet connection is very slow or unstable, the finalization of a machine running from a cloud backup may fail. We recommend to run virtual machines from local backups if you are planning to perform finalization and have the choice.

15.24.2 Working in VMware vSphere

This section describes operations that are specific for VMware vSphere environments.

15.24.2.1 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. 274), 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. 273)
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. 273)
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.

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. 274).
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:

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.

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. 273)
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. 274)
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. 274)
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.

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.

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**.

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. 275).
7. Click **Start recovery**.
8. Confirm your decision.

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. 149).

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. 151).

Pre/Post commands

This option is similar to the backup option "Pre/Post commands" (p. 167).

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. 173).

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. 195).

Performance

This option is similar to the recovery option "Performance" (p. 197).

Pre/Post commands

This option is similar to the recovery option "Pre/Post commands" (p. 197).

VM power management

This option is similar to the recovery option "VM power management" (p. 199).

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. 272).

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.

15.24.2.2 Agent for VMware - 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. This capability is called a LAN-free backup.

The diagram below illustrates a LAN-based and a LAN-free backup. LAN-free access to virtual machines is available if you have a fibre channel (FC) or iSCSI Storage Area Network. To completely eliminate transferring the backed-up data via LAN, store the backups on a local disk of the agent's machine or on a SAN attached storage.



To enable the agent to access a datastore directly

1. Install Agent for VMware on a Windows machine that has network access to the vCenter Server.
2. Connect the logical unit number (LUN) that hosts the datastore to the machine. Consider the following:
 - Use the same protocol (i.e. iSCSI or FC) that is used for the datastore connection to the ESXi.
 - The LUN *must not* be initialized and must appear as an "offline" disk in **Disk Management**. If Windows initializes the LUN, it may become corrupted and unreadable by VMware vSphere.

As a result, the agent will use the SAN transport mode to access the virtual disks, i.e. it will read raw LUN sectors over iSCSI/FC without recognizing the VMFS file system (which Windows is not aware of).

Limitations

- In vSphere 6.0 and later, the agent cannot use the SAN transport mode if some of the VM disks are located on a VMware Virtual Volume (VVol) and some are not. Backups of such virtual machines will fail.
- Encrypted virtual machines, introduced in VMware vSphere 6.5, will be backed up via LAN, even if you configure the SAN transport mode for the agent. The agent will fall back on the NBD transport because VMware does not support SAN transport for backing up encrypted virtual disks.

Example

If you are using an iSCSI SAN, configure the iSCSI initiator on the machine running Windows where Agent for VMware is installed.

To configure the SAN policy

1. Log on as an administrator, open the command prompt, type **diskpart**, and then press **Enter**.
2. Type **san**, and then press **Enter**. Ensure that **SAN Policy : Offline All** is displayed.
3. If another value for SAN Policy is set:
 - a. Type **san policy=offlineall**.
 - b. Press **Enter**.
 - c. To check that the setting has been applied correctly, perform step 2.
 - d. Restart the machine.

To configure an iSCSI initiator


1. Go to **Control Panel > Administrative Tools > iSCSI Initiator**.

Tip. To find the **Administrative Tools** applet, you may need to change the **Control Panel** view to something other than **Home** or **Category**, or use search.

2. If this is the first time that Microsoft iSCSI Initiator is launched, confirm that you want to start the Microsoft iSCSI Initiator service.
3. On the **Targets** tab, type the fully qualified domain name (FQDN) name or the IP address of the target SAN device, and then click **Quick Connect**.
4. Select the LUN that hosts the datastore, and then click **Connect**.

If the LUN is not displayed, ensure that the zoning on the iSCSI target enables the machine running the agent to access the LUN. The machine must be added to the list of allowed iSCSI initiators on this target.
5. Click **OK**.

The ready SAN LUN should appear in **Disk Management** as shown in the screenshot below.



15.24.2.3 Using a locally attached storage

You can attach an additional disk to Agent for VMware (Virtual Appliance) so the agent can back up to this locally attached storage. This approach eliminates the network traffic between the agent and the backup location.

A virtual appliance that is running on the same host or cluster with the backed-up virtual machines has direct access to the datastore(s) where the machines reside. This means the appliance can attach the backed-up disks by using the HotAdd transport, and therefore the backup traffic is directed from one local disk to another. If the datastore is connected as **Disk/LUN** rather than **NFS**, the backup will be completely LAN-free. In the case of NFS datastore, there will be network traffic between the datastore and the host.

Using a locally attached storage presumes that the agent always backs up the same machines. If multiple agents work within the vSphere, and one or more of them use locally attached storages, you need to manually bind (p. 279) each agent to all machines it has to back up. Otherwise, if the machines are redistributed among the agents by the management server, a machine's backups may be dispersed over multiple storages.

You can add the storage to an already working agent or when deploying the agent from an OVF template (p. 65).

To attach a storage to an already working agent

1. In VMware vSphere inventory, right click the Agent for VMware (Virtual Appliance).
2. Add the disk by editing the settings of the virtual machine. The disk size must be at least 10 GB.

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

3. Go to the virtual appliance console. The **Create storage** link is available at the bottom of the screen. If it is not, click **Refresh**.
4. Click the **Create storage** link, select the disk and specify a label for it. The label length is limited to 16 characters, due to file system restrictions.

To select a locally attached storage as a backup destination

When creating a protection plan (p. 109), in **Where to back up**, select **Local folders**, and then type the letter corresponding to the locally attached storage, for example, **D:**.

15.24.2.4 Virtual machine binding

This section gives you an overview of how the Cyber Protection service organizes the operation of multiple agents within VMware vCenter.

The below distribution algorithm works for both virtual appliances and agents installed in Windows.

Distribution algorithm

The virtual machines are automatically evenly distributed between Agents for VMware. By evenly, we mean that each agent manages an equal number of machines. The amount of storage space occupied by a virtual machine is not counted.

However, when choosing an agent for a machine, the software tries to optimize the overall system performance. In particular, the software considers the agent and the virtual machine location. An agent hosted on the same host is preferred. If there is no agent on the same host, an agent from the same cluster is preferred.

Once a virtual machine is assigned to an agent, all backups of this machine are delegated to this agent.

Redistribution

Redistribution takes place each time the established balance breaks, or, more precisely, when a load imbalance among the agents reaches 20 percent. This may happen when a machine or an agent is added or removed, or a machine migrates to a different host or cluster, or if you manually bind a machine to an agent. If this happens, the Cyber Protection service redistributes the machines using the same algorithm.

For example, you realize that you need more agents to help with throughput and deploy an additional virtual appliance to the cluster. The Cyber Protection service will assign the most appropriate machines to the new agent. The old agents' load will reduce.

When you remove an agent from the Cyber Protection service, the machines assigned to the agent are distributed among the remaining agents. However, this will not happen if an agent gets corrupted or is deleted from manually from vSphere. Redistribution will start only after you remove such agent from the web interface.

Viewing the distribution result

You can view the result of the automatic distribution:

- in the **Agent** column for each virtual machine on the **All devices** section

- in the **Assigned virtual machines** section of the **Details** panel when an agent is selected in the **Settings > Agents** section

Manual binding

The Agent for VMware binding lets you exclude a virtual machine from this distribution process by specifying the agent that must always back up this machine. The overall balance will be maintained, but this particular machine can be passed to a different agent only if the original agent is removed.

To bind a machine with an agent

1. Select the machine.
2. Click **Details**.
In the **Assigned agent** section, the software shows the agent that currently manages the selected machine.
3. Click **Change**.
4. Select **Manual**.
5. Select the agent to which you want to bind the machine.
6. Click **Save**.

To unbind a machine from an agent

1. Select the machine.
2. Click **Details**.
In the **Assigned agent** section, the software shows the agent that currently manages the selected machine.
3. Click **Change**.
4. Select **Automatic**.
5. Click **Save**.

Disabling automatic assignment for an agent

You can disable the automatic assignment for Agent for VMware to exclude it from the distribution process by specifying the list of machines that this agent must back up. The overall balance will be maintained between other agents.

Automatic assignment cannot be disabled for an agent if there are no other registered agents, or if automatic assignment is disabled for all other agents.

To disable automatic assignment for an agent

1. Click **Settings > Agents**.
2. Select Agent for VMware for which you want to disable the automatic assignment.
3. Click **Details**.
4. Disable the **Automatic assignment** switch.

Usage examples

- Manual binding comes in handy if you want a particular (very large) machine to be backed up by Agent for VMware (Windows) via a fibre channel while other machines are backed up by virtual appliances.
- It is necessary to bind VMs to an agent if the agent has a locally attached storage.
- Disabling the automatic assignment enables you to ensure that a particular machine is predictably backed up on the schedule you specify. The agent that only backs up one VM cannot be busy backing up other VMs when the scheduled time comes.

- Disabling the automatic assignment is useful if you have multiple ESXi hosts that are separated geographically. If you disable the automatic assignment, and then bind the VMs on each host to the agent running on the same host, you can ensure that the agent will never back up any machines running on the remote ESXi hosts, thus saving network traffic.

15.24.2.5 Support for VM migration

This section informs you about what to expect when virtual machines migrate within a vSphere environment, including migration between ESXi hosts that are part of a vSphere cluster.

vMotion

vMotion moves a virtual machine's state and configuration to another host while the machine's disks remain in the same location on shared storage.

- vMotion of Agent for VMware (Virtual Appliance) is not supported and is disabled.
- vMotion of a virtual machine is disabled during a backup. Backups will continue to run after the migration is completed.

Storage vMotion

Storage vMotion moves virtual machine disks from one datastore to another.

- Storage vMotion of Agent for VMware (Virtual Appliance) is not supported and is disabled.
- Storage vMotion of a virtual machine is disabled during a backup. Backups will continue to run after the migration.

15.24.2.6 Managing virtualization environments


You can view the vSphere, Hyper-V, and Virtuozzo environments in their native presentation. Once the corresponding agent is installed and registered, the **VMware**, **Hyper-V**, or **Virtuozzo** tab appears under **Devices**.

In the **VMware** tab, you can back up the following vSphere infrastructure objects:

- Data center
- Folder
- Cluster
- ESXi host
- Resource pool

Each of these infrastructure objects works as a group object for virtual machines. When you apply a protection plan to any of these group objects, all virtual machines included in it, will be backed up. You can back up either the selected group machines by clicking **Protect**, or the parent group machines in which the selected group is included by clicking **Protect group**.

For example, you have selected the San Stefano cluster and then selected the resource pool inside it. If you click **Protect**, all virtual machines included in the selected resource pool will be backed up. If you click **Protect group**, all virtual machines included in the San Stefano cluster will be backed up.



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 **Details**.
5. Under **Credentials**, click the user name.
6. Specify the new access credentials, and then click **OK**.

15.24.2.7 Viewing backup status in vSphere Client

You can view backup status and the last backup time of a virtual machine in vSphere Client.

This information appears in the virtual machine summary (**Summary > Custom attributes/Annotations/Notes**, depending on the client type and vSphere version). You can also enable the **Last backup** and **Backup status** columns on the **Virtual Machines** tab for any host, datacenter, folder, resource pool, or the entire vCenter Server.

To provide these attributes, Agent for VMware must have the following privileges in addition to those described in "Agent for VMware - necessary privileges" (p. 282):

- **Global > Manage custom attributes**
- **Global > Set custom attribute**

15.24.2.8 Agent for VMware - necessary privileges


To perform any operations with vCenter objects, such as virtual machines, ESXi hosts, clusters, vCenter, and more, Agent for VMware authenticates on vCenter or ESXi host by using the vSphere

credentials provided by a user. The vSphere account, used for connection to vSphere by Agent for VMware, must have the required privileges on all levels of vSphere infrastructure starting from the vCenter level.

Specify the vSphere 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. 281) section.

To assign the permissions to a vSphere user on the vCenter level, do the following:

1. Log in to vSphere web client.
2. Right-click on vCenter and then click **Add permission**.
3. Select or add a new user with the required role (the role must include all the required permissions from the table below).
4. Select the **Propagate to children** option.



Object	Privilege	Operation			
		Back up a VM	Recover to a new VM	Recover to an existing VM	Run VM from backup
Cryptographic operations (starting with vSphere 6.5)	Add disk	+	*		
	Direct Access	+	*		
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
	Low level file operations				+
Global	Licenses	+	+	+	+
	Disable methods	+	+	+	
	Enable methods	+	+	+	
	Manage custom attributes	+	+	+	
	Set custom attribute	+	+	+	
Host > Configuration	Storage partition configuration				+
Host > Local operations	Create VM				+
	Delete VM				+
	Reconfigure VM				+
Network	Assign network		+	+	+
Resource	Assign VM to resource pool		+	+	+
Virtual machine > Configuration	Add existing disk	+	+		+
	Add new disk		+	+	+
	Add or remove device		+		+
	Advanced	+	+	+	
	Change CPU count		+		
	Disk change tracking	+		+	
	Disk lease	+		+	
	Memory		+		
	Remove disk	+	+	+	+
	Rename		+		
	Set annotation				+
	Settings		+	+	+
Virtual machine > Guest Operations	Guest Operation Program Execution	+**			
	Guest Operation Queries	+**			
	Guest Operation Modifications	+**			

		Operation			
Object	Privilege	Back up a VM	Recover to a new VM	Recover to an existing VM	Run VM from backup
Virtual machine > Interaction	Acquire guest control ticket (in vSphere 4.1 and 5.0)				+
	Configure CD media		+	+	
	Guest operating system management by VIX API (in vSphere 5.1 and later)				+
	Power off			+	+
	Power on		+	+	+
Virtual machine > Inventory	Create from existing		+	+	+
	Create new		+	+	+
	Register				+
	Remove		+	+	+
	Unregister				+
Virtual machine > Provisioning	Allow disk access		+	+	+
	Allow read-only disk access	+		+	
	Allow virtual machine download	+	+	+	+
Virtual machine > State	Create snapshot	+		+	+
	Remove snapshot	+		+	+
vApp	Add virtual machine				+

* This privilege is required for backing up encrypted machines only.

** This privilege is required for application-aware backups only.

15.24.3 Backing up clustered Hyper-V machines

In a Hyper-V cluster, virtual machines may migrate between cluster nodes. Follow these recommendations to set up a correct backup of clustered Hyper-V machines:

1. A machine must be available for backup no matter what node it migrates to. To ensure that Agent for Hyper-V can access a machine on any node, the agent service must run under a domain user account that has administrative privileges on each of the cluster nodes.

We recommend that you specify such an account for the agent service during the Agent for Hyper-V installation.

2. Install Agent for Hyper-V on each node of the cluster.
3. Register all of the agents in the Cyber Protection service.

High Availability of a recovered machine

When you recover backed-up disks to an *existing* Hyper-V virtual machine, the machine's High Availability property remains as is.

When you recover backed-up disks to a *new* Hyper-V virtual machine, the resulting machine is not highly available. It is considered as a spare machine and is normally powered off. If you need to use the machine in the production environment, you can configure it for High Availability from the **Failover Cluster Management** snap-in.

15.24.4 Limiting the total number of simultaneously backed-up virtual machines

The **Scheduling** (p. 170) backup option defines how many virtual machines an agent can back up simultaneously when executing the given protection plan.

When multiple protection plans overlap in time, the numbers specified in their backup options are added up. Even though the resulting total number is programmatically limited to 10, overlapping plans can affect the backup performance and overload both the host and the virtual machine storage.

You can further reduce the total number of virtual machines that an Agent for VMware or Agent for Hyper-V can back up simultaneously.

To limit the total number of virtual machines that Agent for VMware (Windows) or Agent for Hyper-V can back up

1. On the machine running the agent, 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\MMS\Configuration\ManagedMachine\SimultaneousBackupsLimits]
"MaxNumberOfSimultaneousBackups"=dword:00000001
```

3. Replace 00000001 with the hexadecimal value of the limit that you want to set. For example, 00000001 is 1 and 0000000A is 10.
4. Save the document as **limit.reg**.
5. Run the file as an administrator.
6. Confirm that you want to edit the Windows registry.
7. 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
```

To limit the total number of virtual machines that Agent for VMware (Virtual Appliance) can back up

1. To start the command shell, press CTRL+SHIFT+F2 while in the virtual appliance UI.
2. Open the file **/etc/Acronis/MMS.config** in a text editor, such as **vi**.
3. Locate the following section:

```
<key name="SimultaneousBackupsLimits">  
  <value name="MaxNumberOfSimultaneousBackups" type="Tdwor">"10"</value>  
</key>
```

4. Replace **10** with the decimal value of the limit that you want to set.
5. Save the file.
6. Execute the **reboot** command to restart the agent.

15.24.5 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	Virtuozzo virtual machine	Virtuozzo container
Physical machine	+	+	+	-	-
VMware ESXi virtual machine	+	+	+	-	-
Hyper-V virtual machine	+	+	+	-	-
Virtuozzo virtual machine	+	+	+	+	-
Virtuozzo container	-	-	-	-	+

Note You cannot recover MacOS virtual machines to Hyper-V hosts, because Hyper-V does not support MacOS. You can recover MacOS virtual machines to a VMware host that is installed on Mac hardware.

For instructions on how to perform migration, refer to the following sections:

- Physical-to-virtual (P2V) - "Physical machine to virtual"
- Virtual-to-virtual (V2V) - "Virtual machine" (p. 182)
- Virtual-to-physical (V2P) - "Virtual machine" (p. 182) or "Recovering disks by using bootable media" (p. 183)

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:

- Perform P2V migration or V2P migration or V2V migration from Virtuozzo, of a Linux machine containing logical volumes (LVM). Use Agent for Linux or bootable media to create the backup and bootable media to recover.
- Provide drivers for specific hardware that is critical for the system bootability.

15.24.6 Windows Azure and Amazon EC2 virtual machines

To back up a Windows Azure or Amazon EC2 virtual machine, install a protection 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.

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. 178). When configuring the recovery, select the new machine as the target machine.

16 Disaster recovery

Note *This functionality is available only with the Disaster Recovery add-on of the Cyber Protection service.*

17 About Cyber Disaster Recovery Cloud

Cyber Disaster Recovery Cloud (DR) – a part of Cyber Protection that provides disaster recovery as a service (DRaaS) oriented mostly on the SMB clients. Cyber Disaster Recovery Cloud provides you with a fast and stable solution to launch the exact copies of your machines on the cloud site and switch the workload from the corrupted original machines to the recovery servers in the cloud in case of a man-made or a natural disaster.

The key functionality

- Manage the Cyber Disaster Recovery Cloud service from a single console
- Extend up to five local networks to the cloud, by using a secure VPN tunnel
- Establish the connection to the cloud site without any VPN appliance deployment (the cloud-only mode)
- Establish the point-to-site connection to your local and cloud sites
- Protect your machines by using recovery servers in the cloud
- Protect applications and appliances by using primary servers in the cloud
- Perform automatic disaster recovery operations for encrypted backups
- Perform a test failover in the isolated network
- Use runbooks to spin up the production environment in the cloud

18 Software requirements

Supported operating systems

Protection with a recovery server has been tested for the following operating systems:

- CentOS 6.6, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6
- Debian 9
- Ubuntu 16.04, 18.04
- Windows Server 2008/2008 R2
- Windows Server 2012/2012 R2
- Windows Server 2016 – all installation options, except for Nano Server
- Windows Server 2019 – all installation options, except for Nano Server

Windows desktop operating systems are not supported due to Microsoft product terms.

The software may work with other Windows operating systems and Linux distributions, but this is not guaranteed.

Supported virtualization platforms

Protection of virtual machines with a recovery server has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016
- Kernel-based Virtual Machines (KVM)
- Red Hat Enterprise Virtualization (RHEV) 3.6
- Red Hat Virtualization (RHV) 4.0
- Citrix XenServer: 6.5, 7.0, 7.1, 7.2

The VPN appliance has been tested for the following virtualization platforms:

- VMware ESXi 5.1, 5.5, 6.0, 6.5, 6.7
- Windows Server 2008 R2 with Hyper-V
- Windows Server 2012/2012 R2 with Hyper-V
- Windows Server 2016 with Hyper-V – all installation options, except for Nano Server
- Windows Server 2019 with Hyper-V – all installation options, except for Nano Server
- Microsoft Hyper-V Server 2012/2012 R2
- Microsoft Hyper-V Server 2016

The software may work with other virtualization platforms and versions, but this is not guaranteed.

Limitations

The following platforms and configurations are not supported in Cyber Disaster Recovery Cloud:

1. Unsupported platforms:

- Agents for Virtuozzo
- macOS

2. Unsupported configurations:

Microsoft Windows:

- Dynamic disks are not supported
- Windows desktop operating systems are not supported (due to Microsoft product terms)
- Active Directory service with FRS replication is not supported
- Removable media without either GPT or MBR formatting (so-called "superfloppy") are not supported

Linux:

- Linux machines that have logical volumes (LVM) or volumes formatted with the XFS file system
- File system without a partition table

3. Unsupported backup types:

- Continuous data protection (CDP) recovery points (p. 118) are incompatible.

Important *If you create a recovery server from a backup having a CDP recovery point, then during the failback or creating backup of a recovery server, you will lose the data contained in the CDP recovery point.*

- Forensic backups (p. 154) cannot be used for creating recovery servers.

A recovery server has one network interface. If the original machine has several network interfaces, only one is emulated.

Cloud servers are not encrypted.

19 Setting up the disaster recovery functionality

To set up the disaster recovery functionality

1. Configure the connectivity type to the cloud site:
 - Point-to-site connection (p. 308)
 - Site-to-site connection (p. 301)
 - Cloud-only mode (p. 304)
2. Create a protection plan (p. 95) with the backup module enabled and select the entire machine or system plus boot volumes for backing up. At least one protection plan is required for creating a recovery server.
3. Apply the protection plan to the local servers to be protected.
4. Create the recovery servers (p. 314) for each of your local servers that you want to protect.
5. Perform a test failover (p. 316) to check how it works.
6. [Optional] Create the primary servers (p. 321) for application replication.

As a result, you have set up the disaster recovery functionality to protect your local servers from a disaster.

If a disaster occurs, you can fail over the workload (p. 317) to the recovery servers in the cloud. At least one recovery point must be created before failing over to recovery servers. When your local site is recovered from a disaster, you can switch the workload back to your local site by performing a failback (p. 319).

20 Setting up connectivity

This section explains the network concepts necessary for you to understand how it all works in Cyber Disaster Recovery Cloud. You will learn how to configure different types of connectivity to the cloud site, depending on your needs. Finally, you will learn how to manage your networks in the cloud and manage the settings of the VPN appliance and VPN gateway.

20.1.1 Networking concepts

Cyber Disaster Recovery Cloud allows you to define the connectivity type to the cloud site:

- **Site-to-site connection**

This type of connection requires a VPN appliance deployment on the local site.

Your local site is connected to the cloud site by means of a secure VPN tunnel. This type of connection is suitable in case you have tightly dependent servers on the local site, such as a web server and a database server. In case of partial failover, when one of these servers is recreated on the cloud site while the other stays on the local site, they will still be able to communicate with each other via a VPN tunnel.

Cloud servers on the cloud site are accessible through the local network, point-to-site VPN, and public IP addresses (if assigned).

- **Cloud-only mode**

This type of connection does not require a VPN appliance deployment on the local site.

The local and cloud networks are independent networks. This type of connection implies either the failover of all the local site's protected servers or partial failover of independent servers that do not need to communicate with the local site.

Cloud servers on the cloud site are accessible through the point-to-site VPN, and public IP addresses (if assigned).

- **Point-to-site remote VPN access**

A secure point-to-site remote VPN access to your cloud and local site workloads from outside by using your endpoint device.


For a local site access, this type of connection requires a VPN appliance deployment on the local site.

This option is not available in the initial connectivity configuration screen. It is accessible after you establish the connection to the Cyber Disaster Recovery Cloud site.

20.1.1.1 Site-to-site connection

To understand how networking works in Cyber Disaster Recovery Cloud, we will consider a case when you have three networks with one machine each in the local site. You are going to configure the protection from a disaster for the two networks – Network 10 and Network 20.

On the diagram below, you can see the local site where your machines are hosted and the cloud site where the cloud servers are launched in case of a disaster. The Cyber Disaster Recovery Cloud solution allows you to fail over all the workload from the corrupted machines in the local site to the cloud servers in the cloud. A maximum of five networks can be protected with Cyber Disaster Recovery Cloud.



To establish a site-to-site communication between the local and cloud sites, **VPN appliance** and **VPN gateway** are used. First, when you start configuring the site-to-site connection in the service console, the VPN gateway is automatically deployed in the cloud site. Then, you must deploy the VPN appliance on your local site, add the networks to be protected, and register the appliance in the cloud. Cyber Disaster Recovery Cloud creates a replica of your local network in the cloud. A secure VPN tunnel is established between the VPN appliance and the VPN gateway. It provides your local network extension to the cloud. The production networks in the cloud are bridged with your local networks. The local and cloud servers can communicate via this VPN tunnel as if they are all in the same Ethernet segment.

For each source machine to be protected, you must create a recovery server on the cloud site. It stays in the **Standby** state until a failover event happens. If a disaster happens and you start a failover process (in the **production mode**), the recovery server representing the exact copy of your protected machine is launched in the cloud. It may be assigned the same IP address as the source machine has and launched in the same Ethernet segment. Your clients can continue working with the server, without noticing any background changes.

You can also launch a failover process in the **test mode**. This means that the source machine is still working and at the same time the respective recovery server with the same IP address is launched in the cloud. To prevent IP address conflicts, a special virtual network is created in the cloud – **test network**. The test network is isolated to prevent duplication of the source machine IP address in one Ethernet segment. To access the recovery server in the test failover mode, you must assign the **Test IP address** to the recovery server when creating it. There are other parameters for the recovery server that can be specified, they will be considered in the respective sections below.

How routing works

In case of the site-to-site connection established, routing between cloud networks is performed with your local router. The VPN server does not perform routing between the cloud servers located in

different cloud networks. If a cloud server from one network wants to communicate to a server from another cloud network, the traffic goes through the VPN tunnel to the local router on the local site, then the local router routes it to another network, and it goes back through the tunnel to the destination server on the cloud site.

VPN gateway

The major component that allows communication between the local and cloud sites is the **VPN gateway**. It is a virtual machine in the cloud on which the special software is installed, and the network is specifically configured. The VPN gateway provides the following functions:

- Connecting the Ethernet segments of your local network and production network in the cloud in the L2 mode.
- Providing iptables and ebtables rules.
- Working as a default router and NAT for the machines in the test and production networks.
- Working as a DHCP server. All machines in the production and test networks get the network configuration (IP addresses, DNS settings) via DHCP. Every time a cloud server will get the same IP address from the DHCP server. If you need to set up the custom DNS configuration, you should contact the support team.
- Working as a caching DNS.

VPN gateway network configuration

The VPN gateway has several network interfaces:

- External interface, connected to the Internet
- Production interfaces, connected to the production networks
- Test interface, connected to the test network

In addition, two virtual interfaces are added for point-to-site and site-to-site connections.

When the VPN gateway is deployed and initialized, the bridges are created – one for the external interface, and one for the client and production interfaces. Though the client-production bridge and the test interface use the same IP addresses, the VPN gateway can route packages correctly by using a specific technique.

VPN appliance

The **VPN appliance** is a virtual machine on the local site with Linux and the special software installed, and the special network configuration. It allows communications between the local and cloud sites.

Recovery servers

A **recovery server** – a replica of the original machine based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers in case of a disaster.

When creating a recovery server, you must specify the following network parameters:


- **Cloud network** (required): a cloud network to which a recovery server will be connected.
- **IP address in production network** (required): an IP address with which a virtual machine for a recovery server will be launched. This address is used in both the production and test networks. Before launching, the virtual machine is configured for getting the IP address via DHCP.
- **Test IP address** (optional): this IP address is needed to access a recovery server from the client-production network during the test failover, to prevent the production IP address from

being duplicated in the same network. This IP address is different from the IP address in the production network. Servers in the local site can reach the recovery server during the test failover via the test IP address, while access in the reverse direction is not available. Internet access from the recovery server in the test network is available if the **Internet access** option was selected during the recovery server creation.


- **Public IP address** (optional): an IP address used to access a recovery server from the Internet. If a server has no public IP address, it can be reached only from the local network.
- **Internet access** (optional): it allows a recovery server to access the Internet (in both the production and test failover cases).

Public and test IP address

If you assign the public IP address when creating a recovery server, it becomes available from the Internet via this IP address. When a packet comes from the Internet with the destination public IP address, the VPN gateway remaps it to the respective production IP address by using NAT, and then sends it to the corresponding recovery server.



If you assign the test IP address when creating a recovery server, it becomes available in the test network via this IP address. When you perform the test failover, the original machine is still running while the recovery server with the same IP address is launched in the test network in the cloud. There is no IP address conflict as the test network is isolated. The recovery servers in the test network are reachable by their test IP addresses, which are remapped to the production IP addresses via NAT.



Primary servers

A **primary server** – a virtual machine that does not have a linked machine on the local site if compared with a recovery server. Primary servers are used for protecting an application by means of replication or running various auxiliary services (such as a web server).

Typically, a primary server is used for real-time data replication across servers running crucial applications. You set up the replication by yourself, using the application's native tools. For example, Active Directory replication or SQL replication can be configured among the local servers and the primary server.

Alternatively, a primary server can be included in an AlwaysOn Availability Group (AAG) or Database Availability Group (DAG).

Both methods require a deep knowledge of the application and the administrator rights for it. A primary server constantly consumes computing resources and space on the fast disaster recovery storage. It needs maintenance on your side: monitoring the replication, installing software updates, backing up. The benefits are the minimal RPO and RTO with a minimal load on the production environment (as compared to backing up entire servers to the cloud).

Primary servers are always launched only in the production network and have the following network parameters:


- **Cloud network** (required): a cloud network to which a primary server will be connected.
- **IP address in production network** (required): an IP address that the primary server will have in the production network. By default, the first free IP address from your production network is set.
- **Public IP address** (optional): an IP address used to access a primary server from the Internet. If a server has no public IP address, it can be reached only from the local network, not via the Internet.
- **Internet access** (optional): allows a primary server to access the Internet.

20.1.1.2 Cloud-only mode

The cloud-only mode does not require a VPN appliance deployment on the local site. It implies that you have two independent networks: one on the local site, another on the cloud site.

How routing works

In case the cloud-only mode is established, routing is performed with the router on the cloud site so that servers from different cloud networks can communicate with each other.




20.1.1.3 Point-to-site remote VPN access

The point-to-site connection is a secure connection from the outside by using your endpoint devices (such as computer or laptop) to the cloud and local sites via a VPN. It is available after you establish a connection to the Cyber Disaster Recovery Cloud site. This type of connection can be used for the following cases:

- In many companies, the corporate services and web resources are available only from the corporate network. The point-to-site connection allows you to securely connect to the local site.
- In case of a disaster, when a workload is switched to the cloud site and your local network is down, you may need direct access to your cloud servers. This is possible via the point-to-site connection to the cloud site.

For the point-to-site connection to the local site, you need to install the VPN appliance on the local site, configure the site-to-site connection, and then the point-to-site connection to the local site. Thus, your remote employees will have access to the corporate network via L2 VPN.

The scheme below shows the local site, cloud site, and communications between servers highlighted in green. The L2 VPN tunnel connects your local and cloud sites. When a user establishes a point-to-site connection, the communications to the local site are performed through the cloud site.



The point-to-site configuration uses certificates to authenticate to the VPN client. Additionally user credentials are used for authentication. Note the following about the point-to-site connection to the local site:

- Users should use their Cyber Cloud credentials to authenticate in the VPN client. They must have either "Company Administrator" or "Cyber Protection" user role.
- If you re-generated the OpenVPN configuration (p. 310), you need to provide the updated configuration to all of the users using the point-to-site connection to the cloud site.

20.1.1.4 Automatic deletion of unused customer environment on cloud site

The Disaster Recovery service tracks usage of customer environment created for disaster recovery purposes and automatically deletes it if unused.

The following criteria are used to define that the customer tenant is active:

- Currently, there is at least one cloud server or there were cloud server(s) in the last seven days.
- OR
- The **VPN access to local site** option is enabled and either the site-to-site VPN tunnel is established or there are data reported from the VPN appliance for the last 7 days.

All the rest of the tenants are considered as inactive tenants. Therefore, for such tenants the system performs the following:

- The VPN gateway is deleted and all cloud resources related to the tenant are also deleted
- The VPN appliance is unregistered

Such tenants are rolled back to the state when there was no connectivity type configured.

20.1.2 Initial connectivity configuration

This section describes connectivity configuration scenarios.

20.1.2.1 Site-to-site connection

Requirements for the VPN appliance

System requirements

- 1 CPU
- 1 GB RAM
- 8 GB disk space

Ports

- TCP 443 (outbound) – for VPN connection
- TCP 80 (outbound) – for automatic update of the appliance (p. 307)


Ensure that your firewalls and other components of your network security system allow connections through these ports to any IP address.

Configuring site-to-site connection

The VPN appliance extends your local network to the cloud via a secure VPN tunnel. This kind of connection is often referred to as a "site-to-site" (S2S) connection. You can follow the procedure below or watch the video tutorial.


To set up a connection via the VPN appliance

1. In the service console, go to **Disaster Recovery**.




2. Click **Use Site-to-site VPN connection** and click **Deploy**.

The system starts deploying the VPN gateway in the cloud. This will take some time. Meanwhile, you can proceed to the next step.



Note The VPN gateway is provided without additional charge. It will be deleted if the disaster recovery functionality is not used, i.e. no primary or recovery server is present in the cloud for seven days.

3. In the **VPN appliance** block, click **Download and deploy**. Depending on the virtualization platform you are using, download the VPN appliance for VMware vSphere or Microsoft Hyper-V.




4. Deploy the appliance and connect it to the production networks.
In vSphere, ensure that **Promiscuous mode** and **Forged transmits** are enabled and set to **Accept** for all virtual switches that connect the VPN appliance to the production networks. To access these settings, in vSphere Client, select the host > **Summary** > **Network**, and then select the switch > **Edit settings...** > **Security**.
In Hyper-V, create a **Generation 1** virtual machine with 1024 MB of memory. We also recommend enabling **Dynamic Memory** for the machine. Once the machine is created, go to **Settings** > **Hardware** > **Network Adapter** > **Advanced Features** and select the **Enable MAC address spoofing** check box.
5. Power on the appliance.
6. Open the appliance console and log in with the "admin"/"admin" user name and password.

7. [Optional] Change the password.
8. [Optional] Change the network settings if needed. Define which interface will be used as the WAN for Internet connection.
9. Register the appliance in the Cyber Protection service by using the credentials of the company administrator.

These credentials are only used once to retrieve the certificate. The datacenter URL is predefined.

Note If two-factor authentication is configured for your account, you will also be prompted to enter the TOTP code. If two-factor authentication is enabled but not configured for your account, you cannot register the VPN appliance. First, you must go to the service console login page and complete the two-factor authentication configuration for your account. For more details on two-factor authentication, go to the Management Portal Administrator's Guide.


Once the configuration is complete, the appliance will have the **Online** status. The appliance connects to the VPN gateway and starts to report information about networks from all active interfaces to the Cyber Disaster Recovery Cloud service. The service console shows the interfaces, based on the information from the VPN appliance.



20.1.2.2 Cloud-only mode


To set up a connection in the cloud-only mode

1. In the service console, go to **Disaster Recovery**.



2. Click **Use Cloud-only connection**, and click **Deploy**.

As a result, the VPN gateway and cloud network with the defined address and mask will be deployed on the cloud site.



To learn how to manage your networks in the cloud and set up the VPN gateway settings, refer to "Managing cloud networks (p. 305)".

20.1.3 Network management


This section describes network management scenarios.

20.1.3.1 Managing networks

Site-to-site connection

To add a network on the local site and extend it to the cloud

1. On the VPN appliance, set up the new network interface with the local network that you want to extend in the cloud.
2. Log in to the VPN appliance console.
3. In the **Networking** section, set up network settings for the new interface.



The VPN appliance starts to report information about networks from all active interfaces to Cyber Disaster Recovery Cloud. The service console shows the interfaces based on the information from the VPN appliance.

To delete a network extended to the cloud

1. Log in to the VPN appliance console.
2. In the **Networking** section, select the interface that you want to delete, and then click **Clear network settings**.
3. Confirm the operation.

As a result, the local network extension to the cloud via a secure VPN tunnel will be stopped. This network will operate as an independent cloud segment. If this interface is used to pass the traffic from (to) the cloud site, all of your network connections from (to) the cloud site will be disconnected.

To change the network parameters

1. Log in to the VPN appliance console.
2. In the **Networking** section, select the interface that you want to edit.
3. Click **Edit network settings**.
4. Select one of the two possible options:
 - For automatic network configuration via DHCP, click **Use DHCP**. Confirm the operation.
 - For manual network configuration, click **Set static IP address**. The following settings are available for editing:
 - **IP address**: the IP address of the interface in the local network.
 - **VPN gateway IP address**: the special IP address which is reserved for the cloud segment of network for the proper Cyber Disaster Recovery Cloud service work.

- **Network mask:** network mask of the local network.
- **Default gateway:** default gateway on the local site.
- **Preferred DNS server:** primary DNS server on the local site.
- **Alternate DNS server:** secondary DNS server on the local site.

```

Disaster Recovery VPN Appliance
Registered by:                                     9.0.1.234
                                                [ldagny@mailinator.com]

Command: Networking \ configure ens160

Usage:
<Up>, <Down> - to select parameter
<Esc> - to cancel the command

IP address:
VPN gateway IP address:
Network mask:
Default gateway:
Preferred DNS server:
Alternate DNS server:

```

Make the necessary changes and confirm them by pressing Enter.

Cloud-only mode

You can have up to five networks in the cloud.

To add a new cloud network

1. Go to **Disaster Recovery > Connectivity**.
2. On **Cloud site**, click **Add cloud network**.
3. Define the cloud network parameters: the network address and mask. When ready, click **Done**.

As a result, the additional cloud network with the defined address and mask will be created on the cloud site.

To delete a cloud network

Note You cannot delete a cloud network if there is at least one cloud server in it. First, delete the cloud server, and then delete the network.

1. Go to **Disaster Recovery > Connectivity**.
2. On **Cloud site**, click the network address that you want to delete.
3. Click **Delete** and confirm the operation.

To change cloud network parameters

1. Go to **Disaster Recovery > Connectivity**.
2. On **Cloud site**, click the network address that you want to edit.
3. Click **Edit**.
4. Define the network address and mask, and click **Done**.

IP address reconfiguration

For proper disaster recovery performance, the IP addresses assigned to the local and cloud servers must be consistent. If there is any inconsistency or mismatch in IP addresses, you will see the exclamation mark next to the corresponding network in **Disaster Recovery > Connectivity**.

Some of the commonly known reasons of IP address inconsistency are listed below:

1. A recovery server was migrated from one network to another or the network mask of the cloud network was changed. As a result, cloud servers have the IP addresses from networks to which they are not connected.
2. The connectivity type was switched from the without site-to-site connection to the site-to-site connection. As a result, a local server is placed in the network different from the one that was created for the recovery server on the cloud site.
3. Editing the following network parameters on the VPN appliance site:
 - Adding an interface via the network settings
 - Editing the network mask manually via the interface settings
 - Editing the network mask via DHCP
 - Editing the network address and mask manually via the interface settings
 - Editing the network mask and address via DHCP

As a result of the actions listed above, the network on the cloud site may become a subset or superset of the local network, or the VPN appliance interface may report the same network settings for different interfaces.

To resolve the issue with network settings

1. Click the network that requires IP address reconfiguration.
You will see a list of servers in the selected network, their status, and IP addresses. The servers whose network settings are inconsistent are marked with the exclamation mark.
2. To change network settings for a server, click **Go to server**. To change network settings for all servers at once, click **Change** in the notification block.
3. Change the IP addresses as needed by defining them in the **New IP** and **New test IP** fields.
4. When ready, click **Confirm**.

20.1.3.2 Managing the VPN appliance settings

In the service console (**Disaster Recovery > Connectivity**), you can:

- Download log files
- Unregister the appliance (if you need to reset the VPN appliance settings or switch to the cloud-only mode)

To access these settings, click the **i** icon in the **VPN appliance** block.

In the VPN appliance console, you can:

- Change the password for the appliance
- View/change the network settings and define which interface to use as the WAN for the Internet connection
- Register/change the registration account (by repeating the registration)
- Restart the VPN service
- Reboot the VPN appliance
- Run the Linux shell command (only for advanced troubleshooting cases)

20.1.3.3 Enabling and disabling the site-to-site connection

You can enable the site-to-site connection in the following cases:

- If you need the cloud servers on the cloud site to communicate with servers on the local site.

- After a failover to the cloud, the local infrastructure is recovered, and you want to fail back your servers to the local site.

To enable the site-to-site connection

1. Go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**, and then enable the **Site-to-site connection** option.

As a result, the site-to-site VPN connection is enabled between the local and cloud sites. The Cyber Disaster Recovery Cloud service gets the network settings from the VPN appliance and extends the local networks to the cloud site.

If you do not need cloud servers on the cloud site to communicate with servers on the local site, you can disable the site-to-site connection.

To disable the site-to-site connection

1. Go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**, and then disable the **Site-to-site connection** option.

As a result, the local site is disconnected from the cloud site.

20.1.3.4 Configuring local routing

In addition to your local networks that are extended to the cloud via the VPN appliance, you may have other local networks that are not registered in the VPN appliance but the servers in them need to communicate with cloud servers. To establish the connectivity between such local servers and cloud servers, you need to configure the local routing settings.

To configure local routing

1. Go to **Disaster Recovery > Connectivity**.
2. Click **Show properties**, and then click **Local routing**.
3. Specify the local networks in the CIDR notation.
4. When ready, click **Save**.

As a result, the servers from the specified local networks will be able to communicate with the cloud servers.


20.1.3.5 Point-to-site remote VPN access

If you need to connect to your local site remotely, you can use the point-to-site connection to the local site. Follow the procedure below.

To configure the point-to-site connection

1. In the service console, go to **Disaster Recovery > Connectivity**.

2. Click **Show properties**.



3. Under **Point-to-site**, select **VPN access to local site**.

The system will automatically enable the site-to-site connection between the local and cloud sites and enable the point-to-site access to the local site.

4. Deploy the VPN appliance by clicking **Download VPN appliance**.

5. Ensure that your user who needs to establish the point-to-site connection to the local site has a user account in Cyber Cloud. These credentials are used for authentication in the VPN client. Otherwise, create a user account in Cyber Cloud. Ensure that a user has the "Company Administrator" or "Cyber Protection" user role.

6. Configure the OpenVPN client:

1. Download the OpenVPN client from the following location
<https://openvpn.net/community-downloads/> <https://openvpn.net/community-downloads/>.
The supported OpenVPN client versions: 2.4.0 and later.
2. Install the OpenVPN client on the machine from which you want to connect to the local site.
3. Click **Download configuration for OpenVPN**. The configuration file is valid for users in your organization with the "Company Administrator" or "Cyber Protection" user role.
4. Import the downloaded configuration to OpenVPN.
5. Log in to the OpenVPN client by using the Cyber Cloud user credentials (see step 4 above).
6. [Optional] If two-factor authentication is enabled for your organization, then you should provide the one-time generated TOTP code.

Important If you enabled two-factor authentication for your account, you need to re-generate the configuration file and renew it for your existing OpenVPN clients. Users must re-log in to Cyber Cloud to set up two-factor authentication for their accounts.


As a result, your user will be able to connect to machines on the local site.

To configure the point-to-site connection for a cloud-only connection

If you configured a cloud-only connection, and then want to switch to point-to-site connection, click **Show properties** and enable the **VPN access to local site** option. Follow the procedure above to complete the configuration.

20.1.3.6 Managing point-to-site connection settings

In the service console, go to **Disaster Recovery > Connectivity** and then click **Show properties** in the upper right corner.



VPN access to local site

This option is used for managing VPN access to the local site. By default it is enabled. If it is disabled, then the point-to-site access to the local site will be not allowed.

Download configuration for OpenVPN

This will download the configuration file for the OpenVPN client. The file is required to establish a point-to-site connection to the cloud site.

Re-generate configuration

You can re-generate the configuration file for the OpenVPN client.

This is required in the following cases:

- If you suspect that the configuration file is compromised.
- If two-factor authentication was enabled for your account.

As soon as the configuration file is updated, connecting by means of the old configuration file becomes not possible. Make sure to distribute the new file among the users who are allowed to use the point-to-site connection.

20.1.3.7 Active point-to-site connections

You can view all active point-to-site connections in **Disaster recovery > Connectivity**. Click the machine icon on the blue **Point-to-site** line and you will see the detailed information about active point-to-site connections grouped by the user name.

Connectivity


Active point-to-site connections

User name	Connections	Login at	Inbound traffic	Outbound traffic
> [redacted]@acronis.com	4	Jan, 10, 08:39 PM	11.2 GB	11.2 GB
▼ superadmin@acronis.com	2	—	4.6 GB	4.6 GB
10.96.77.16 - 8800		Jan, 09, 10:39 PM	1.6 GB	1.6 GB
10.96.77.16 - 8800		Jan, 09, 10:39 PM	4.6 GB	4.6 GB
> user@mail.com	1	Jan, 10, 08:39 PM	1.2 GB	1.2 GB
> 34get_2@hotmail.com	5	Jan, 10, 08:39 PM	3.1 GB	3.1 GB
> admin@acronis.com	1	Jan, 10, 08:39 PM	2 GB	2 GB
> man-23@yandex.com	5	Jan, 10, 08:39 PM	21.4 GB	21.4 GB

Show properties

1

Add cloud network



21 Setting up recovery servers

This section describes the concepts of failover and failback, a recovery server lifecycle, creation of a recovery server, and the disaster recovery operations.

21.1.1 How failover and failback work


Failover and failback

When a recovery server is created, it stays in the **Standby** state. The corresponding virtual machine does not exist until you initiate the failover. Before starting the failover process, you need to create at least one disk image backup (with bootable volume) of your original machine.

When starting the failover process, you select the recovery point of the original machine from which a virtual machine with the predefined parameters is created. The failover operation uses the "run VM from a backup" functionality. The recovery server gets the transition state **Finalization**. This process implies transferring the server's virtual disks from the backup storage ("cold" storage) to the disaster recovery storage ("hot" storage). During the finalization, the server is accessible and operable although the performance is lower than normal. When the finalization is completed, the server performance reaches its normal value. The server state changes to **Failover**. The workload is now switched from the original machine to the recovery server in the cloud site.

If the recovery server has a protection agent inside, the agent service is stopped in order to avoid interference (such as starting a backup or reporting outdated statuses to the backup component).

On the diagram below, you can see both the failover and failback processes.




Test failover

During a **test failover**, a virtual machine is not finalized. This means that the agent reads the virtual disks' content directly from the backup – that is, performs random access to different parts of the backup.

21.1.2 Recovery server lifecycle

On the diagram below, you can see a recovery server lifecycle, which shows server permanent states and transitional states. Each block shows a recovery server state, a corresponding virtual machine

state, and the actions that are available to a user at this stage. Each arrow is an event or user action that leads to the next state.



Failover and failback workflow

1. User action: Create a recovery server for the selected machine to be protected.
2. **Standby** state. The recovery server configuration is defined, but the corresponding virtual machine is not ready.
3. User action: The failover is initiated in the production mode and the recovery server is being created from the selected recovery point.
4. **Finalization** state. Virtual machine disks are finalized from the mounted recovery point to the high-performance storage. The recovery server is operational, though its performance is lower than normal until finalization is completed.
5. Event: Finalization is successful.
6. **Failover** state. The workload is switched from the original machine to the recovery server.
7. User actions:

- Initiate a failback. As a result, the recovery server is turned off and backed up to the cloud storage.
 - OR
 - If a user cancels the failover, then the workload is switched back to the original machine and the recovery server returns back to the **Standby** state.
8. **Ready for failback** state. The recovery server backup is created. You must recover your local server from this backup by using the regular recovery process.
 9. User actions:
 - Confirm failback. As a result, cloud resources that were allocated to the recovery server are released.
 - OR
 - Cancel failback. The failback is canceled by your request. The recovery server returns to the **Failover** state.

Test failover workflow

1. User action: Create a recovery server for the selected machine to be protected.
2. **Standby** state. The recovery server configuration is defined, but the respective virtual machine is not ready.
3. User action: Start testing the failover.
4. **Testing failover** state. In this state, a temporary virtual machine is created for testing purposes.
5. User action: Stop testing the failover.

21.1.3 Creating a recovery server

You can follow the instructions below or watch the video tutorial.

Prerequisites

- A protection plan must be applied to the original machine that you want to protect. This plan must back up the entire machine, or only the disks, required for booting up and providing the necessary services, to a cloud storage.
- One of the connectivity types to the cloud site must be set.

To create a recovery server

1. On the **All machines** tab, select the machine that you want to protect.
2. Click **Disaster recovery**, and then click **Create recovery server**.
3. Select the number of virtual cores and the size of RAM.
Be aware of the compute points next to every option. The number of compute points reflects the cost of running the recovery server per hour.
4. Specify the cloud network to which the server will be connected.
5. Specify the IP address that the server will have in the production network. By default, the IP address of the original machine is set.

Note If you use a DHCP server, add this IP address to the server exclusion list in order to avoid IP address conflicts.

6. [Optional] Select the **Test IP address** check box, and then specify the IP address.
This will give you the capability to test a failover in the isolated test network and to connect to the recovery server via RDP or SSH during a test failover. In the test failover mode, the VPN

gateway will replace the test IP address with the production IP address by using the NAT protocol.

If you leave the check box cleared, the console will be the only way to access the server during a test failover.

Note *If you use a DHCP server, add this IP address to the server exclusion list, in order to avoid IP address conflicts.*

You can select one of the proposed IP addresses or type in a different one.

7. [Optional] Select the **Internet access** check box.

This will enable the recovery server to access the Internet during a real or test failover.

8. [Optional] Set the **RPO threshold**.

The RPO threshold defines the maximum time interval allowed between the last suitable recovery point for a failover and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

9. [Optional] Select the **Use public IP address** check box.

Having a public IP address makes the recovery server available from the Internet during a failover or test failover. If you leave the check box cleared, the server will be available only in your production network. The **Use public IP address** option requires the **Internet access** option to be enabled.

The public IP address will be shown after you complete the configuration. The following ports are open for inbound connections to public IP addresses:


TCP: 80, 443, 8088, 8443

UDP: 1194

If you need other ports to be open, contact the support team.

10. [Optional] If the backups for the selected machine are encrypted, you can specify the password that will be automatically used when creating a virtual machine for the recovery server from the encrypted backup. Click **Specify**, and then define the credential name and password. By default, you will see the most recent backup in the list. To view all the backups, select **Show all backups**.
11. [Optional] Change the recovery server name.
12. [Optional] Type a description for the recovery server.
13. Click **Create**.

The recovery server appears in the **Disaster Recovery > Servers > Recovery servers** tab of the service console. You can also view its settings by selecting the original machine and clicking **Disaster recovery**.



21.1.4 Performing a test failover

Testing a failover means starting a recovery server in a test VLAN that is isolated from your production network. You can test several recovery servers at a time in order to check their interaction. In the test network, the servers communicate using their production IP addresses, but they cannot initiate TCP or UDP connections to the machines in your local network.


Though testing a failover is optional, we recommend that you make it a regular process with a frequency that you find adequate in terms of cost and safety. A good practice is creating a runbook – a set of instructions describing how to spin up the production environment in the cloud.

It is recommended to create a recovery server (p. 314) in advance to protect your devices from a disaster. You will be able to perform the test failover from any of the recovery points generated after the recovery server was created for the device.

To run a test failover

1. Select the original machine or select the recovery server that you want to test.
2. Click **Disaster Recovery**.
The description of the recovery server opens.
3. Click **Failover**.
4. Select the failover type **Test failover**.
5. Select the recovery point, and then click **Test failover**.

When the recovery server starts, its state changes to **Testing failover**.



6. Test the recovery server by using any of the following methods:

- In **Disaster Recovery > Servers**, select the recovery server, and then click **Console**.
- Connect to the recovery server by using RDP or SSH, and the test IP address that you specified when creating the recovery server. Try the connection from both inside and outside the production network (as described in "Point-to-site connection").
- Run a script within the recovery server.
The script may check the login screen, whether applications are started, the Internet connection, and the ability of other machines to connect to the recovery server.
- If the recovery server has access to the Internet and a public IP address, you may want to use TeamViewer.

7. When the test is complete, click **Stop testing**.

The recovery server is stopped. All changes made to the recovery server during the test failover are not preserved.

21.1.5 Performing a failover

A failover is a process of moving a workload from your premises to the cloud, and also the state when the workload remains in the cloud.

When you initiate a failover, the recovery server starts in the production network. All protection plans are revoked from the original machine. A new protection plan is automatically created and applied to the recovery server.

At least one recovery point must be created before failing over to a recovery server.

A good practice is to create a recovery server (p. 314) in advance to protect your devices from a disaster. You will be able to perform the production failover from any of the recovery points generated after the recovery server was created for the device.


You can follow the instructions below or watch the video tutorial.

To perform a failover

1. Ensure that the original machine is not available on the network.

2. In the service console, go to **Disaster recovery > Servers > Recovery servers** and select the recovery server.
3. Click **Failover**.
4. Select the type of failover **Production failover**.
5. Select the recovery point, and then click **Start production failover**.

When the recovery server starts, its state changes to **Finalization**, and after some time to **Failover**. It is critical to understand that the server is available in both states, despite the spinning progress indicator. For details, refer to "How failover and failback work" (p. 312).



6. Ensure that the recovery server is started by viewing its console. Click **Disaster Recovery > Servers**, select the recovery server, and then click **Console**.
7. Ensure that the recovery server can be accessed using the production IP address that you specified when creating the recovery server.

Once the recovery server is finalized, a new protection plan is automatically created and applied to it. This protection plan is based on the protection plan that was used for creating the recovery server, with certain limitations. In this plan, you can change only the schedule and retention rules. For more information, refer to "Backing up the cloud servers" (p. 324).

If you want to cancel failover, select the recovery server and click **Cancel failover**. All changes starting from the failover moment except the recovery server backups will be lost. The recovery server will return back to the **Standby** state.

If you want to perform failback (p. 319), select the recovery server and click **Failback**.

How to perform failover of servers using local DNS

If you use DNS servers on the local site for resolving machine names, then after a failover the recovery servers, corresponding to the machines relying on the DNS, will fail to communicate because the DNS servers used in the cloud are different. By default, the DNS servers of the cloud site are used for the newly created cloud servers. If you need to apply custom DNS settings, contact the support team.

How to perform failover of a DHCP server

Your local infrastructure may have the DHCP server located on a Windows or Linux host. When such a host is failed over to the cloud site, the DHCP server duplication issue occurs because the VPN gateway in the cloud also performs the DHCP role. To resolve this issue, do one of the following:

- If only the DHCP host was failed over to the cloud, while the rest local servers are still on the local site, then you must log in to the DHCP host in the cloud and turn off the DHCP server on it. Thus, there will be no conflicts and only the VPN gateway will work as the DHCP server.
- If your cloud servers already got the IP addresses from the DHCP host, then you must log in to the DHCP host in the cloud and turn off the DHCP server on it. You must also log in to the cloud servers and renew the DHCP lease to assign new IP addresses allocated from the correct DHCP server (hosted on the VPN gateway).


21.1.6 Performing a failback

A failback is a process of moving the workload from the cloud back to your premises.

During this process, the server being moved is unavailable. The length of the maintenance window is approximately equal to the duration of a backup and the subsequent recovery of the server.

To perform a failback

1. Select the recovery server that is in the **Failover** state.
2. Click **Failback**.



3. Click **Prepare failback**.

The recovery server will be stopped and backed up to the cloud storage. Wait for the backup to complete.

At this time, two actions become available: **Cancel failback** and **Confirm failback**. If you click **Cancel failback**, the recovery server will start and the failover will continue.

4. Recover the server from this backup to hardware or to a virtual machine on your premises.
 - When using bootable media, proceed as described in "Recovering disks by using bootable media" (p. 183) in the Cyber Protection User Guide. Ensure that you sign in to the cloud by using the account for which the server is registered and that you select the most recent backup.
 - If the target machine is online or is a virtual machine, you can use the service console. On the **Backup storage** tab, select the cloud storage. In **Machine to browse from**, select the target physical machine or the machine running the agent, if the target machine is virtual. The selected machine must be registered for the same account for which the server is registered. Find the most recent backup of the server, click **Recover entire machine**, and then set up

other recovery parameters. For the detailed instructions, refer to "Recovering a machine" (p. 178) in the Cyber Protection User Guide.

Ensure that the recovery is completed and the recovered machine works properly.

5. Return to the recovery server in the service console, and then click **Confirm failback**.

The recovery server and recovery points become ready for the next failover. To create new recovery points, apply a protection plan to the new local server.

21.1.7 Working with encrypted backups

You can create recovery servers from the encrypted backups. For your convenience, you can set up an automatic password application to an encrypted backup during the failover to a recovery server.

When creating a recovery server, you can specify the password to be used for automatic disaster recovery operations (p. 314). It will be saved to the Credentials store, a secure storage of credentials that can be found in **Settings > Credentials** section.

One credential can be linked to several backups.

To manage the saved passwords in the Credentials store

1. Go to **Settings > Credentials**.
2. To manage a specific credential, click the icon in the last column. You can view the items linked to this credential.
 - To unlink the backup from the selected credential, click the recycle bin icon near the backup. As a result, you will have to specify the password manually during the failover to the recovery server.
 - To edit the credential, click **Edit**, and then specify the name or password.
 - To delete the credential, click **Delete**. Note that you will have to specify the password manually during the failover to the recovery server.

22 Setting up primary servers

This section describes how to create and manage your primary servers.

22.1.1 Creating a primary server

Prerequisites

- One of the connectivity types to the cloud site must be set.

To create a primary server

1. Go to **Disaster Recovery > Servers > Primary servers** tab.
2. Click **Create**.
3. Select a template for the new virtual machine.
4. Select the number of virtual cores and the size of RAM.

Pay attention to the compute points next to every option. The number of compute points reflects the cost of running the primary server per hour.

5. [Optional] Change the virtual disk size. If you need more than one hard disk, click **Add disk**, and then specify the new disk size. Currently, you can add no more than 10 disks for a primary server.
6. Specify the cloud network in which the primary server will be included.
7. Specify the IP address that the server will have in the production network. By default, the first free IP address from your production network is set.

Note If you use a DHCP server, add this IP address to the server exclusion list in order to avoid IP address conflicts.

8. [Optional] Select the **Internet access** check box.

This will enable the primary server to access the Internet.

9. [Optional] Select the **Use public IP address** check box.

Having a public IP address makes the primary server available from the Internet. If you leave the check box cleared, the server will be available only in your production network.

The public IP address will be shown after you complete the configuration. The following ports are open for inbound connections to public IP addresses:

TCP: 80, 443, 8088, 8443

UDP: 1194


If you need other ports to be open, contact the support team.

10. [Optional] Select **Set RPO threshold**.

RPO threshold defines the maximum allowable time interval between the last recovery point and the current time. The value can be set within 15 – 60 minutes, 1 – 24 hours, 1 – 14 days.

11. Define the primary server name.
12. [Optional] Specify a description for the primary server.
13. Click **Create**.

The primary server becomes available in the production network. You can manage the server by using its console, RDP, SSH, or TeamViewer.



22.1.2 Operations with a primary server

The primary server appears in the **Disaster Recovery > Servers > Primary servers** tab in the service console.

To start or stop the server, click **Power on** or **Power off** on the primary server panel.

To edit the primary server settings, stop the server, and then click **Edit**.

To apply a protection plan to the primary server, select it and on the **Plan** tab click **Create**. You will see a predefined protection plan where you can change only the schedule and retention rules. For more information, refer to "Backing up the cloud servers" (p. 324).

23 Managing the cloud servers

To manage the cloud servers, go to **Disaster Recovery > Servers**. There are two tabs there: **Recovery servers** and **Primary servers**. To show all optional columns in the table, click the gear icon.

You can find the following information about each cloud server by selecting it.

Column name	Description
Name	A cloud server name defined by you
Status	The status reflecting the most severe issue with a cloud server (based on the active alerts)
State	A cloud server state according to its lifecycle (p. 312)
VM state	The power state of a virtual machine associated with a cloud server
Active location	The location where a cloud server is hosted. For example, Cloud .
RPO threshold	The maximum time interval allowed between the last suitable recovery point for failover and the current time. The value can be set within 15-60 minutes, 1-24 hours, 1-14 days.
RPO compliance	The RPO compliance is the ratio between the actual RPO and RPO threshold. The RPO compliance is shown if the RPO threshold is defined. It is calculated as follows: $\text{RPO compliance} = \text{Actual RPO} / \text{RPO threshold}$ where $\text{Actual RPO} = \text{current time} - \text{last recovery point time}$ RPO compliance statuses Depending on the value of the ratio between the actual RPO and RPO threshold, the following statuses are used: ▪ Compliant. The RPO compliance < 1x. A server meets the RPO threshold. ▪ Exceeded. The RPO compliance <= 2x. A server violates the RPO threshold. ▪ Severely exceeded. The RPO compliance <= 4x. A server violates the RPO threshold more than 2x times. ▪ Critically exceeded. The RPO compliance > 4x. A server violates the RPO threshold more than 4x times. ▪ Pending (no backups). The server is protected with the protection plan but the backup is being created and not completed yet.
Actual RPO	The time passed since the last recovery point creation
Last recovery point	The date and time when the last recovery point was created

23.1 Backing up the cloud servers

Primary and recovery servers are backed up by Agent for VMware, which is installed on the cloud site. In the initial release, this backup is somewhat restricted in functionality as compared to a backup performed by local agents. These limitations are temporary and will be removed in future releases.

- The only possible backup location is the cloud storage.
- A protection plan cannot be applied to multiple servers. Each server must have its own protection plan, even if all of the protection plans have the same settings.
- Only one protection plan can be applied to a server.
- Application-aware backup is not supported.
- Encryption is not available.
- Backup options are not available.

When you delete a primary server, its backups are also deleted.

A recovery server is backed up only in the failover state. Its backups continue the backup sequence of the original server. When a failback is performed, the original server can continue this backup sequence. So, the backups of the recovery server can only be deleted manually or as a result of applying the retention rules. When a recovery server is deleted, its backups are always kept.

Note *The protection plans for cloud servers are performed according to UTC time.*

24 Orchestration (runbooks)

A runbook is a set of instructions describing how to spin up the production environment in the cloud. You can create runbooks in the service console. To access the **Runbooks** tab, select **Disaster recovery** > **Runbooks**.

Why use runbooks?

Runbooks let you:

- Automate a failover of one or multiple servers
- Automatically check the failover result by pinging the server IP address and checking the connection to the port you specify
- Set the sequence of operations for servers running distributed applications
- Include manual operations in the workflow
- Verify the integrity of your disaster recovery solution, by executing runbooks in the test mode

24.1.1 Creating a runbook

You can follow the instruction below or watch the video tutorial.

To start creating a runbook, click **Create runbook** > **Add step** > **Add action**. You can use drag and drop to move actions and steps. Do not forget to give a distinctive name to the runbook. While creating a long runbook, click **Save** from time to time. Once you are finished, click **Close**.

The screenshot shows the 'New runbook' interface. On the left, 'Step 1' contains an action named 'Failover server' with a sub-action 'recovery' and the instruction 'Continue if already done'. Below the step is an 'Add step' button. On the right, the configuration panel for the selected action is visible. It includes a dropdown for 'Action' set to 'Failover server', checkboxes for 'Continue if already done' (checked) and 'Continue if failed' (unchecked), a 'Server' dropdown, a 'Completion check' section with 'Ping IP address' (checked) and 'Connect to port' (checked), and a 'Timeout in minutes' field set to 10.

Steps and actions

A runbook consists of steps that are executed consecutively. A step consists of actions that start simultaneously. An action may consist of:

- An operation to be performed with a cloud server (**Failover server**, **Start server**, **Stop server**, **Failback server**). To define this operation, you need to choose the operation, the cloud server, and the operation parameters.
- A manual operation that you need to describe verbally. Once the operation is completed, a user must click the confirmation button to allow the runbook to proceed.
- Execution of another runbook. To define this operation, you need to choose the runbook.
A runbook can include only one execution of a given runbook. For example, if you added the action "execute Runbook A", you can add the action "execute Runbook B", but cannot add another action "execute Runbook A".

Note In this product version a user has to perform a failback manually. A runbook shows the prompt when it is required.

Action parameters

All operations with cloud servers have the following parameters:

- **Continue if already done** (enabled by default)
This parameter defines the runbook behavior when the required operation is already done (for example, a failover has already been performed or a server is already running). When enabled, the runbook issues a warning and proceeds. When disabled, the operation fails and the runbook fails.
- **Continue if failed** (disabled by default)
This parameter defines the runbook behavior when the required operation fails. When enabled, the runbook issues a warning and proceeds. When disabled, the operation fails and the runbook fails.

Completion check

You can add completion checks to the **Failover server** and **Start server** actions, to ensure that the server is available and provides the necessary services. If any of the checks fail, the action is considered failed.

- **Ping IP address**
The software will ping the production IP address of the cloud server until the server replies or the timeout expires, whichever comes first.
- **Connect to port** (443 by default)
The software will try to connect to the cloud server by using its production IP address and the port you specify, until the connection is established or the timeout expires, whichever comes first. This way, you can check if the application that listens on the specified port is running.

The default timeout is 10 minutes. You can change it if you wish.

24.1.2 Operations with runbooks

To access the list of operations, hover on a runbook and click the ellipsis icon. When a runbook is not running, the following operations are available:

- **Execute**
- **Edit**
- **Clone**
- **Delete**

Executing a runbook

Every time you click **Execute**, you are prompted for the execution parameters. These parameters apply to all failover and failback operations included in the runbook. The runbooks specified in the **Execute runbook** operations inherit these parameters from the main runbook.

- **Failover and failback mode**

Choose whether you want to run a test failover (by default) or a real (production) failover. The failback mode will correspond to the chosen failover mode.

- **Failover recovery point**

Choose the most recent recovery point (by default) or select a point in time in the past. If the latter is the case, the recovery points closest before the specified date and time will be selected for each server.

Stopping a runbook execution

During a runbook execution, you can select **Stop** in the list of operations. The software will complete all of the already started actions except for those that require user interaction.

Viewing the execution history

When you select a runbook on the **Runbooks** tab, the software displays the runbook details and execution history. Click the line corresponding to a specific execution to view the execution log.

Runbooks

Name ↑
Failback 3-2
Rb0 000
Runbook with ConfirmManualOperation
Runbook with ConfirmManualOperation
jk one server with checking port
New runbook (10)
Failover/Failback (centos-1) (Clone)
New runbook (9)
Runbook #009.
Runbook #010.

Rb0 000

Execute

Edit

Clone

Delete

Details

Name

Rb0 000

Description

-

Execution history

Start and end time	Result	Mode
Aug 14, 5:30 PM - Aug 14, 10:27 PM	Failed	Production
Aug 14, 5:23 PM - Aug 14, 5:25 PM	Failed	Production
Aug 4, 2:45 AM - Aug 4, 2:46 AM	Completed	Test
Jul 30, 4:18 PM - Jul 30, 4:18 PM	Completed	Test
Jul 30, 4:16 PM - Jul 30, 4:16 PM	Completed	Test

25 Antimalware and web protection

Antimalware protection in Cyber Protection provides you with the following benefits:

- Top protection on all the stages: proactive, active, and reactive.
- Four different antimalware technologies inside to provide the best of the breed multi-layered protection.
- Management of Microsoft Security Essentials and Windows Defender Antivirus.

25.1 Antivirus and Antimalware protection

The Antivirus and Antimalware protection module allows you to protect your Windows and macOS machines from all recent malware threats. Note that the Active Protection functionality that is part of the antimalware protection is not supported on macOS machines. See the full list of supported antimalware features: Supported features by operating system (p. 13).

Antivirus and Antimalware protection is supported and registered in Windows Security Center.

The following antimalware capabilities are available to you:

- Detection of malware in files in the real-time protection and on-demand modes (for Windows, macOS)
- Detection of malicious behavior in processes (for Windows)
- Blocking access to malicious URLs (for Windows)
- Placing dangerous files to the quarantine
- Adding trusted corporate applications to the allowlist

The Antivirus and Antimalware protection module provides you with two types of scanning:

- Real-time protection scan
- On-demand antimalware scan

Real-time protection

Real-time protection checks all the files that are being executed or opened on a machine to prevent malware threats.

Real-time protection cannot work in parallel with other antivirus solutions that also use real-time protection features in order to prevent potential compatibility and performance issues. The statuses of other installed antivirus solutions are determined through Windows Security Center. If the Windows machine is already protected by another antivirus solution, real-time protection is automatically turned off.

To enable real-time protection, disable or uninstall the other antivirus solution. Real-time protection can replace Windows Defender real-time protection automatically.

You can choose one of the following types of scanning:

- On-access detection means that the antimalware program runs in the background and actively and constantly scans your machine system for viruses and other malicious threats for the entire duration that your system is powered on. Malware will be detected in both cases when a file is being executed and during various operations with the file such as opening it for reading/editing.

- On-execution detection means that only executable files will be scanned at the moment they are run to ensure they are clean and will not cause any damage to your machine or data. Copying of an infected file will remain unnoticed.

On-demand antimalware scan

Antimalware scanning is performed according to a schedule.

You can monitor the results of antimalware scanning in **Dashboard > Overview > Recently affected** (p. 391) widget.

25.1.1 Antivirus and Antimalware protection settings

To learn how to create a protection plan with the Antivirus and antimalware protection module, refer to "Creating a protection plan (p. 95)".

The following settings can be specified for the Antivirus and antimalware protection module.

Active Protection

Active Protection protects a system from ransomware and cryptocurrency mining malware. Ransomware encrypts files and demands a ransom for the encryption key. Cryptomining malware performs mathematical calculations in the background, thus stealing the processing power and network traffic.

Active Protection is available for machines running Windows 7 and later, Windows Server 2008 R2 and later. Agent for Windows must be installed on the machine.

Active Protection is available for agents starting with version 12.0.4290. To update an agent, follow the instructions in "Updating agents" (p. 77).

How it works

Active Protection monitors processes running on the protected machine. When a third-party process tries to encrypt files or mine cryptocurrency, Active Protection generates an alert and performs additional actions, if those are specified by the configuration.

In addition, Active Protection prevents unauthorized changes to the backup software's own processes, registry records, executable and configuration files, and backups located in local folders.

To identify malicious processes, Active Protection uses behavioral heuristics. Active Protection compares the chain of actions performed by a process with the chains of events recorded in the database of malicious behavior patterns. This approach enables Active Protection to detect new malware by its typical behavior.

Default setting: **Enabled**.

Active Protection settings

In **Action on detection**, select the action that the software will perform when detecting a ransomware activity, and then click **Done**.

You can select one of the following:

- **Notify only**
The software will generate an alert about the process.
- **Stop the process**

The software will generate an alert and stop the process.

- **Revert using cache**

The software will generate an alert, stop the process, and revert the file changes by using the service cache.

Default setting: **Revert using cache**.

Behavior detection

Acronis Cyber Protection protects your system by using behavioral heuristics to identify malicious processes: it compares the chain of actions performed by a process with the chains of actions recorded in the database of malicious behavior patterns. Thus, a new malware is detected by its typical behavior.

Default setting: **Enabled**.

Behavior detection settings

In **Action on detection**, select the action that the software will perform when detecting a malware activity, and then click **Done**.

You can select one of the following:

- **Notify only**

The software will generate an alert about the process suspected of malware activity.

- **Stop the process**

The software will generate an alert and stop the process suspected of malware activity.

- **Quarantine**

The software will generate an alert, stop the process, and move the executable file to the quarantine folder.

Default setting: **Quarantine**.

Exploit prevention

Exploit prevention detects and prevents infected processes from spreading and exploiting the software vulnerabilities on Windows systems. When an exploit is detected, the software can generate an alert and stop the process suspected of exploit activities.

Exploit prevention is available only with agent versions 20.08 or later.

Default setting: **Enabled** for newly created protection plans, and **Disabled** for existing protection plans, created with previous agent versions.

Exploit prevention settings

You can select what should the program do when an exploit is detected, and which exploit prevention methods are applied by the program.

Under **Enabled Action on detection**, select what to do when an exploit is detected, and then click **Done**.

- **Notify only**

The software will generate an alert about the process suspected of malware activity.

- **Stop the process**

The software will generate an alert and stop the process suspected of malware activity.

Default setting: **Stop the process**

Under **Enabled exploit prevention techniques**, enable or disable the methods that you want to be applied, and then click **Done**.

You can select one of the following:

- **Memory protection**
Detects and prevents suspicious modifications of the execution rights on memory pages. Malicious processes apply such modifications to page properties, to enable the execution of shellcodes from non-executable memory areas like stack and heaps.
- **Privilege escalation protection**
Detects and prevents attempts for elevation of privileges made by an unauthorized code or application. Privilege escalation is used by malicious code to gain full access of the attacked machine, and then perform critical and sensitive tasks. Unauthorized code is not allowed to access critical system resources or modify system settings.
- **Code injection protection**
Detects and prevents malicious code injection into remote processes. Code injection is used to hide malicious intent of an application behind clean or benign processes, to evade detection by antimalware products.

Default setting: **All methods are enabled**.

Note Processes that are listed as trusted processes in the Exclusions list will not be scanned for exploits.

Self-protection

Self-protection prevents unauthorized changes to the software's own processes, registry records, executable and configuration files, and backups located in local folders. We do not recommend disabling this feature.

Default setting: **Enabled**.

Allowing processes to modify backups

The **Allow specific processes to modify backups** option is effective when **Self-protection** is enabled.

It applies to files that have extensions .tibx, .tib, .tia, and are located in local folders.

This option lets you specify the processes that are allowed to modify the backup files, even though these files are protected by self-protection. This is useful, for example, if you remove backup files or move them to a different location by using a script.

If this option is disabled, the backup files can be modified only by processes signed by the backup software vendor. This allows the software to apply retention rules and to remove backups when a user requests this from the web interface. Other processes, no matter suspicious or not, cannot modify the backups.

If this option is enabled, you can allow other processes to modify the backups. Specify the full path to the process executable, starting with the drive letter.

Default setting: **Disabled**.

Network folder protection

The **Protect network folders mapped as local drives** option defines whether Active protection protects from local malicious processes network folders that are mapped as local drives.

This option applies to folders shared via SMB or NFS protocols.

If a file was originally located on a mapped drive, it cannot be saved to the original location when extracted from the cache by the **Revert using cache** action. Instead, it will be saved to the folder specified in this option's settings. The default folder is **C:\ProgramData\Acronis\Restored Network Files**. If this folder does not exist, it will be created. If you want to change this path, specify a local folder. Network folders, including folders on mapped drives, are not supported.

Default setting: **Enabled**.

Server-side protection

This option defines whether Active protection protects network folders that are shared by you from the external incoming connections from other servers in the network that may potentially bring threats.

Default setting: **Disabled**.

Setting trusted and blocked connections

On the **Trusted** tab, you can specify the connections that are allowed to modify any data. You should define the user name and IP address.

On the **Blocked** tab, you can specify the connections that will not be able to modify any data. You should define the user name and IP address.

Cryptomining process detection

This option defines whether Active protection detects potential cryptomining malware.

Cryptomining malware degrades performance of useful applications, increases electricity bills, may cause system crashes and even hardware damage due to abuse. We recommend that you add cryptomining malware to the **Harmful** processes list to prevent it from running.

Default setting: **Enabled**.

Cryptomining process detection settings

In **Action on detection**, select the action that the software will perform when a cryptomining activity is detected, and then click **Done**.

You can select one of the following:

- **Notify only**
The software generates an alert about the process suspected of cryptomining activities.
- **Stop the process**
The software generates an alert and stops the process suspected of cryptomining activities.

Default setting: **Stop the process**.

Real-time protection scan

Real-time protection scan constantly checks your machine system for viruses and other threats for the entire time that you system is powered on.

Default setting: **Enabled**.

Configuring the action on detection for real-time protection

In **Action on detection**, select the action that the software will perform when a virus or other malicious threat is detected, and then click **Done**.

You can select one of the following:

- **Block and notify**
The software blocks the process and generates an alert about the process suspected of malware activities.
- **Quarantine**
The software generates an alert, stops the process, and moves the executable file to the quarantine folder.

Default setting: **Quarantine**.

Configuring the scan mode for real-time protection

In **Scan mode**, select the action that the software will perform when a virus or other malicious threat is detected, and then click **Done**.

You can select one of the following:

- **Smart on-access** – Monitors all system activities and automatically scans files when they are accessed for reading or writing, or whenever a program is launched.
- **On-execution** – Automatically scans only executable files when they are launched to ensure that they are clean and will not cause any damage to your computer or data.

Default setting: **Smart on-access**.

Schedule scan

You can define schedule according to which your machine will be checked for malware. Enable the **Schedule scan** option.

Default setting: **Enabled**.

Action on detection:

- **Quarantine**
The software generates an alert and moves the executable file to the quarantine folder.
- **Notify only**
The software generates an alert about the process that is suspected to be malware.

Default setting: **Quarantine**.

Scan mode:

- **Full**
The full scan takes much longer to finish in comparison to the quick scan because every file will be checked.
- **Quick**
The quick scan only scans the common areas where malware normally resides on the machine.

You can schedule both **Quick** and **Full** scan in one protection plan.

Default setting: **Quick** and **Full** scan are scheduled.

Schedule the task run using the following events:

- **Schedule by time** – The task will run according to the specified time.
- **When user logs in to the system** – By default, login of any user will initiate a task run. You can this setting so that only a specific user account can trigger the task.
- **When user logs off the system** – By default, logoff of any user will make the task run. You can this setting so that only a specific user account can trigger the task.

***Note** The task will not run at system shutdown. Shutting down and logging off are different actions.*

- **On the system startup** – The task will run when the operating system starts.
- **On the system shutdown** – The task will run when the operating system shuts down.

Default setting: **Schedule by time**.

Schedule type:

- **Monthly** – Select the months and the weeks or days of the month when the task will run.
- **Daily** – Select the days of the week when the task will run.
- **Hourly** – Select the days of the week, repetition number, and the time interval in which the task will run.

Default setting: **Daily**.

Start at – Select the exact time when the task will run.

Run within a date range – Set a range in which the configured schedule will be effective.

Start conditions – Define all the conditions that must be simultaneously met for the task to run. They are similar to the start conditions for the Backup module that are described in "Start conditions (p. 131)". Also, the following additional start conditions can be defined:

- **Distribute task start time within a time window** – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM.
- **If the machine is turned off, run missed tasks at the machine startup**
- **Prevent the sleep or hibernate mode during task running** – This option is effective only for machines running Windows.
- **If start conditions are not met, run the task anyway after** – Specify the period after which the task will run, regardless of the other start conditions.

Scan only new and changed files – only newly created and modified files will be scanned.

Default setting: **Enabled**.

When scheduling a **Full scan**, you have two additional options:

Scan archive files

Default setting: **Enabled**.

- **Max recursion depth**
How many levels of embedded archives can be scanned. For example, MIME document > ZIP archive > Office archive > document content.
Default setting: **16**.

- **Max size**
Maximum size of an archive file to be scanned.
Default setting: **Unlimited**.

Scan removable drives

Default setting: **Disabled**.

- **Mapped (remote) network drives**
- **USB storage devices** (such as pens and external hard-drives)
- **CDs/DVDs**

Quarantine

Quarantine (p. 345) is a folder for keeping suspicious (probably infected) or potentially dangerous files isolated.

Remove quarantined files after – Defines the period in days after which the quarantined files will be removed.

Default setting: **30 days**.

Exclusions

To minimize the resources used by the heuristic analysis and to eliminate the so-called false positives when a trusted program is considered as ransomware, you can define the following settings:

On the **Trusted** tab, you can specify:

- Processes that will never be considered as malware. Processes signed by Microsoft are always trusted.
- Folders in which file changes will not be monitored.
- Files and folders in which the scheduled scan will not be performed.

On the **Blocked** tab, you can specify:

- Processes that will always be blocked. These processes will not be able to start as long as Active Protection is enabled on the machine.
- Folders in which any processes will be blocked

Default setting: No exclusions are defined by default.

25.2 Active Protection

Active Protection protects a system from ransomware and cryptocurrency mining malware. Ransomware encrypts files and demands a ransom for the encryption key. Cryptomining malware performs mathematical calculations in the background, thus stealing the processing power and network traffic.

In Cyber Backup Standard edition, Active Protection is a separate module in the protection plan. Thus, it can be configured separately and applied to different devices or group of devices.

The Active Protection module has the following settings:

- Action on detection
- Self-protection

- Network folder protection
- Server-side protection
- Cryptomining process detection
- Exclusions

In all other editions of the Cyber Protection service, Active Protection is part of the Antivirus and Antimalware protection module.

To learn more about Active Protection and its settings, refer to "Antivirus and Antimalware protection settings (p. 329)".

Active Protection is available for machines running Windows 7 and later, Windows Server 2008 R2 and later. Agent for Windows must be installed on the machine.

Active Protection is available for agents starting with version 12.0.4290. To update an agent, follow the instructions in "Updating agents" (p. 77).

25.3 Windows Defender Antivirus

Windows Defender Antivirus is a built-in antimalware component of Microsoft Windows that is delivered starting from Windows 8.

The Windows Defender Antivirus (WDA) module allows you to configure Windows Defender Antivirus security policy and track its status via the Cyber Protection service console.

This module is applicable for the machines on which Windows Defender Antivirus is installed.

Schedule scan

Specify the schedule for scheduled scanning.

Scan mode:

- **Full** – a full check of all files and folders additionally to the items scanned in the quick scan. It required more machine resources for execution compared to the quick scan.
- **Quick** – a quick check of the in-memory processes and folders where malware is typically found. It required less machine resources for execution.

Define the time and day of week when the scan will be performed.

Daily quick scan – define the time for the daily quick scan.

You can set the following options depending on your needs:

Start the scheduled scan when the machine is on but not in use

Check for the latest virus and spyware definitions before running a scheduled scan

Limit CPU usage during the scan to

For more details about the WDA settings, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#scheduled-scans-settings>

Default actions

Define the default actions to be performed for the detected threats of different severity levels:

- **Clean** – clean up the detected malware on a machine.
- **Quarantine** – put the detected malware in the quarantine folder but do not remove it.
- **Remove** – remove the detected malware from a machine.
- **Allow** – do not remove or quarantine the detected malware.
- **User defined** – a user will be prompted to specify the action to be performed with the detected malware.
- **No action** – no actions will be taken.
- **Block** – block the detected malware.

For more details about the WDA settings, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#default-actions-settings>

Real-time protection

Enable **Real-time protection** to detect and stop malware from installing or running on machines.

Scan all downloads – if selected, scanning is performed for all downloaded files and attachments.

Enable behavior monitoring – if selected, behavior monitoring will be enabled.

Scan network files – if selected, network files will be scanned.

Allow full scan on mapped network drives – if selected, mapped network drives will be fully scanned.

Allow email scanning – if enabled, the engine will parse the mailbox and mail files, according to their specific format, in order to analyze the mail bodies and attachments.

For more details about the WDA settings, refer to <https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#real-time-protection-settings>

Advanced

Specify the advanced scan settings:

- **Scan archive files** – include archived files such as .zip or .rar files into scanning.
- **Scan removable drives** – scan removable drives during full scans.
- **Create a system restore point** – in some cases an important file or registry entry could be removed as "false positive", then you will be able to recover from a restore point.
- **Remove quarantined files after** – define the period after which the quarantined files will be removed.
- **Send file samples automatically when a further analysis is required:**
 - **Always prompt** – you will be asked for confirmation before file sending.
 - **Send safe samples automatically** – most samples will be sent automatically except files that may contain personal information. Such files will require additional confirmation.
 - **Send all samples automatically** – all samples will be sent automatically.
- **Disable Windows Defender Antivirus GUI** – if selected, the WDA user interface will not be available to a user. You can manage the WDA policies via Cyber Protection service console.
- **MAPS (Microsoft Active Protection Service)** – online community that helps you choose how to respond to potential threats.

- **I don't want to join MAPS** – no information will be sent to Microsoft about the software that was detected.
- **Basic membership** – basic information will be sent to Microsoft about the software that was detected.
- **Advanced membership** – more detailed information will be sent to Microsoft about the software that was detected.

For more details, refer to

<https://www.microsoft.com/security/blog/2015/01/14/maps-in-the-cloud-how-can-it-help-your-enterprise/>

For more details about the WDA settings, refer to

<https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#advanced-settings>

Exclusions

You can define the following files and folders to be excluded from scanning:

- **Processes** – any file that the defined process reads from or writes to will be excluded from scanning. You need to define a full path to the executable file of the process.
- **Files and folders** – the specified files and folders will be excluded from scanning. You need to define a full path to a folder or file, or define the file extension.

For more details about the WDA settings, refer to

<https://docs.microsoft.com/en-us/sccm/protect/deploy-use/endpoint-antimalware-policies#exclusion-settings>

25.4 Microsoft Security Essentials

Microsoft Security Essentials is a built-in antimalware component of Microsoft Windows that is delivered with Windows versions earlier than 8.

The Microsoft Security Essentials module allows you to configure Microsoft Security Essentials security policy and track its status via the Cyber Protection service console.

This module is applicable for the machines on which Microsoft Security Essentials is installed.

The Microsoft Security Essentials settings are almost the same as Microsoft Windows Defender Antivirus except the absence of the real-time protection settings and inability to define exclusions via the Cyber Protection service console.

25.5 URL filtering

Malware is often distributed by malicious or infected sites and uses the so called Drive-by download method of infection.

The URL filtering functionality allows you to protect machines from threats like malware and phishing coming from the Internet. You can protect your organization by blocking user access to the websites that may have malicious content. The URL filtering database also includes data about the websites having disputed information about COVID-19, scam and phishing URLs. Thus, such websites will be automatically blocked by the system when a user tries to open them.


The URL filtering also allows you to control web usage to comply with the external regulations and internal company policies. You can configure access to the websites depending on the category they

relate to. The URL filtering supports currently 44 website categories and allows to manage access to them.

Currently, the HTTP/HTTPS connections on Windows machines will be checked by the protection agent.

How it works

A user enters a URL link in a browser. The Interceptor gets the link and sends it to the protection agent. The agent gets the URL, parses it, and then checks the verdict. The Interceptor redirects a user to the page with the message with available actions to manually proceed to the requested page.




URL filtering configuration workflow

Generally, the URL filtering configuration consists of the following steps:

1. You create a protection plan (p. 95) with the enabled **URL filtering** module.
2. Specify the URL filtering settings (see below).
3. Assign the protection plan to the machines.

To check which URLs have been blocked, go to **Dashboard > Alerts**.



URL filtering settings

The following settings can be specified for the URL filtering module.

Malicious website access

Specify which action will be performed when a user opens a malicious website:

- **Block** – block access to the malicious website. A user will not be able to access the website and a warning alert will be generated.
- **Always ask user** – ask a user whether to proceed to the website anyway or return back.

Categories to filter

There are 44 website categories for which you can configure access:

- **Allow** – allow access to websites related to the selected category.
- **Deny** – deny access to websites related to the selected category.

By default all categories are allowed.

Show all notifications for blocked URLs by categories – if enabled, you will get all notifications shown in the tray for blocked URLs by categories. If a website has several sub-domains, then the system also generates notifications for them, therefore the number of notifications may be big.

In the table below, you can find category descriptions:

	Website category	Description
1	Advertising	This category covers domains whose main purpose is to serve advertisements.
2	Message boards	This category covers forums, discussion boards, and question-answer type websites. This category does not cover the specific sections on company websites where customers ask questions.

3	Personal websites	This category covers personal websites, as well as all types of blogs: individual, group, and even company ones. A blog is a journal published on the World Wide Web. It consists of entries (“posts”), typically displayed in reverse chronological order so that the most recent post appears first.
4	Corporate/business websites	This is a broad category that covers corporate websites that typically do not belong to any other category.
5	Computer software	This category covers websites offering computer software, typically either open-source, freeware, or shareware. It may also cover some online software stores.
6	Medical drugs	This category covers websites related to medicine/alcohol/cigars that have discussions on the use or selling of (legal) medical drugs or paraphernalia, alcohol, or tobacco products. Note that illegal drugs are covered in the Narcotics category.
7	Education	This category covers websites belonging to official educational institutions, including those that are outside of the .edu domain. It also includes educational websites, such as an encyclopedia.
8	Entertainment	This category covers websites that provide information related to artistic activities and museums, as well as websites that review or rate content such as movies, music, or art.
9	File sharing	This category covers file-sharing websites where a user can upload files and share them with others. It also covers torrent-sharing websites and torrent trackers.
10	Finance	This category covers websites belonging to all banks around the world that provide online access. Some credit unions and other financial institutions are covered as well. However, some local banks may be left uncovered.
11	Gambling	This category covers gambling websites. These are the “online casino” or “online lottery” type website, which typically requires payment before a user can gamble for money in online roulette, poker, blackjack, or similar games. Some of them are legitimate, meaning there is a chance to win; and some are fraudulent, meaning that there is no chance to win. It also detects “beating tips and cheats” websites that describe the ways to make money on gambling and online lottery websites.
12	Games	This category covers websites that provide online games, typically based on Adobe Flash or Java applets. It does not matter for detection whether the game is free or requires a subscription, however, casino-style websites are detected in the Gambling category. This category does not cover: ▪ Official websites of companies that develop video games (unless they produce online games) ▪ Discussion websites where games are discussed ▪ Websites where non-online games can be downloaded (some of them are covered in the Illegal category) ▪ Games that require a user to download and run an executable, like World of Warcraft; those can be prevented by different means like a firewall
13	Government	This category covers government websites, including government institutions, embassies, and office websites.

14	Hacking	This category covers websites that provide the hacking tools, articles, and discussion platforms for hackers. It also covers websites offering exploits for common platforms that facilitate Facebook or Gmail account hacking.
15	Illegal activities	This category is a broad category related to hate, violence and racism, and it is intended to block the following categories of websites: ▪ Websites belonging to terrorist organizations ▪ Websites with racist or xenophobic content ▪ Websites discussing aggressive sports, and/or promoting violence
16	Health and fitness	This category covers websites associated with medical institutions, websites related to disease prevention and treatment, websites that offer information or products about weight loss, diets, steroids, anabolic or HGH products, as well as websites providing information on plastic surgery.
17	Hobbies	This category covers websites that present resources related to activities typically performed during an individual's free time, such as collecting, arts and crafts, and cycling.
18	Web hosting	This category covers free and commercial website hosting services that allow private users and organizations to create and publish web pages.
19	Illegal downloads	This category covers websites related to software piracy, including: ▪ Peer-to-peer (BitTorrent, emule, DC++) tracker websites that are known in helping to distribute copyrighted content without the copyright holder's consent ▪ Warez (pirated commercial software) websites and discussion boards ▪ Websites providing users with cracks, key generators, and serial numbers to facilitate the use of software illegally Some of these websites may also be detected as pornography or alcohol/cigars, since they often use porn or alcohol advertisements to earn money.
20	Instant messaging	This category covers instant messaging and chat websites that allow users to chat in real-time. It will also detect yahoo.com and gmail.com since they both contain an embedded instant messenger service.
21	Jobs/employment	This category covers websites presenting job boards, job-related classified advertisements, and career opportunities, as well as aggregators of such services. It does not cover recruiting agencies or the "jobs" pages on regular company websites.
22	Mature content	This category covers the content that was labeled by a website creator as requiring a mature audience. It covers a wide range of websites from the Kama Sutra book and sex education websites, to hardcore pornography.
23	Narcotics	This category covers websites sharing information about recreational and illegal drugs. This category also covers websites covering development or growing drugs.
24	News	This category covers news websites that provide text and video news. It strives to cover both global and local news websites; however, some small local news websites may not be covered.

25	Online dating	This category covers online dating websites – paid and free - where users can search for other people by using some criteria. They may also post their profiles to let others search them. This category includes both free and paid online dating websites. Because most of the popular social networks can be used as online dating websites, some popular websites like Facebook are also detected in this category. It's recommended to use this category with the Social networks category.
26	Online payments	This category covers websites offering online payments or money transfers. It detects popular payment websites like PayPal or Moneybookers. It also heuristically detects the webpages on the regular websites that ask for the credit card information, allowing detection of hidden, unknown, or illegal online stores.
27	Photo sharing	This category covers photo-sharing websites whose primary purpose is to let users upload and share photos.
28	Online stores	This category covers known online stores. A website is considered an online store if it sells goods or services online.
29	Pornography	This category covers websites containing erotic content and pornography. It includes both paid and free websites. It covers websites that provide pictures, stories, and videos, and it will also detect pornographic content on mixed-content websites.
30	Portals	This category covers websites that aggregate information from multiple sources and various domains, and that usually offer features such as search engines, e-mail, news, and entertainment information.
31	Radio	This category covers websites that offer Internet music streaming services, from online radio stations to websites that provide on-demand (free or paid) audio content.
32	Religion	This category covers websites promoting religion or a sect. It also covers the discussion forums related to one or multiple religions.
33	Search engines	This category covers search engine websites, such as Google, Yahoo, and Bing.
34	Social networks	This category covers social network websites. This includes MySpace.com, Facebook.com, Bebo.com, etc. However, specialized social networks, like YouTube.com, will be listed in the Video/Photo category.
35	Sport	This category covers websites that offer sports information, news, and tutorials.
36	Suicide	This category covers websites promoting, offering, or advocating suicide. It does not cover suicide prevention clinics.
37	Tabloids	This category is mainly designed for soft pornography and celebrity gossip websites. A lot of the tabloid-style news websites may have subcategories listed here. Detection for this category is also based on heuristics.
38	Waste of time	This category covers websites where individuals tend to spend a lot of time. This can include websites from other categories such as social networks or entertainment.
39	Traveling	This category covers websites that present travel offers and travel equipment, as well as travel destination reviews and ratings.

40	Videos	This category covers websites that host various videos or photos, either uploaded by users or provided by various content providers. This includes websites like YouTube, Metacafe, Google Video, and photo websites like Picasa or Flickr. It will also detect videos embedded in other websites or blogs.
41	Violent cartoons	This category covers websites discussing, sharing, and offering violent cartoons or manga that may be inappropriate for minors due to violence, explicit language, or sexual content. This category doesn't cover the websites that offer mainstream cartoons such as "Tom and Jerry".
42	Weapons	This category covers websites offering weapons for sale or exchange, manufacture, or usage. It also covers the hunting resources and the usage of air and BB guns, as well as melee weapons.
43	Email	This category covers websites that provide email functionality as a web application.
44	Web proxy	This category covers websites that provide web proxy services. This is a "browser inside a browser" type website when a user opens a web page, enters the requested URL into a form, and clicks "Submit". The web proxy site downloads the actual page and shows it inside the user browser. These are the following reasons this type is detected (and might need to be blocked): ▪ For anonymous browsing. Since requests to the destination web server are made from the proxy web server, only its IP address is visible and if the server administrators trace the user, the trace will end on web proxy – which may or may not keep logs necessary to locate the original user. ▪ For location spoofing. User IP addresses are often used for profiling the service by the source location (some national government websites may only be available from local IP addresses), and using those services might help the user to spoof their true location. ▪ For accessing prohibited content. If a simple URL filter is used, it will only see the web proxy URLs and not the actual servers that the user visits. ▪ For avoiding company monitoring. A business policy might require monitoring employee Internet usage. By accessing everything through a web proxy, a user might escape monitoring that will not provide correct information. Since the SDK analyzes the HTML page (if provided), and not just URLs, for some categories the SDK will still be able to detect the content. Other reasons, however, cannot be avoided just by using the SDK.

Exclusions

URLs that are known as safe can be added to the list of trusted URLs. URLs that represent a threat can be added to the list of blocked URLs.

To add a domain to the trusted URLs, click **Add** on the **Trusted** tab and specify the URL by using a specific domain name or IP address.

To add a domain to the blocked URLs, click **Add** on the **Blocked** tab and specify the URL by using a specific domain name or IP address.

Note All addresses from the domain that you entered will be treated as trusted or blocked. For example, if you entered xyz.com as a trusted domain, all paths or sub-domains under xyz.com are treated as trusted.

25.6 Quarantine

Quarantine is a special isolated folder on a machine's hard disk where the suspicious files detected by Antivirus and Antimalware protection are placed to prevent further spread of threats.

Quarantine allows you to review suspicious and potentially dangerous files from all machines and decide whether they should be removed or restored. The quarantined files are automatically removed if the machine is removed from the system.

How do files get into the quarantine folder?

1. You configure the protection plan and define the default action for infected files – to place in Quarantine.
2. The system during the scheduled or on-access scanning detects malicious files, places them in the secure folder - Quarantine.
3. The system updates the quarantine list on machines.
4. Files are automatically cleaned up from the quarantine folder after the time period defined in the **Remove quarantined files after** setting in the protection plan.


Managing quarantined files

To manage the quarantined files, go to **Antimalware protection > Quarantine**. You will see a list with quarantined files from all machines.

Name	Description
File	The file name.
Date quarantined	The date and time when the file was placed in Quarantine.
Device	The device on which the infected file was found.
Threat name	The threat name.
Protection plan	The protection plan according to which the suspicious file was placed in Quarantine.

You have two possible actions with quarantined files:

- **Delete** – permanently remove a quarantined file from all machines.
- **Restore** – restore a quarantined file to the original location without any modifications. If currently there is a file with the same name in the original location, then it will be overwritten with the restored file. Note that the restored file will be added to the allowlist and skipped during further antimalware scans.



Quarantined files		
File	Date quarantined	Device
test-malware.exe	Oct 23, 2019 1:50 PM	Win-10-MC-red
test-2-malware.exe	Oct 23, 2019 1:56 PM	Win-10-MC-red
358a5079b824548ef877cf89d3e4b5284e780edc4de8a450f3e51878d1290eca	Oct 23, 2019 2:00 PM	Win-10-MC-red
240387329dee4f03f98a89a2eff9b730dcb61fc614cdac24129da54442762	Oct 23, 2019 2:00 PM	Win-10-MC-red

Quarantine location on machines

The default location for quarantined files is:

For a Windows machine: %ProgramData%\%product_name%\Quarantine

For a Mac/Linux machine: /usr/local/share/%product_name%/quarantine

The quarantine storage is under the service provider's self-defense protection.

25.7 Corporate allowlist

Companies usually have their corporate specific applications that can be recognized and detected as false positive by antivirus solutions. It is also time-consuming to add manually the trusted applications to the allowlist.

Cyber Protection can automate the process of adding such applications to the allowlist. Backups are scanned by the Antivirus and Antimalware protection module and the scanned data are analyzed to allowlist such applications and prevent the false positive detections.

Adding trusted applications to the allowlist on a company-wide level enables to increase the further scanning performance. The allowlist is composed for each customer basing only on their data.

Automatic adding to the allowlist

1. First, you should run cloud scanning of backups on at least two machines. This can be done by means of backup scanning plans (p. 377).
2. Then, you should enable **Automatic generation of allowlist** in the allowlist settings.

Manual adding to the allowlist

Manual adding of applications to the allowlist is available only if the **Automatic generation of allowlist** option is enabled.

1. In the service console, go to **Antimalware protection > Allowlist**.
2. Click **Add file**.
3. Specify the path to the file and click **Add**.

Allowlist settings

You can enable the **Automatic generation of allowlist** option.

Then, you can specify one of the three trusted rule levels that define the aggression level of heuristics:

- **Low** – corporate applications will be added to the allowlist only after a significant amount of time and checks are done. Such applications are more trusted. However, this approach increases the possibility of false positive detections. The criteria to consider a file as clear and trusted are high.
- **Default** – corporate applications will be added to the allowlist according to the recommended protection level, to reduce possible false positive detections. The criteria to consider a file as clear and trusted are medium.
- **High** – corporate applications will be added to the allowlist faster, to reduce possible false positive detections. However, this does not guarantee that software is clean and it might later be recognized as suspicious/malware. The criteria to consider a file as clear and trusted are low.


25.8 Antimalware scan of backups

The backup scanning functionality allows you to prevent restoring infected files from backups. By using this functionality, you can check if your backups are clean (not infected by malware). The backup scanning functionality is supported only for Windows operating systems.

Backup scanning is performed by the cloud agent in the environment outside of an end-user machine – in the Acronis cloud. Every new backup scanning plan creates a new scanning task, the task is put in the common queue for the current data center and processed according to its order in the queue. The time required for scanning depends on a backup size, thus, you may experience some delay after creating a backup scanning plan and its execution.

If the backup scanning was not performed, then the backups stay in the **Not scanned** status. After backup scanning was performed, the backups get one of the following statuses:

- **No malware**
- **Malware detected**



The backup scanning can be configured by using a backup scanning plan.

How to configure backup scanning in the cloud

Note the following:

- The supported backup types are "Entire machine" or "Disks/volumes" backups.
- Only volumes with the NTFS file system with GPT and MBR partitioning will be scanned.
- The supported backup location is cloud storage (currently, only Acronis hosted).
- The backups that have CDP recovery points (p. 118) can be selected for scanning but only regular recovery points (excluding CDP recovery points) will be scanned.
- When the CDP backup was selected for safe recovery of an entire machine, the machine will be safely recovered without the data in the CDP recovery point. To restore the CDP data, start the Files/folders recovery activity.

To configure backup scanning in the cloud, create a backup scanning plan (p. 377).

The results of backup scanning can be found on the dashboard in the "Backup scanning details (p. 390)" widget.

26 Protection of collaboration and communication applications

Zoom, Cisco Webex Meetings, Citrix Workspace, and Microsoft Teams are now widely used for video/web conferencing and communications. The Cyber Protection service allows you to protect your collaboration tools.

The protection configuration for Zoom, Cisco Webex Meetings, Citrix Workspace, and Microsoft Teams is similar. In the example below, we will consider configuration for Zoom.

To set up Zoom protection

1. Install the protection agent on the machine where the collaboration application is installed.
2. Log in to the service console and apply a protection plan (p. 101) that has one of the following modules enabled:
 - **Antivirus and Antimalware protection** (p. 329) (with the **Self-Protection** and **Active Protection** settings enabled) – if you have one of the Cyber Protect editions.
 - **Active Protection** (p. 335) (with the **Self-Protection** setting enabled) – if you have one of the Cyber Backup editions.
1. [Optional] For automatic update installation, configure the **Patch management** module (p. 354) in the protection plan.

As a result, your Zoom application will be under protection that includes the following activities:

- Installing Zoom client updates automatically
- Protecting Zoom processes from code injections
- Preventing suspicious operations by Zoom processes
- Protecting the "hosts" file from adding the domains related to Zoom

27 Vulnerability assessment and patch management

27.1 Supported Microsoft and third-party products

Supported Microsoft products

Windows OS

- Windows 7 (Enterprise, Professional, Ultimate)
- Windows 8
- Windows 8.1
- Windows 10

Windows Server OS

- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2

Microsoft Office and related components

- Microsoft Office 2019 (x64, x86)
- Microsoft Office 2016 (x64, x86)
- Microsoft Office 2013 (x64, x86)
- Microsoft Office 2010 (x64, x86)

Windows OS related components

- Internet Explorer
- Microsoft EDGE
- Windows Media Player
- .NET Framework
- Visual Studio and Applications
- Components of operating system

Server applications

- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2012
- Microsoft SQL Server 2014
- Microsoft SQL Server 2016
- Microsoft SQL Server 2017
- Microsoft SQL Server 2019
- Microsoft Exchange Server 2013
- Microsoft Sharepoint Server 2016
- Microsoft Sharepoint Server 2016

Supported third-party products for Windows OS

Remote work becomes more and more wide-spread across the world, therefore collaboration and communication tools, VPN clients are now important to be always up-to-date and checked on possible vulnerabilities. The Cyber Protection service supports the vulnerability assessment and patch management for such applications.

Collaboration and communication tools, VPN clients

- Microsoft Teams
- Zoom
- Skype
- Slack
- Webex
- NordVPN
- TeamViewer

For more information about the supported third-party products for Windows OS, refer to <https://kb.acronis.com/content/62853>.

27.2 Vulnerability assessment

Vulnerability assessment (VA) is a process of identifying, quantifying, and prioritizing found vulnerabilities in the system. The vulnerability assessment module allows you to scan your machines for vulnerabilities and to ensure that all of the installed applications and operating systems are up-to-date and work properly.

Currently, only Windows and Linux (CentOS 7/Virtuozzo/Acronis Cyber Infrastructure) machines are supported for VA scanning. For more details about configurations for Linux machines, refer to "Vulnerability assessment for Linux machines (p. 353)".

How it works

1. You create a protection plan (p. 95) with the enabled vulnerability assessment module, specify the VA settings (p. 351), and assign the plan to machines (p. 95).
2. The system, by schedule or on demand, sends a command to run the VA scanning to the protection agents installed on machines.
3. The agents get the command, start scanning machines for vulnerabilities, and generate the scanning activity.
4. After the VA scanning is completed, the agents generate the results and send them to the monitoring service.
5. The monitoring service processes the data from the agents and shows the results in the vulnerability assessment widgets (p. 387) and list of found vulnerabilities.
6. When you get a list of found vulnerabilities (p. 352), you can process it and decide which of the found vulnerabilities must be fixed.

You can monitor the results of the vulnerability assessment scanning in **Dashboard > Overview > Vulnerabilities / Existing vulnerabilities** (p. 387) widgets.

27.2.1 Vulnerability assessment settings

To learn how to create a protection plan with the vulnerability assessment module, refer to "Creating a protection plan (p. 95)". Vulnerability assessment scanning can be performed by schedule or on demand (by using the **Run now** action in a protection plan).

The following settings can be specified for the vulnerability assessment module.

What to scan

Define which software products you want to scan for vulnerabilities:

- Windows machines:
 - **Microsoft products**
 - **Windows third-party products** (for more information about the supported third-party products for Windows OS, refer to <https://kb.acronis.com/content/62853>)
- Linux machines:
 - **Scan Linux packages**

Schedule

Define the schedule according to which the vulnerability assessment scan will be performed on the selected machines:

Schedule the task run using the following events:

- **Schedule by time** – The task will run according to the specified time.
- **When user logs in to the system** – By default, login of any user will initiate a task run. You can this setting so that only a specific user account can trigger the task.
- **When user logs off the system** – By default, logoff of any user will make the task run. You can this setting so that only a specific user account can trigger the task.

Note *The task will not run at system shutdown. Shutting down and logging off are different actions.*

- **On the system startup** – The task will run when the operating system starts.
- **On the system shutdown** – The task will run when the operating system shuts down.

Default setting: **Schedule by time**.

Schedule type:

- **Monthly** – Select the months and the weeks or days of the month when the task will run.
- **Daily** – Select the days of the week when the task will run.
- **Hourly** – Select the days of the week, repetition number, and the time interval in which the task will run.

Default setting: **Daily**.

Start at – Select the exact time when the task will run.

Run within a date range – Set a range in which the configured schedule will be effective.

Start conditions – Define all the conditions that must be simultaneously met for the task to run. They are similar to the start conditions for the Backup module that are described in "Start conditions (p. 131)". Also, the following additional start conditions can be defined:

- **Distribute task start time within a time window** – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM.
- **If the machine is turned off, run missed tasks at the machine startup**
- **Prevent the sleep or hibernate mode during task running** – This option is effective only for machines running Windows.

If start conditions are not met, run the task anyway after – **Specify the period after which the task will run, regardless of the other start conditions.**

27.2.2 Managing found vulnerabilities

If the vulnerability assessment was performed at least once and some vulnerabilities were found, you will see them in **Software management > Vulnerabilities**. The list of vulnerabilities shows both vulnerabilities that have patches to be installed and those that do not have suggested patches. You can use the filter to show only vulnerabilities with patches.

Name	Description
Name	The name of vulnerability.
Affected products	Software products for which the vulnerabilities were found.
Machines	The number of affected machines.
Severity	The severity of found vulnerability. The following levels can be assigned according to the Common Vulnerability Scoring System (CVSS): ▪ Critical: 9 - 10 CVSS ▪ High: 7 - 9 CVSS ▪ Medium: 3 - 7 CVSS ▪ Low: 0 - 3 CVSS ▪ None
Patches	The number of appropriate patches.
Published	The date and time when the vulnerability was published in Common Vulnerabilities and Exposures (CVE).
Detected	The first date when an existing vulnerability was detected on machines.

You can find the description of found vulnerability by clicking its name in the list.

Acronis Cyber Cloud		Vulnerabilities				
Manage account		Filter Search				Loaded: 30 / Total: 82
ANTI-MALWARE PROTECTION						
SOFTWARE MANAGEMENT						
Patches						
Vulnerabilities						
BACKUP STORAGE						
REPORTS						
SETTINGS						
Powered by Acronis AnyData Engine						

☐	Name ↓	Affected products ↓	Machines ↓	Severity ↓	Patches ↓	
	CVE-2015-16723	Microsoft Windows 8.1	1	CRITICAL	2	
	CVE-2015-0016	Microsoft Windows 8.1	1	CRITICAL	1	
	CVE-2014-4073	Microsoft Windows 8.1	1	CRITICAL	1	
☐	CVE-2010-3190	Microsoft Visual Studio 2008	1	CRITICAL	1	
	CVE-2015-1756	Microsoft Windows 8.1	1	CRITICAL	1	
	CVE-2014-4121	Microsoft Windows 8.1	1	CRITICAL	1	
	CVE-2016-3236	Microsoft Windows 8.1	1	CRITICAL	1	
	CVE-2014-6324	Microsoft Windows 8.1	1	CRITICAL	1	

To start the vulnerability remediation process

1. In the service console, go to **Software management > Vulnerabilities**.
2. Select the vulnerability in the list and click **Install patches**. The vulnerability remediation wizard will open.
3. Select the patches to be installed on the selected machines. Click **Next**.
4. Select the machines that you want to install patches for.
5. Select if the machine reboot must be performed after patch installation:
 - **No** – reboot will never be initiated after the update installation.
 - **If required** – reboot is done only if it is required for applying the updates.
 - **Yes** – reboot will be always initiated after the updates. You can always specify the reboot delay.

Do not reboot until backup is finished – if the backup process is running, the machine reboot will be delayed until the backup is completed.

When ready, click **Install patches**.

As a result, the selected patches will be installed on the selected machines.

27.2.3 Vulnerability assessment for Linux machines

Vulnerability assessment is also supported for Linux machines. You can scan Linux machines for application-level and kernel-level vulnerabilities.

The following Linux distributions and versions are supported:

- Virtuozzo 7.0.11
- Virtuozzo 7.0.10 (320)
- Virtuozzo 7.0.9 (539)
- Virtuozzo 7.0.8 (524)
- CentOS 7.x
- Acronis Cyber Infrastructure 3.x
- Acronis Storage 2.4.0
- Acronis Storage 2.2.0

To configure the vulnerability assessment for Linux machines

1. Install Agent for Linux on the Acronis Cyber Infrastructure (or Virtuozzo) host or into a virtual machine with CentOS.
2. In the service console, create a protection plan (p. 95) and enable the **Vulnerability assessment** module.
3. Specify the vulnerability assessment settings:
 - **What to scan** – select **Scan Linux packages**.
 - **Schedule** – define the schedule for performing the vulnerability assessment.
4. Assign the plan to the machines (p. 95).

As a result of vulnerability assessment, you will get a list of found vulnerabilities (p. 352). You can process it and decide which of the found vulnerabilities must be fixed.

You can monitor the results of the vulnerability assessment in the **Dashboard > Overview > Vulnerabilities / Existing vulnerabilities** (p. 387) widgets.

27.3 Patch management

Patch management (PM) provides you with capabilities to manage patches/updates for applications and operating systems installed on your machines and keep your systems up-to-date. The patch management module allows you to automatically or manually approve update installation on your machines. Currently, only Windows machines are supported by the patch management functionality.


The patch management functionality allows you:

- To install OS-level and application level updates
- To approve patches manually or automatically
- To install patches on-demand and according to a schedule
- To precisely define which patches to apply by different criteria: severity, category, and approval status
- To perform pre-update backup to prevent possible unsuccessful updates
- To define the reboot option to be applied after patch installation

Cyber Protection introduces peer-to-peer technology to minimize network bandwidth traffic. You can choose one or more dedicated agents that will download updates from the Internet and distribute them among other agents in the network. All agents will also share updates with each other as peer-to-peer agents.

How it works

You can configure either automatic or manual patch approval. In the scheme below, you can see both automatic and manual patch approval workflows.



1. First, you need to perform at least one vulnerability assessment scan (p. 350) by using the protection plan with the **Vulnerability assessment** module enabled. After the scan was performed, the lists of found vulnerabilities (p. 352) and available patches are composed by the system.
2. Then, you can configure the automatic patch approval (p. 360) or use manual patch approval (p. 362) approach.
3. Define how to install patches – according to a schedule or on-demand. On-demand patch installation can be done in three ways according to your preferences:
 - Go to the list of patches (**Software management > Patches**) and install the necessary patches.

Acronis Cyber Cloud Patches									
Filter Search									
Loaded: 3 / Total: 3 Settings									
☐	Name ↑	Severity ↓	Product ↓	Installed versions	Version ↓	Microsoft KB ↓	Machines ↓	Approval status ↓	
☐	2020-03 Preview of Monthly Quality Rollup f.	MEDIUM	Windows Server ...	—	—	KB4541334	1	Not defined	⚙️
☐	Mozilla Thunderbird	MEDIUM	Thunderbird	68.5.0	68.6.0	—	1	Not defined	⚙️
☐	Notepad++ Team Notepad++	MEDIUM	Notepad++	7.8.4	7.8.5	—	1	Not defined	⚙️

- Go to the list of vulnerabilities (**Software management > Vulnerabilities**) and start the remediation process which includes patch installation as well.
- Go to the list of devices (**Devices > All devices**), select the particular machines that you want to update, and install patches on them.

You can monitor the results of the patch installation in **Dashboard > Overview > Patch installation history** widget.

27.3.1 Patch management settings

To learn how to create a protection plan with the patch management module, refer to "Creating a protection plan (p. 95)". By using the protection plan, you can specify what updates for Microsoft products and other third-party products for Windows OS to automatically install on the defined machines.

The following settings can be specified for the patch management module.

Microsoft products

To install the Microsoft updates on the selected machines, enable the **Update Microsoft products** option.

Select what updates you want to be installed:

- **All updates**
- **Only Security and Critical updates**
- **Updates of specific products:** you can define custom settings for different products. If you want to update specific products, for each product you can define which updates to install by category, severity, or approval status.

Updates of specific products ✕

	Products ↓	Category	Severity	Approval status
☐	Windows Server 2012 R2 L...	Custom	Custom	Custom
☒	Windows Server 2012 R2	ServicePacks, Upd...	Critical, High, Medi...	Approved
☒	Windows Server 2012	CriticalUpdates	Critical, High	Approved
☐	Windows Server 2016 and ...	—	—	—
☒	Windows Server 2016	SecurityUpdates	Critical	Approved

[Reset to default](#)

[Cancel](#)
[Save](#)

Windows third-party products

To install the third-party updates for Windows OS on the selected machines, enable the **Windows third-party products** option.

Select what updates you want to be installed:

- **Only last major updates** allows you to install the latest available version of the update.

- **Only last minor updates** allows you to install the minor version of the update.
- **Updates of specific products:** you can define custom settings for different products. If you want to update specific products, for each product you can define which updates to install by category, severity, or approval status.

Updates of specific products ✕

	Products ↓	Category	Severity	Approval
☐	Adobe Reader	Custom	Custom	Approved
☐	Adobe Flash Player for Chr...	—	—	—
☐	Adobe Flash Player for Fire...	—	—	—
☒	Oracle Java Runtime Envir...	Major updates	Critical	Approved
☒	Mozilla Firefox	Minor updates	All	Approved
☐	Google Chrome	—	—	—

Reset to default Cancel Save

Schedule

Define the schedule according to which the updates will be installed on the selected machines.

Schedule the task run using the following events:

- **Schedule by time** – The task will run according to the specified time.
- **When user logs in to the system** – By default, login of any user will initiate a task run. You can this setting so that only a specific user account can trigger the task.
- **When user logs off the system** – By default, logoff of any user will make the task run. You can this setting so that only a specific user account can trigger the task.

***Note** The task will not run at system shutdown. Shutting down and logging off are different actions.*

- **On the system startup** – The task will run when the operating system starts.
- **On the system shutdown** – The task will run when the operating system shuts down.

Default setting: **Schedule by time**.

Schedule type:

- **Monthly** – Select the months and the weeks or days of the month when the task will run.
- **Daily** – Select the days of the week when the task will run.
- **Hourly** – Select the days of the week, repetition number, and the time interval in which the task will run.

Default setting: **Daily**.

Start at – Select the exact time when the task will run.

Run within a date range – Set a range in which the configured schedule will be effective.

Start conditions – Define all the conditions that must be simultaneously met for the task to run. They are similar to the start conditions for the Backup module that are described in "Start conditions (p. 131)". Also, the following additional start conditions can be defined:

- **Distribute task start time within a time window** – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM.
- **If the machine is turned off, run missed tasks at the machine startup**
- **Prevent the sleep or hibernate mode during task running** – This option is effective only for machines running Windows.

If start conditions are not met, run the task anyway after – Specify the period after which the task will run, regardless of the other start conditions.

Reboot after update – define whether reboot is initiated after installing updates:

- **Never** – reboot will never be initiated after the updates.
- **If required** – reboot is done only if it is required for applying the updates.
- **Always** – reboot will be always initiated after the updates. You can always specify the reboot delay.

Do not reboot until backup is finished – if the backup process is running, the machine reboot will be delayed until the backup is completed.

Pre-update backup

Run backup before installing software updates – the system will create an incremental backup of machine before installing any updates on it. If there were no backups created earlier, then a full backup of machine will be created. It allows you to prevent such cases when the installation of updates was unsuccessful and you need to get back to the previous state. For the **Pre-update backup** option to work, the corresponding machines must have both the patch management and the backup module enabled in a protection plan and the items to back up – entire machine or boot+system volumes. If you select inappropriate items to back up, then the system will not allow you to enable the **Pre-update backup** option.

27.3.2 Managing list of patches

After vulnerability assessment scanning was done, you will find the available patches in **Software management > Patches**.

Name	Description
Name	The name of the patch
Severity	The severity of the patch: ▪ Critical ▪ High ▪ Medium ▪ Low ▪ None
Vendor	The vendor of the patch

Product	Product for which the patch is applicable
Installed versions	Product versions that are already installed
Version	Version of the patch
Category	The category to which the patch belongs: ▪ Critical update – broadly released fixes for specific problems addressing critical, non-security related bugs. ▪ Security update – broadly released fixes for specific products addressing security issues. ▪ Definition update – updates to virus or other definition files. ▪ Update rollup – cumulative set of hotfixes, security updates, critical updates, and updates packaged together for easy deployment. A rollup generally targets a specific area, such as security, or a specific component, such as Internet Information Services (IIS). ▪ Service pack – cumulative sets of all hotfixes, security updates, critical updates, and updates created since the release of the product. Service packs might also contain a limited number of customer-requested design changes or features. ▪ Tool – utilities or features that aid in accomplishing a task or set of tasks. ▪ Feature pack – new feature releases, usually rolled into products at the next release. ▪ Update – broadly released fixes for specific problems addressing non-critical, non-security related bugs. ▪ Application – patches for an application.
Microsoft KB	If the patch for Microsoft product, the KB article ID is provided
Release date	The date when the patch was released
Machines	Number of affected machines
Approval status	The approval status is mainly needed for automatic approval scenario and to be able to define in the protection plan which updates to install by status. You can define one of the following statuses for a patch: ▪ Approved – the patch was installed on at least one machine and validated as ok ▪ Declined – the patch is not safe and may corrupt a machine system ▪ Not defined – the patch status is unclear and should be validated
License agreement	▪ Read and accept ▪ Disagreed. If you disagree with the license agreement, then the patch status becomes Declined and it will not be installed
Vulnerabilities	The number of vulnerabilities. If you click on it, you will be redirected to the list of vulnerabilities.
Size	The average size of the patch
Language	The language which is supported by the patch
Vendor site	The official site of the vendor

27.3.3 Automatic patch approval

Automatic patch approval allows you to make the process of installing updates on machines easier. Let's consider the example how it works.

How it works

You should have two environments: test and production. The test environment is used for testing the patch installation and ensuring that they do not break anything. After you tested patch installation on the test environment, you can automatically install these safe patches on the production environment.

Configuring automatic patch approval

To configure automatic patch approval

1. For each vendor whose products you are planning to update, you must read and accept the license agreements. Otherwise, automatic patch installation will not be possible.
2. Configure the settings for automatic approval.
3. Prepare the protection plan (p. 95) (for example, "Test patching") with the enabled **Patch management** module and apply it to the machines in the test environment. Specify the following condition of patch installation: the patch approval status must be **Not defined**. This step is needed to validate the patches and check if the machines work properly after patch installation.
4. Prepare the protection plan (p. 95) (for example, "Production patching") with the enabled **Patch management** module and apply it to the machines in the production environment. Specify the following condition of patch installation: the patch status must be **Approved**.
5. Run the Test patching plan and check the results. The approval status for those machines that have no issues can be preserved as **Not defined** while the status for machines working incorrectly must be set to **Declined**.
6. According to the number of days set in the **Automatic approval** option, those patches that were **Not defined** will become **Approved**.
7. When the Production patching plan is launched, only those patches that are **Approved** will be installed on the production machines.

The manual steps are listed below.

Step 1. Read and accept the license agreements for the products that you want to update

1. In the service console, go to **Software management > Patches**.
2. Select the patch, then read and accept the license agreement.

Step 2. Configure the settings for automatic approval

1. In the service console, go to **Software management > Patches**.
2. Click **Settings**.
3. Enable the **Automatic approval** option and specify the number of days. This means that after the specified number of days starting from the first attempt of patch installation, the patches with the status **Not defined** will become **Approved** automatically.

For example, you specified 10 days. You performed the Test patching plan for test machines and installed patches. Those patches that broke the machines, you marked as **Declined** while the rest of patches stay as **Not defined**. After 10 days passed, the patches in the **Not defined** status will be automatically switched to **Approved**.

4. Enable the **Automatically accept the license agreements** option. This is needed for automatic license acceptance during patch installation, no confirmation is required from a user.

Step 3. Prepare the Test patching protection plan

1. In the service console, go to **Plans > Protection**.
2. Click **Create plan**.
3. Enable the **Patch management** module.
4. Define which updates to install for Microsoft and third-party products, schedule, and pre-update backup. For more details about these settings, refer to "Patch management settings (p. 356)".

Important: For all the products to be updated, define **Approval status** as **Not defined**. When the time to update comes, the agent will install only **Not defined** patches on the selected machines in the test environment.

Updates of specific products ✕

	Products ↓	Category	Severity	Approval status
☒		Custom	Custom	Not defined
☒	Active Directory Rights Ma...	CriticalUpdates, Se...	Critical	Not defined
☒	Antigen for Exchange/SMTP	None	All	Not defined
☒	ASP.NET Web Frameworks	Updates	Critical, High, Medi...	Not defined
☒	Azure File Sync agent upda...	None	All	Not defined
☒	Azure File Sync agent upda...	None	All	Not defined

Reset to default Cancel Save

Step 4. Prepare the Production patching protection plan

1. In the service console, go to **Plans > Protection**.
2. Click **Create plan**.
3. Enable the **Patch management** module.
4. Define which updates to install for Microsoft and third-party products, schedule, and pre-update backup. For more details about these settings, refer to "Patch management settings (p. 356)".

Important: For all the products to be updated, define **Approval status** as **Approved**. When the time to update comes, the agent will install only **Approved** patches on the selected machines in the production environment.

Updates of specific products

☒	Products ↓	Category	Severity	Approval status
☒	Active Directory Rights Ma...	Custom	Custom	Approved
☒	Antigen for Exchange/SMTP	CriticalUpdates, Se...	Critical	Approved
☒	ASP.NET Web Frameworks	All	All	Approved
☒	Azure File Sync agent upda...	Updates	Critical, High, Medi...	Approved
☒	Azure File Sync agent upda...	All	All	Approved
☒	Azure File Sync agent upda...	All	All	Approved

Reset to default

Cancel Save

Step 5. Run the Test patching protection plan and check the results

1. Run the Test patching protection plan (by schedule or on-demand).
2. After that, check which of the installed patches are safe and which are not.
3. Go to **Software management > Patches** and set the **Approval status** as **Declined** for those patches that are not safe.

27.3.4 Manual patch approval

The manual patch approval process is the following:

1. In the service console, go to **Software management > Patches**.
2. Select the patches that you want to install, then read and accept the license agreements.
3. Set **Approval status** to **Approved** for the patches that you approve for installation.
4. Create a protection plan with the enabled Patch management (p. 356) module. You can either configure the schedule or launch the plan on-demand by clicking **Run now** in the patch management module settings.

As a result, only the approved patches will be installed on the selected machines.

27.3.5 On-demand patch installation

On-demand patch installation can be done in three ways according to your preferences:

- Go to the list of patches (**Software management > Patches**) and install the necessary patches.
- Go to the list of vulnerabilities (**Software management > Vulnerabilities**) and start the remediation process which includes patch installation as well.
- Go to the list of devices (**Devices > All devices**), select the particular machines that you want to update, and install patches on them.

Let's consider patch installation from the list of patches:

1. In the service console, go to **Software management > Patches**.

2. Accept the license agreements for the patches that you want to install.
3. Select the patches that you want to install and click **Install**.
4. Select the machines on which patches must be installed.

If you want to have a rollback option in case of patch installation damages the system, select the **Run backup before installing software updates** option. The system checks right away if there is a protection plan with the Backup module enabled (the entire machine backup is required). If there is no such protection plan assigned to the machine, then such machines are marked with the red icon. You can unselect these machines and proceed.

5. Define whether reboot is initiated after installing patches:
 - **Never** – reboot will never be initiated after the patches.
 - **If required** – reboot is done only if it is required for applying the patches.
 - **Always** – reboot will be always initiated after the patches. You can always specify the reboot delay.

Do not reboot until backup is finished – if the backup process is running, the machine reboot will be delayed until the backup is completed.
6. Click **Install patches**.

The selected patches will be installed on the selected machines.

27.3.6 Patch lifetime in the list

To keep the list of patches up-to-date, go to **Software management > Patches > Settings** and specify the **Lifetime in list** option.

The **Lifetime in list** option defines how long will the detected available patch be kept in the list of patches. Generally, the patch is removed from the list if it is successfully installed on all the machines where its absence is detected or the defined time lapses.


- **Forever** – the patch always stays in the list.
- **7 days** – the patch is removed if after its first installation seven days passed.

For example, you have two machines where patches must be installed. One of them is online, another – offline. The patch was installed on the first machine. After 7 days, the patch will be removed from the list of patches even if it is not installed on the second machine because it was offline.
- **30 days** – the patch is removed if after its first installation thirty days passed.

28 Remote desktop access

28.1 Remote access (RDP and HTML5 clients)


Cyber Protection provides you with remote access capability. You can remotely connect and manage your end user machines. You can copy and paste text to and from the remote machine with the HTML5 client. With the RDP client you can copy and paste text as well as files. This allows you to easily assist your end users in resolving issues on their machines.



Prerequisites:

- A remote machine is registered in Cyber Protection and the protection agent is installed.
- The Cyber Protect quota exists or was already acquired for a machine.
- For RDP connections, the Remote Desktop Connection client is installed on a machine from which the connection is launched.

An RDP session can be established from both Windows and macOS machines. An HTML5 remote connection session can be established from any browser with HTML5 support.




The remote access functionality can be used for connections to Windows machines with the Windows Remote Desktop feature available. Thus, remote access cannot be used, for example, for a connection to Windows 10 Home and macOS systems.

To establish a connection from a macOS machine to a remote machine, ensure that the following applications are installed on the macOS machine:

- The Remote Desktop Connection client
- The Microsoft Remote Desktop application

How it works

When you try to connect to a remote machine, the system first checks that this machine has the Cyber Protect quota. Then, the system checks that the connection via the HTML5 or RDP client is possible. You initiate a connection via the RDP or HTML5 client. The system establishes a tunnel to the remote machine and checks that the remote desktop connections are enabled on the remote machine. Then, you enter the credentials and, if their validation is ok, get access to the machine.



How to connect to a remote machine

To connect to a remote machine, do the following:

1. In the service console, go to **Devices > All devices**.

2. Click on the machine to which you want to connect remotely and then click **Cyber Protection Desktop > Connect via RDP client / Connect via HTML5 client**.
3. [Optional, only for connection via RDP client] Download and install the Remote Desktop Connection Client. Initiate the connection to the remote machine.
4. Specify the login and password to access the machine and click **Connect**.

As a result, you are connected to the remote machine and can manage it.

How to run a remote assistance session

Remote assistance allows concurrent access to the same remote desktop session. For example, when you need to fix a problem on a remote user computer, you can use remote assistance to connect to the computer. The user and the remote administrator share one session and the user can share and reproduce an issue.

1. In the service console, go to **Devices > All devices**.
2. Click on the machine to which you want to connect remotely and then click **Cyber Protection Desktop > Run remote assistance**.
3. Copy the remote assistance session password and click **Connect**. If the session does not start, download and install the connectivity agent on your machine, and retry the connection.
4. If there are ongoing interactive sessions, click **Connect to session**.
5. Enter the remote assistance session password.

As a result, you have remote desktop access to the remote machine and can assist the user.

28.2 Sharing remote connection with end-users

Employees that are working from home may need to have access to their work machines, but it is possible that your organization may not have a configured VPN or other tools for remote connection.

The Cyber Protection service provides you with the capability to share an RDP link with end-users, thus providing them with the remote access to their machines.

How to enable the sharing remote connection functionality

To enable the sharing remote connection functionality, do the following:

1. In the service console, go to **Settings > Protection > Remote connection**.
2. Enable **Share remote desktop connection**.

As a result, the new option **Share remote connection** will appear in the right menu when you select a device.

How to share a remote connection with your users

To share a remote connection with a user, do the following:

1. In the service console, go to **Devices > All devices**.
2. Select the device to which you want provide the remote connection and then click **Cyber Protection Desktop > Share remote connection**.
3. Click **Get link**. In the opened window, copy the generated link. This link can be shared with a user who needs the remote access to the device. The link is valid for 10 hours.


After getting the link, you can share it via email or other means of communication with your user. When the user clicks this link, they will be redirected to the page where the connection type must be selected:

- Connect via RDP client. This connection will prompt downloading and installing the Remote Connection Client.
- Connect via HTML5 client. This connection does not require the installation of any RDP client on the user's machine. The user will be redirected to the login screen where the user's credentials to the remote machine have to be entered.

29 Remote wipe

Remote wipe allows a Cyber Protection service administrator and a machine owner to delete the data on a managed machine – for example, if it gets lost or stolen. Thus, any unauthorized access to sensitive information will be prevented.

Remote wipe is only available for machines running Windows 10. To receive the wipe command, the machine must be turned on and connected to the Internet.



To wipe data from a machine

1. In the service console, go to **Devices > All devices**.
2. Select the machine whose data you want to wipe.

Note You can wipe data from one machine at a time.

3. Click **Details**, and then click **Wipe data**.
If the machine that you selected is offline, the **Wipe data** option is inaccessible.
4. Confirm your choice.
5. Enter the credentials of this machine's local administrator, and then click **Wipe data**.

Tip You can check the details about the wiping process and who started it in **Dashboard > Activities**.

30 Smart protection

30.1 Threat feed


Acronis Cyber Protection Operations Center (CPOC) generates security alerts that are sent only to the related geographic regions. These security alerts provide information about malware, vulnerabilities, natural disasters, public health, and other types of global events that may affect your data protection. The threat feed informs you about all the potential threats and allows you to prevent them.

A security alert can be resolved with the number of specific actions that are provided by the security experts. There are some alerts that are used just for notifying you about the upcoming threats but no recommended actions are available.

How it works

Acronis Cyber Protection Operations Center monitors external threats and generates alerts about malware, vulnerability, natural disaster, and public health threats. You will be able to see all these alerts in the service console, in the **Threat feed** section. You can perform respective recommended actions depending on the type of alert.

The main workflow of the threat feed is illustrated in the diagram below.



To run the recommended actions on received alerts from Acronis Cyber Protection Operations Center, do the following:

1. In the service console, go to **Dashboard > Threat feed** to review if there are any existing security alerts.
2. Select an alert in the list and review the provided details.
3. Click **Start** to launch the wizard.
4. Enable the actions that you want to be performed and machines to which these actions must be applied. The following actions can be suggested:

- **Vulnerability assessment** – to scan machines for vulnerabilities
- **Patch management** – to install patches on the selected machines
- **Antimalware Protection** – to run full scan of the selected machines
- **Backup of protected or unprotected machines** – to back up protected/unprotected machines

5. Click **Start**.

6. On the **Activities** page, verify that the activity was successfully performed.

Name	Severity	Type	Date
Warning over powerful Smominru crypto mining botnet	MEDIUM	Malware	Dec 13, 2019
Acronis discovers new Autoit Cryptominer campaign injecting Windows process	HIGH	Malware	Dec 11, 2019
Manila vulnerable to major earthquake	LOW	Natural Disaster	Dec 11, 2019
Snatch ransomware reboots PCs into Safe Mode to bypass protection	HIGH	Malware	Dec 10, 2019
Caution! Ryuk ransomware decrypter damages larger files, even if you pay	MEDIUM	Malware	Dec 10, 2019
5.3 earthquake shakes New Zealand's North Island	LOW	Natural Disaster	Dec 10, 2019
Town hit by ransomware: System shut down to limit damage	MEDIUM	Malware	Dec 9, 2019
5.0M earthquake strikes Gunungkidul, Yogyakarta	LOW	Natural Disaster	Dec 9, 2019
Beware: Windows 10 update email is a ransomware trap	LOW	Malware	Dec 4, 2019
Dexphot malware uses fileless techniques to install cryptominer	LOW	Malware	Dec 4, 2019
New Chrome Password Stealer Sends Stolen Data to a MongoDB Database	LOW	Malware	Dec 2, 2019
New Malware Campaign Targets the Hospitality Industry	LOW	Malware	Dec 2, 2019
New DeathRansomware started encrypting files for real	HIGH	Malware	Nov 28, 2019
Docker platforms are targeted by hackers to deliver cryptomining malware	MEDIUM	Malware	Nov 28, 2019
Fake software update tries to download malware	MEDIUM	Malware	Nov 25, 2019
New malware DePrimon registers as Default Print Monitor	MEDIUM	Malware	Nov 22, 2019

Deleting all alerts

Automatic clean-up from the threat feed is made after the following time periods:

- Natural disaster – 1 week
- Vulnerability – 1 month
- Malware – 1 month
- Public health – 1 week

30.2 Data protection map

The Data protection map functionality allows you

- To get detailed information about stored data (classification, locations, protection status, and additional information) on your machines.
- To detect whether data are protected or not. The data are considered protected if they are protected with backup (a protection plan with the backup module enabled).
- To perform actions for data protection.

How it works

1. First, you create a protection plan with the Data protection map module (p. 373) enabled.
2. Then, after the plan was performed and your data were discovered and analyzed, you will get the visual representation of data protection on the Data protection map (p. 386) widget.

3. You can also go to **Devices > Data protection map** and find there information about unprotected files per device.
4. You can take actions to protect the detected unprotected files on devices.

Managing the detected unprotected files

To protect the important files that were detected as unprotected, do the following:

1. In the service console, go to **Devices > Data protection map**.
In the list of devices, you can find general information about the number of unprotected files, size of such files per device, and the last data discovery.
To protect files on a particular machine, click the Ellipsis icon and then **Protect all files**. You will be redirected to the list of plans where you can create a protection plan with the backup module enabled.
To delete the particular device with unprotected files from the list, click **Hide until next data discovery**.
2. To view a more detailed information about the unprotected files on a particular device, click on the name of the device.
You will see the number of unprotected files per extension and per location. Define the extensions in the search field, for which you want to get the information about unprotected files.
3. To protect all unprotected files, click **Protect all files**. You will be redirected to the list of plans where you can create a protection plan with the backup module enabled.

To get the information about the unprotected files in the form of report, click **Download detailed report in CSV**.

30.2.1 Data protection map settings

To learn how to create a protection plan with the Data protection map module, refer to "Creating a protection plan (p. 95)".

The following settings can be specified for the Data protection map module.

Schedule

You can define different settings to create the schedule according to which the task for data protection map will be performed.

Schedule the task run using the following events:

- **Schedule by time** – The task will run according to the specified time.
- **When user logs in to the system** – By default, login of any user will initiate a task run. You can this setting so that only a specific user account can trigger the task.
- **When user logs off the system** – By default, logoff of any user will make the task run. You can this setting so that only a specific user account can trigger the task.

Note *The task will not run at system shutdown. Shutting down and logging off are different actions.*

- **On the system startup** – The task will run when the operating system starts.
- **On the system shutdown** – The task will run when the operating system shuts down.

Default setting: **Schedule by time**.

Schedule type:

- **Monthly** – Select the months and the weeks or days of the month when the task will run.
- **Daily** – Select the days of the week when the task will run.
- **Hourly** – Select the days of the week, repetition number, and the time interval in which the task will run.

Default setting: **Daily**.

Start at – Select the exact time when the task will run.

Run within a date range – Set a range in which the configured schedule will be effective.

Start conditions – Define all the conditions that must be simultaneously met for the task to run. They are similar to the start conditions for the Backup module that are described in "Start conditions (p. 131)". Also, the following additional start conditions can be defined:

- **Distribute task start time within a time window** – This option allows you to set the time frame for the task in order to avoid network bottlenecks. You can specify the delay in hours or minutes. For example, if the default start time is 10:00 AM and the delay is 60 minutes, then the task will start between 10:00 AM and 11:00 AM.
- **If the machine is turned off, run missed tasks at the machine startup**
- **Prevent the sleep or hibernate mode during task running** – This option is effective only for machines running Windows.

If start conditions are not met, run the task anyway after – Specify the period after which the task will run, regardless of the other start conditions.

Extensions and exception rules

On the **Extensions** tab, you can define the list of file extensions that will be considered as important during data discovery and checked whether they are protected. Use the following format for defining extensions:

.html, .7z, .docx, .zip, .pptx, .xml

On the **Exception rules** tab, you can define which files and folders not to check on protection status during data discovery.

- **Hidden files and folders** – if selected, hidden files and folders will be skipped during data examination.
- **System files and folders** – if selected, system files and folders will be skipped during data examination.

31 The Plans tab

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

You can manage protection plans and other plans by using the **Plans** tab.

Each section of the **Plans** tab contains all plans of a specific type. The following sections are available:

- **Protection** (p. 375)
- **Backup scanning** (p. 377)
- **Cloud applications backup** (p. 378)
- **VM replication** (p. 271)

31.1 Protection plan

To create the protection plan

1. In the service console, go to **Plans > Protection**.
2. Click **Create plan**.
3. Select the machines that you want to protect.

4. Click **Protect**. You will see the protection plan with the default settings.

Create protection plan ✕

+ Add devices

New protection plan (1)

Cancel>Create

Backup Entire machine to Cloud storage, Monday to Friday at 11:45 AM	☒	>
Anti-malware Protection Self-protection on, Real-time protection on, at 01:55 PM, Sunday through Saturday	☒	>
URL filtering Always ask user	☒	>
Windows Defender Antivirus Full scan, Real-time protection on, at 12:00 PM, only on Friday	☐	>
Microsoft Security Essentials Full scan, at 12:00 PM, only on Friday	☐	>
Vulnerability assessment Microsoft products, Windows third-party products, Linux packages, at 02:00 PM, ...	☒	>
Patch management Microsoft and other third-party products, at 02:35 PM, only on Monday	☒	>
Data protection map 66 extensions, at 03:50 PM, Monday through Friday	☒	>

5. [Optional] To modify the protection plan name, click on the pencil icon next to the name.

6. [Optional] To enable or disable the plan module, click the switch next to the module name.
7. [Optional] To configure the module parameters, click the corresponding section of the protection plan.
8. Click **Add devices** to select the machines to which you want to apply the plan.
9. When ready, click **Create**.

As a result, the selected devices will be protected with the protection plan.

You can perform the following operations with protection plans:

- Create, view, edit, clone, disable, enable, and delete a protection plan
- View activities related to each protection plan
- View alerts related to each protection plan
- Export a plan to a file
- Import a previously exported plan

31.2 Backup scanning plan

If you need to scan backups on malware, you can create a backup scanning plan.

Note the following:

- The backups that have CDP recovery points (p. 118) can be selected for scanning but only regular recovery points (excluding CDP recovery points) will be scanned.
- When the CDP backup was selected for safe recovery of an entire machine, the machine will be safely recovered without the data in the CDP recovery point. To restore the CDP data, start the Files/folders recovery activity.

To create a backup scanning plan

1. In the service console, go to **Plans > Backup scanning**.
2. Click **Create plan**.
3. Specify the name of the plan and the following parameters:
 - **Scan type:**
 - **Cloud** – this option cannot be redefined. The backups will be scanned in the cloud data center by the cloud agent. The system automatically selects the cloud agent that will perform scanning.
 - **Backups to scan:**
 - **Locations** – select locations with backups that you want to scan.
 - **Backups** – select backups that you want to scan.
 - **Scan for:**
 - **Malware** – this option cannot be redefined. It checks backups on malware presence.
 - **Encryption** – provide a password to scan encrypted backups. If a vault or multiple backups are selected, then you can specify a single password for all backups. If the password does not suit to a backup, the system will create an alert.
 - **Schedule** – this option cannot be redefined. The scan activity will be started in the cloud storage automatically.
4. When ready, click **Create**.

As a result, the backup scanning plan is created. The specified locations or backups will be scanned by the cloud agent automatically.

31.3 Backup plans for cloud applications

The **Plans > Cloud applications backup** section shows cloud-to-cloud backup plans. These plans back up applications running in the cloud by means of agents that run in the cloud and use the cloud storage as a backup location.

In this section, you can perform the following operations:

- Create, view, run, stop, edit, and delete a backup plan
- View activities related to each backup plan
- View alerts related to each backup plan

For more information about cloud applications backup, refer to:

- Protecting Office 365 data (p. 230)
- Protecting G Suite data (p. 252)

Running cloud-to-cloud backups manually

To prevent disrupting the Cyber Protection service, the number of manual cloud-to-cloud backup runs is limited to 10 runs per Office 365 or GSuite organization during an hour. After this number has been reached, the number of runs allowed is reset to one per hour, and then an additional run becomes available each hour thereafter (e.g. hour 1, 10 runs; hour 2, 1 run; hour 3, 2 runs) until a total of 10 runs per hour is reached.

Backup plans applied to groups of devices (mailboxes, drives, sites) or containing more than 10 devices cannot be run manually.


32 Monitoring

The **Overview** dashboard provides a number of customizable widgets that give an overview of operations related to the Cyber Protection service. Widgets for other services will be available in future releases.

The widgets are updated every five minutes. The widgets have clickable elements that enable you to investigate and troubleshoot issues. You can download the current state of the dashboard or send it via email in the .pdf or/and .xlsx format.

You can choose from a variety of widgets, presented as tables, pie charts, bar charts, lists, and tree maps. You can add multiple widgets of the same type with different filters.

The buttons **Download** and **Send** in **Dashboard > Overview** are not available in the Standard editions of the Cyber Protection service.



To rearrange the widgets on the dashboard

Drag and drop the widgets by clicking on their names.

To edit a widget

Click the pencil icon next to the widget name. Editing a widget enables you to rename it, change the time range, set filters, and group rows.

To add a widget

Click **Add widget**, and then do one of the following:

- Click the widget that you want to add. The widget will be added with the default settings.
- To edit the widget before adding it, click the pencil icon when the widget is selected. After editing the widget, click **Done**.

To remove a widget

Click the X sign next to the widget name.

32.1 Cyber Protection

This widget shows the overall information about the size of backups, blocked malware, blocked URLs, found vulnerabilities, and installed patches.

Cyber Protection				
				
Backed up today	Malware blocked	Malicious URLs blocked	Existing vulnerabilities	Patches ready to install
1.60 GB	0	0	347	114
overall compressed size	overall blocked	overall blocked	overall found	overall installed
2.43 GB	14	4	819	5

The upper row shows the current statistics:

- **Backed up today** – the sum of recovery point sizes for the last 24 hours
- **Malware blocked** – the number of currently active alerts about malware blocked
- **URLs blocked** – the number of currently active alerts about URLs blocked
- **Existing vulnerabilities** – the number of currently existing vulnerabilities
- **Patches ready to install** – the number of currently available patches to be installed

The lower row shows the overall statistics:

- The compressed size of all backups
- The accumulated number of blocked malware across all machines
- The accumulated number of blocked URLs across all machines
- The accumulated number of discovered vulnerabilities across all machines
- The accumulated number of installed updates/patches across all machines

32.2 Protection status


Protection status

This widget shows the current protection status for all machines.

A machine can be in one of the following statuses:

- **Protected** – machines with applied protection plan.
- **Unprotected** – machines without applied protection plan. These include both discovered machines and managed machines with no protection plan applied.
- **Managed** – machines with installed protection agent.
- **Discovered** – machines without installed protection agent.

If you click on the machine status, you will be redirected to the list of machines with this status for more details.



Discovered machines

This widget shows the list of discovered machines during the specified time range.

Discovered machines					
Device name ↑	IP address	OS	Organizational unit	Discovery type	⚙
Windows Server 2012 R2					
win-6g34mv70qa3	10.248.90.221	Windows Server 2012 R2	-	Local Network	
Windows 10 Enterprise 2016 LTSC					
device1	10.248.90.238	Windows 10 Enterprise 2016 LTSC	OU1	Active Directory, Local Network	
device2	-	Windows 10 Enterprise 2016 LTSC	OU1	Active Directory	
device3	10.248.91.243	Windows 10 Enterprise 2016 LTSC	OU1	Active Directory, Manual, Loc...	
device4	10.248.91.125	Windows 10 Enterprise 2016 LTSC	-	Active Directory, Local Network	
device5	-	Windows 10 Enterprise 2016 LTSC	-	Active Directory, Manual	
-					
-	10.250.41.189	-	-	Manual	
-	10.248.44.199	-	-	Manual	

32.3 #CyberFit Score by machine

This widget shows for each machine the total #CyberFit Score, its compound scores, and findings for each of the assessed metrics:

- Antimalware
- Backup
- Firewall
- VPN
- Encryption
- NTLM traffic

To improve the score of each of the metrics, you can view the recommendations that are available in the report.

For more details about the #CyberFit Score, refer to "#CyberFit Score for machines (p. 103)".

#CyberFit Score by machine ?			
Metric	#CyberFit Score	Findings	
▼  DESKTOP-2N2TRE8	 625 / 850		
Anti-malware	 275 / 275	You have anti-malware protection enabled	
Backup	 175 / 175	You have a backup solution protecting your data	
Firewall	 175 / 175	You have a firewall enabled for public and private networks	
VPN	 0 / 75	No VPN solution was found, your connection to public and shared networks is n...	
Encryption	 0 / 125	No disk encryption was found, your device is at risk from physical tampering	
NTLM traffic	 0 / 25	Outgoing NTLM traffic to remote servers is not denied, your credentials may be ...	

32.4 Disk health forecast

The disk health control feature allows you to monitor the current disk health status and get a forecast of disk health. This information allows you to prevent any problems with data loss related to disk crashes. Both HDD and SSD types of disk are supported.

Limitations:

1. Disk health forecast is supported only for Windows machines.
2. Only the disks of physical machines can be monitored. The disks of virtual machines cannot be monitored and shown in the widget.

Disk health can be in one of the following statuses:

- **OK** – disk health is 70-100%
- **Warning** – disk health is 30-70%
- **Critical** – disk health is 0-30%
- **Calculating disk data** – the current disk status and forecast are being calculated

How it works


The Disk Health Prediction Service uses the artificial intelligence based prediction model.

1. The agent collects the SMART parameters of disks and passes this data to Disk Health Prediction Service:
 - SMART 5 – reallocated sectors count
 - SMART 9 – power-on hours
 - SMART 187 – reported uncorrectable errors
 - SMART 188 – command timeout
 - SMART 197 – current pending sector count
 - SMART 198 – offline uncorrectable sector count
 - SMART 200 – write error rate
2. Disk Health Prediction Service processes the received SMART parameters, makes forecasts, and provides the following disk health characteristics:
 - Disk health current state: OK, Warning, Critical.

- Disk health forecast: negative, stable, positive.
- Disk health forecast probability in percentage.

The prediction period is always one month.

3. The Monitoring Service gets the disk health characteristics and use this data in disk health widgets shown to a user in the console.



Disk health widgets


The results of the disk health monitoring can be found on the dashboard in the disk health related widgets:

- **Disk health overview** – a treemap widget that has two levels of details that can be switched by drilling down:

- Machine level – shows summarized information about disk status per the selected customer machines. The widget represents the most critical disk status data, other statuses are shown in the tooltip when you hover over the particular block. The machine block size depends on the total size of all disks of this machine. The machine block color depends on the most critical disk status found.

Disk health overview

Resources




- Disk level – shows the current disk status of all disks for the selected machine. Each disk block shows a forecast of disk status change:
 - Will be degraded (disk health forecast probability in %)
 - Will stay stable (disk health forecast probability in %)

- Will be improved (disk health forecast probability in %)



- **Disk health status** – a pie chart widget showing the number of disks for each status.



Disk health status alerts

Disk health check runs every 30 minutes while the corresponding alert is generated once a day. When the disk health has changed from Warning to Critical, you will also get the alert even if you already got another alert during a day.

Alert name	Severity	Disk health status	Description
Disk failure is possible	Warning	[30;70)	The [disk_name] disk on [machine_name] machine is likely to fail in the future. Please run a full image backup of this disk as soon as possible, replace it and then recover the image to the new disk.
Disk failure is imminent	Critical	(0;30)	The [disk_name] disk on [machine_name] machine is in a critical state and will most likely fail very soon. An image backup of this disk is not recommended at this point as the added stress can cause the disk to fail. Please back up all the most important files on this disk right now and replace it.

32.5 Data protection map

The data protection map feature allows you to discover all data that are important for you and get detailed information about number, size, location, protection status of all important files in a treemap scalable view.


Each block size depends on the total number/size of all important files that belong to a customer/machine.

Files can have one of the following protection statuses:

- **Critical** – there are 51-100% of unprotected files with the extensions specified by you that are not being backed up and will not be backed up with the existing backup settings for the selected machine/location.
- **Low** – there are 21-50% of unprotected files with the extensions specified by you that are not being backed up and will not be backed up with the existing backup settings for the selected machine/location.
- **Medium** – there are 1-20% of unprotected files with the extensions specified by you that are not being backed up and will not be backed up with the existing backup settings for the selected machine/location.
- **High** – all files with the extensions specified by you are protected (backed up) for the selected machine/location.

The results of the data protection examination can be found on the dashboard in the Data Protection Map widget, a treemap widget that shows details on a machine level:

- Machine level – shows information about the protection status of important files per machines of the selected customer.



To protect files that are not protected, hover over the block and click **Protect all files**. In the dialog window, you can find information about the number of unprotected files and their location. To protect them, click **Protect all files**.

You can also download a detailed report in CSV format.

32.6 Vulnerability assessment widgets

Vulnerable machines

This widget shows the vulnerable machines by the vulnerability severity.

The found vulnerability can have one of the following severity levels according to the Common Vulnerability Scoring System (CVSS) v3.0:

- Secured: no vulnerabilities are found
- Critical: 9.0 - 10.0 CVSS
- High: 7.0 - 8.9 CVSS
- Medium: 4.0 - 6.9 CVSS
- Low: 0.1 - 3.9 CVSS

- None: 0.0 CVSS



Existing vulnerabilities

This widget shows currently existing vulnerabilities on machines. In the **Existing vulnerabilities** widget, there are two columns showing timestamps:

- **First detected** – date and time when a vulnerability was detected initially on the machine.
- **Last detected** – date and time when a vulnerability was detected the last time on the machine.

Machine name	Vendor	Product	Vulnerability name/ID	Severity ↓	Last detected	First detected	⚙
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-7096	Critical	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0856	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0688	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0739	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0752	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0753	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0806	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0810	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0812	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
DESKTOP-NU01945	Microsoft	Windows 10 LTSB	CVE-2019-0829	High	06/12/2020 5:16 PM	06/12/2020 5:15 PM	
More							

32.7 Patch installation widgets


There are four widgets related to the patch management functionality.

Patch installation status

This widget shows the number of machines grouped by the patch installation status.

- **Installed** – all available patches are installed on a machine
- **Reboot required** – after patch installation reboot is required for a machine

- **Failed** – patch installation failed on a machine



Patch installation summary

This widget shows the summary of patches on machines by the patch installation status.

Patch installation summary							
Installation status	Total number of machines	Total number of updates	Microsoft updates	Application updates	Critical severity	High severity	Medium severity
Installed	1	2	1	1	2	0	0

Patch installation history

This widget shows the detailed information about patches on machines.


Patch installation history						
Machine name	Update name	Version	Severity	Approval status	Installation status	Installation date ↓
NIKITATIKHOC4E5	FastStone Soft FastStone I...	5.9	Medium	New	Installed	02/05/2020
NIKITATIKHOB524	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020
NIKITATIKHOB524	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Oracle Java Runtime Envir...	8.0.2410.7	High	New	Failed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Installed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020
NIKITATIKHOC4E5	Mozilla Firefox	72.0.1	Critical	Approved	Failed	02/04/2020

Missing updates by categories

This widget shows the number of missing updates per category. The following categories are shown:

- Security updates
- Critical updates

- Other



32.8 Backup scanning details

This widget shows the detailed information about the detected threats in backups.

Device name	Plan name	Backup Date and time	Contents type	Location	Threat name	Affected files	Date and time	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Gen:Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 11:40 AM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 11:40 AM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Gen:Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 11:45 AM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 11:45 AM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Gen:Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 11:50 AM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 11:50 AM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Gen:Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 1:10 PM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 1:10 PM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Gen:Heur.PonyStealer.lm0@c05cs0dG	F:\882a04265361d588801b35...	01/21/2020 1:33 PM	
NIKITATIKHOC4E5	New protection plan (1)	01/20/2020 11:00 AM	Full	\\...	Trojan.GenericKD.3947747	F:\2f2b2e30abe71f9a93d6ad7...	01/21/2020 1:33 PM	

More

32.9 Recently affected

This widget shows the detailed information about recently infected machines. You can find information about what threat was detected and how many files were infected.

Recently affected					
Machine name	Protection plan	Threat	Affected files	Detection time	⚙
Ubuntu_14.04_x64-1	Total protection	Adware.DealPlyIgen2	15	27.12.2	Folder Customer ✓ Machine name ✓ Protection plan Detected by ✓ Threat File name File path ✓ Affected files ✓ Detection time
Ubuntu_14.04_x64-1	Protection plan	Bloodhound.MalMacroIlg1	274	27.12.2	
dc_w2k12_r2	Protection plan	Backdoor:Win32/Caphaw...	13	27.12.2	
Win2012_r2-Hyper-V	Protection plan	W97M.DownloaderIlg32	5	27.12.2	
HyperV_for12A	Total protection	Miner.XMRigIlg1	68	27.12.2	
vm-sql_2012	Total protection	Backdoor:Win32/Caphaw...	61	27.12.2	
vm-sql_2012	Protection plan	Adware.DealPlyIgen2	9	27.12.2	
MF_2012_R2	Total protection	MSH.DownloaderIlg8	73	27.12.2	
MF_2012_R2	Total protection	Bloodhound.MalMacroIlg1	182	27.12.2	
MF_2012_R2	Protection plan	Bloodhound.MalMacroIlg1	18	27.12.2	
ESXirestore	Protection plan	MSH.DownloaderIlg8	682	27.12.2017 11:23 AM	
MF_2012_R2	Protection plan	Miner.XMRigIlg1	13	27.12.2017 11:23 AM	
Ubuntu_14.04_x64-1	Total protection	Adware.DealPlyIgen2	3	27.12.2017 11:23 AM	
Win2012_r2-Hyper-V	Total protection	W97M.DownloaderIlg32	27	27.12.2017 11:23 AM	
More Show all 556					

32.10 Cloud applications

This widget shows detailed information about cloud-to-cloud resources:

- Office 365 users (mailbox, OneDrive)
- Office 365 groups (mailbox, group site)
- Office 365 public folders
- Office 365 site collections
- Office 365 Teams
- G Suite users (Gmail, GDrive)

- G Suite shared drives

Cloud applications  				
Device name	Protection status 	Last successful backup	Next backup	Number of backups 
 HR - Onboarding	 OK	06/17/2020 10:48 AM	06/18/2020 7:34 AM	1
 Sales and Marketing	 OK	06/17/2020 10:49 AM	06/18/2020 4:48 AM	1
 HR Leadership Team	 OK	06/17/2020 10:48 AM	06/18/2020 6:51 AM	1
 Retail	 OK	06/17/2020 10:47 AM	06/18/2020 2:53 AM	1
 Contoso	 OK	06/17/2020 10:47 AM	06/17/2020 3:23 PM	1
 U.S. Sales	 OK	06/17/2020 10:48 AM	06/18/2020 3:30 AM	1
 IT	 OK	06/17/2020 10:48 AM	06/17/2020 10:35 PM	1
 Mark 8 Project Team	 Warning	06/17/2020 10:49 AM	06/18/2020 3:06 AM	1
 Finance	 OK	06/17/2020 10:47 AM	06/17/2020 4:38 PM	1
 Sales	 Warning	06/17/2020 10:47 AM	06/17/2020 2:06 PM	1
More				

Additional information about cloud-to-cloud resources is also available in the following widgets:


- Activities
- Activity list
- 5 latest alerts
- Alerts history
- Active alerts summary
- Historical alerts summary
- Active alert details
- Locations summary

33 Reports

Note This functionality is available only in the Advanced edition of the Cyber Protection service.

A report about operations can include any set of the dashboard widgets (p. 379). All of the widgets show the summary information for the entire company. All of the widgets show the parameters for the same time range. You can change this range in the report settings.

You can use default reports or create a custom report.



The set of default reports depends on the Cyber Protection service edition that you have. The default reports are listed below:

Report name	Description
#CyberFit Score by machine	Shows the #CyberFit Score, based on the evaluation of security metrics and configurations for each machine, and recommendations for improvements.
Alerts	Shows alerts that occurred during a specified time period.
Backup scanning details	Shows the detailed information about detected threats in the backups.
Daily activities	Shows the summary information about activities performed during a specified time period.
Data protection map	Shows the detailed information about the number, size, location, protection status of all important files on machines.
Detected threats	Shows the details of the affected machines by number of blocked threats and the healthy and vulnerable machines.
Discovered machines	Shows all found machines in the organization network.
Disk health prediction	Shows predictions when your HDD/SSD will break down and current disk status.
Existing vulnerabilities	Shows the existing vulnerabilities for OS and applications in your organization. The report also displays the details of the affected machines in your network for every product that is listed.

Patch management summary	Shows the number of missing patches, installed patches, and applicable patches. You can drill down the reports to get the missing/installed patch information and details of all the systems.
Summary	Shows the summary information about the protected devices for a specified time period.
Weekly activities	Shows the summary information about activities performed during a specified time period.

To view a report, click its name.

To access operations with a report, click the ellipsis icon on the report line. The same operations are available from within the report.

Adding a report

1. Click **Add report**.
2. Do one of the following:
 - To add a predefined report, click its name.
 - To add a custom report, click **Custom**, click the report name (the names assigned by default look like **Custom(1)**), and then add widgets to the report.
3. [Optional] Drag and drop the widgets to rearrange them.
4. [Optional] Edit the report as described below.

Editing a report

To edit a report, click its name, and then click **Settings**. When editing a report, you can:

- Rename the report
- Change the time range for all widgets included in the report

- Schedule sending the report via email in the .pdf or/and .xlsx format

General

Name

Backup scanning details

☐ Set one tenant for all widgets

Range

7 days

▼

Scheduled

Recipients

user1@example.com; user2@example.com

File format

Excel and PDF

▼

Language

English

▼

Days of week

Monthly

SUN

MON

TUE

WED

THU

FRI

SAT

Send at

12:00 AM

▼

Scheduling a report

1. Click the report name, and then click **Settings**.
2. Enable the **Scheduled** switch.
3. Specify the recipients' email addresses.
4. Select the report format: .pdf, .xlsx, or both.
5. Select the days and the time when the report will be sent.
6. Click **Save** in the upper right corner.

Note The maximum number of exported items is: in a .pdf file—1000; in an .xlsx file—10 000.

Exporting and importing the report structure

You can export and import the report structure (the set of widgets and the report settings) to a .json file.

To export the report structure, click the report name, click the ellipsis icon in the top-right corner, and then click **Export**.

To import the report structure, click **Add report**, and then click **Import**.

Downloading a report

You can download a report, click **Download** and select the formats needed:

- Excel and PDF
- Excel
- PDF

Dumping the report data

You can send a dump of the report data in a .csv file via email. The dump includes all of the report data (without filtering) for a custom time range. The timestamps in CSV reports are in the UTC format whereas in Excel and PDF reports the timestamps are in the current system time zone.

The software generates the data dump on the fly. If you specify a long period of time, this action may take a long time.

To dump the report data

1. Click the report name.
2. Click the ellipsis icon in the top-right corner, and then click **Dump data**.
3. Specify the recipients' email addresses.
4. In **Time range**, specify the time range.
5. Click **Send**.

Note The maximum number of items exported in a .csv file is 150 000.

34 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.

35 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**).
If a weekly backup is the first backup created after a month starts, this backup is considered monthly. In this case, a weekly backup will be created on the selected day of the next week.
- A daily backup is the first backup created after a day starts, unless this backup falls within the definition of a monthly or weekly backup.
- An hourly backup is the first backup created after an hour starts, unless this backup falls within the definition of a monthly, weekly, or daily backup.

C

Cloud server

[Disaster Recovery] General reference to a recovery or a primary server.

Cloud site (or DR site)

[Disaster Recovery] Remote site hosted in the cloud and used for running recovery infrastructure, in case of a disaster.

D

Differential backup

A differential backup stores changes to the data against the latest full backup (p. 398). You need access to the corresponding full backup to recover the data from a differential backup.

F

Failback

Switching a workload from a spare server (such as a virtual machine replica or a recovery server running in the cloud) back to the production server.

Failover

Switching a workload from a production server to a spare server (such as a virtual machine replica or a recovery server running in the cloud).

Finalization

The operation that makes a temporary virtual machine that is running from a backup into a permanent virtual machine. Physically, this means recovering all of the virtual machine disks, along with the changes that occurred while the machine was running, to the datastore that stores these changes.

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.

L

Local site

[Disaster Recovery] The local infrastructure deployed on your company's premises.

M

Module

Module is a part of protection plan providing a particular data protection functionality, for example, the backup module, the Antivirus and Antimalware protection module, and so on.

P

Physical machine

A machine that is backed up by an agent installed in the operating system.

Point-to-site (P2S) connection

[Disaster Recovery] A secure VPN connection from outside to the cloud and local sites by using your endpoint devices (such as a computer or laptop).

Primary server

[Disaster Recovery] A virtual machine that does not have a linked machine on the local site (such as a recovery server). Primary servers are used for protecting an application or running various auxiliary services (such as a web server).

Production network

[Disaster Recovery] The internal network extended by means of a VPN tunneling and covering both local and cloud sites. Local servers and cloud servers can communicate with each other in the production network.

Protection agent

Protection agent is the agent to be installed on machines for data protection.

Protection plan

Protection plan is a plan that combines the data protection modules including

- Backup
- Antivirus and Antimalware protection
- URL filtering
- Windows Defender Antivirus
- Microsoft Security Essentials
- Vulnerability assessment
- Patch management
- Data protection map

Public IP address

[Disaster Recovery] An IP address that is needed to make cloud servers available from the Internet.

R

Recovery point objective (RPO)

[Disaster Recovery] Amount of data lost from outage, measured as the amount of time from a planned outage or disaster event.

RPO threshold defines the maximum time interval allowed between the last suitable recovery point for a failover and the current time.

Recovery server

[Disaster Recovery] A VM replica of the original machine, based on the protected server backups stored in the cloud. Recovery servers are used for switching workloads from the original servers, in case of a disaster.

Runbook

[Disaster Recovery] Planned scenario consisting of configurable steps that automate disaster recovery actions.

S

Single-file backup format

A backup format, in which the initial full and subsequent incremental backups are saved to a single .tibx file. 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.

The single-file backup format is not available when backing up to locations that do not support random-access reads and writes.

Site-to-site (S2S) connection

[Disaster Recovery] Connection extending the local network to the cloud, via a secure VPN tunnel.

T

Test IP address

[Disaster Recovery] An IP address that is needed in case of a test failover, to prevent duplication of the production IP address.

Test network

[Disaster Recovery] Isolated virtual network that is used to test the failover process.

V

Virtual machine

A virtual machine that is backed up at a hypervisor level by an external agent such as Agent for VMware or Agent for Hyper-V.

A virtual machine with an agent inside is treated as physical from the backup standpoint.

VPN appliance

[Disaster Recovery] A special virtual machine that enables connection between the local network and the cloud site via a secure VPN tunnel. The VPN appliance is deployed on the local site.

VPN gateway (formerly, VPN server or connectivity gateway)

[Disaster Recovery] A special virtual machine providing a connection between the local site and the cloud site networks via a secure VPN tunnel. The VPN gateway is deployed on the cloud site.