

Rapid Recovery 6.8

User Guide



© 2023 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

 **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Contents

Introduction to Rapid Recovery	1
The Core Console	3
Accessing the Rapid Recovery Core Console	3
Understanding the Quick Start Guide	4
Navigating the Rapid Recovery Core Console	6
Understanding the left navigation area	7
Viewing the Rapid Recovery Core Console Home page	8
Understanding the Home page (summary tables view)	8
Understanding Core dashboard reports	11
Viewing the Protected Machines menu	11
Viewing summary information for a protected machine	12
Viewing the Summary pane	13
Viewing Volumes on a protected machine	13
Viewing replication information	13
Viewing the Exchange Server Information pane	13
Viewing the SQL Server Information pane	13
Viewing summary information for a hypervisor or cluster host	14
Viewing recovery points for a machine	14
Viewing events for a protected machine	14
Viewing reports for a protected machine	16
Viewing replicated machines from the navigation menu	16
Viewing the Recovery Points Only menu	17
Viewing the Custom Groups menu	17
Using the Error dialog box	18
Repositories	19
Understanding repositories	19
DVM repositories	20
Azure repositories	21
Deduplication in Rapid Recovery	22
When deduplication occurs	22
Managing a DVM repository	23
Creating a DVM repository	23
Changing DVM repository settings	28
Managing an Azure repository	29
Creating an Azure repository	30
Changing Azure repository settings	32
About repository optimization	33

Optimizing a repository	33
Interrupting or resuming repository optimization	34
Connecting to an existing repository	34
Viewing or modifying repository details	36
Checking a repository	38
Deleting a repository	38
Core settings	39
Core settings key functions	39
Backing up and restoring Core settings	40
Restarting or shutting down the Core service	42
Rapid Recovery Core settings	43
Configuring Core general settings	46
Configuring update settings	48
Understanding nightly jobs	50
Configuring nightly jobs for the Core	53
Modifying transfer queue settings	53
Adjusting client timeout settings	54
Understanding deduplication cache and storage locations	55
Optimizing deduplication	56
Configuring DVM deduplication cache settings	57
Configuring Replay engine settings	58
Configuring deployment settings	60
Core-level tools	61
Viewing system information for the Core	61
Accessing Core logs	63
Configuring database connection settings	64
Modifying local database connection settings	65
Managing SMTP server settings	66
Configuring cloud account connection settings	66
Managing report settings	67
Managing Core SQL attachability settings	68
Understanding Core jobs	70
Core job settings	73
Adding Core jobs to settings	74
Editing Core job settings	74
Understanding SNMP settings	75
Configuring SNMP settings	76
Downloading the SNMP MIB file	77
Configuring vSphere settings	78
Managing VMware proxy settings	79
Configuring vFoglight settings	80
Configuring SAML settings	81

Protecting machines	84
About protecting machines with Rapid Recovery	84
Factors when choosing agent-based or agentless protection	85
General recommendations	85
Release 6.7 license consumption concepts	86
Licensing benefits of using agentless protection	86
About protecting Linux machines with Rapid Recovery	87
About protecting Oracle database servers	88
Entering or editing credentials for Oracle databases	89
Enabling archive log mode and adding VSS writer for protected Oracle databases	91
About truncating Oracle logs	91
Manually truncating Oracle database logs	92
About managing Exchange and SQL servers in Rapid Recovery Core	93
About protecting server clusters	94
Understanding Rapid Snap for Virtual	94
Protecting VMware vCenter/ESXi VMs	95
Protecting VMs on Hyper-V servers and clusters	96
Application support	98
Benefits of installing hypervisor tools for agentless protection	98
Understanding crash-consistent and application-consistent backups	99
Support for Hyper-V guest cluster	99
Understanding the Rapid Recovery Agent software installer	100
Downloading the Rapid Recovery Agent Installer	100
Deploying Agent to multiple machines simultaneously from the Core Console	101
Using the Deploy Agent Software Wizard to deploy to one or more machines	101
Deploying to machines on an Active Directory domain	102
Deploying to machines on a VMware vCenter/ESXi virtual host	103
Deploying an upgrade of the Rapid Recovery Agent to protected machines	104
Deploying to machines manually	105
Verifying the deployment to multiple machines	106
Modifying deploy settings	106
Understanding protection schedules	107
Protecting a machine	108
Protecting a cluster	112
Protecting nodes in a cluster	118
Creating custom protection schedules in Simple Mode	119
Creating multiple protection schedule periods in Advanced Mode	121
Pausing and resuming protection	122
About protecting multiple machines	124
Protecting multiple machines on an Active Directory domain	125
Protecting multiple machines on a VMware vCenter/ESXi virtual host	129
Protecting vCenter/ESXi virtual machines using agentless protection	133
Protecting multiple machines on a Hyper-V virtual host	137
Protecting Hyper-V virtual machines using host-based protection	141

Protecting multiple machines manually	146
Monitoring the protection of multiple machines	149
Enabling application support	149
Settings and functions for protected Exchange servers	150
Setting credentials for an Exchange server machine	151
Forcing log truncation for an Exchange machine	152
About Exchange database mountability checks	152
Forcing a mountability check of an Exchange database	152
Forcing a checksum check of Exchange database files	153
Settings and functions for protected SQL servers	153
Setting credentials for a SQL Server machine	154
Forcing log truncation for a SQL machine	155
About SQL attachability	155
Forcing a SQL Server attachability check	156
Managing protected machines	157
About managing protected machines	157
Viewing protected machines	157
Viewing cluster summary information	158
Configuring machine settings	158
Viewing and modifying protected machine settings	159
Changing the settings for a Hyper-V host or node	165
Changing the settings for a Hyper-V protected virtual machine	166
Changing the vSphere settings for a VMware protected virtual machine	167
Understanding Active Block Mapping	168
Changing ABM settings	168
About modifying transfer settings	169
Throttling transfer speed	170
Customizing nightly jobs for a protected machine	171
Understanding system information for a protected machine	172
Viewing system information for a protected machine	172
Managing machines	173
Removing a machine	173
Removing a cluster from protection	173
Removing cluster nodes from protection	174
Removing all nodes in a cluster from protection	174
Viewing license information on a machine	175
Downloading and viewing the log file for a protected machine	175
Converting a protected cluster node to a protected machine	175
Understanding custom groups	176
Creating custom groups	177
Modifying custom group names	177
Removing custom groups	178

Performing group actions	178
Viewing all machines in a custom group on one page	179
Snapshots and recovery points	180
Managing snapshots and recovery points	180
Viewing the recovery points page of a protected machine	180
Understanding recovery point status indicators	183
Recovery status point colors for Exchange databases	183
Recovery status point colors for SQL databases	183
Mounting a recovery point	184
Dismounting recovery points	185
Working with Linux recovery points	186
Mounting a recovery point volume on a Linux machine	187
Unmounting a recovery point on a Linux machine	188
Forcing a snapshot	189
Removing recovery points	189
Deleting an orphaned recovery point chain	190
Migrating recovery points manually to a different repository	191
Managing privacy	193
General Data Protection Regulation compliance	193
How Rapid Recovery uses personal information	194
Non-phone-home license restrictions	195
Obtaining and using non-phone-home licenses	196
Encryption	198
Understanding encryption keys	198
Encrypting data in transport over a network	199
Applying or removing encryption keys	199
Associating an encryption key with a protected machine	200
Applying an encryption key from the Protected Machines page	201
Disassociating an encryption key from a protected machine	202
Managing encryption keys	203
Adding an encryption key	204
Importing an encryption key	206
Unlocking an encryption key	206
Locking an encryption key	207
Editing an encryption key	208
Changing an encryption key passphrase	208
Exporting an encryption key	209
Removing an encryption key	209
Changing encryption key types	210

Authentication	212
Understanding SAML single sign-on	212
Prerequisite	212
Credentials Vault	212
Understanding the Credentials Vault	213
Adding accounts to the Credentials Vault	213
Viewing or changing accounts saved in the vault	214
Using credentials from the vault	215
Replication	216
Replication with Rapid Recovery	216
Recovery point chains and orphans	220
When replication begins	221
Determining your seeding needs and strategy	221
When seeding data is required	221
Approaches to seeding data	222
Related links	223
Performance considerations for replicated data transfer	223
About replication and encrypted recovery points	225
About retention policies for replication	225
Viewing incoming and outgoing replication	225
Configuring replication	227
Replicating to a self-managed target Core	228
Replicating to a third-party target Core	233
Submitting a replication request to a third-party service provider	233
Reviewing a replication request from a customer	237
Approving a replication request	237
Denying a replication request	237
Ignoring a replication request from a customer	238
Adding a machine to existing replication	238
Consuming the seed drive on a target Core	243
Abandoning a seed drive	244
Managing replication settings	245
Scheduling replication	245
Using the Copy function to create a seed drive	246
Monitoring replication	249
Pausing and resuming replication	251
Forcing replication	252
Managing settings for outgoing replication	253
Changing target Core settings	253
Setting replication priority for a protected machine	254
Removing outgoing replication from the source Core	255

Removing incoming replication from the target Core	255
Recovering replicated data	256
Events	258
Viewing events using tasks, alerts, and journal pages	258
Viewing tasks	259
Viewing running tasks from any Core Console page	260
Suspending or resuming scheduled tasks	261
Viewing alerts	262
Viewing a journal of all logged events	263
Navigating between tasks, alerts, and the events journal	265
Understanding event notifications in Rapid Recovery	266
Configuring notification groups	267
Understanding email notifications	270
Configuring an email server	270
Configuring an email notification template	272
Configuring event settings	274
About repetition reduction	275
Configuring repetition reduction	275
Configuring event retention	276
Reporting	277
About Rapid Recovery reports	277
Generating reports from the Core Console	279
Generating a Core report on demand	279
Generating a protected machine report on demand	282
Managing scheduled reports from the Core Console	284
Scheduling a report	284
Modifying a report schedule	287
Pausing, resuming, or deleting a scheduled report	288
Using the Reports menu	288
Using the Reports toolbar	289
Understanding the Job report	291
Understanding the Job Summary report	291
Core information	291
Protected machines summary	292
Understanding the Failure report	292
Understanding the Summary report	292
Core information	293
Repositories summary	293
Protected machines summary	293
Understanding the Repository report	294
Understanding the Classic Summary report	294

VM export	295
Exporting to virtual machines using Rapid Recovery	295
Exporting data to an ESXi virtual machine	297
Performing a one-time ESXi export	297
Setting up continual export to ESXi	301
Exporting data to a VMware Workstation virtual machine	303
Performing a one-time VMware Workstation export	304
Setting up continual export to VMware Workstation	305
Exporting data to a Hyper-V virtual machine	307
Performing a one-time Hyper-V export	308
Setting up continual export to Hyper-V	310
Exporting data to a VirtualBox virtual machine	314
Performing a one-time VirtualBox export	314
Setting up continual export to VirtualBox	316
Exporting data to an Azure virtual machine	319
Working with Microsoft Azure	320
Azure interface disclaimer	320
Country codes used on the Azure website	320
Before virtual export to Azure	321
About Azure storage accounts	324
Creating an Azure storage account	325
Creating a container in an Azure storage account	327
Creating an Azure Active Directory web application	328
Obtaining the application ID for an Azure web application	329
Obtaining Azure subscription information	330
Obtaining the directory ID for your Azure web application	330
Obtaining a secret key for your Azure web application	331
Microsoft Azure documentation	331
Exporting and deploying VMs for Azure	332
Performing a one-time Azure export	333
Setting up continual export to Azure	338
Deploying a virtual machine in Azure	342
Managing exports	344
Restoring data	348
About restoring data with Rapid Recovery	348
Understanding Live Recovery	349
Restoring data from recovery points	349
VMware VM configuration backup and restore	350
About the file search and restore feature	351
Searching guidelines	351
Restoring guidelines	352
Finding and restoring a file	352
About restoring volumes from a recovery point	354

Restoring volumes from a recovery point	354
Restoring a directory or file using Windows Explorer	357
Restoring a directory or file and preserving permissions using Windows Explorer	358
Restoring clusters and cluster nodes	358
Performing a restore for CCR and DAG (Exchange) clusters	358
Performing a restore for SCC (Exchange, SQL) clusters	359
Restoring from an attached archive	359
Mail Restore in Rapid Recovery	360
Mail Restore prerequisites	361
Granting the required permissions in Microsoft Exchange Server	361
Opening an Exchange database in Rapid Recovery Core	362
Restoring a mail item in Rapid Recovery	363
Bare metal restore	367
About bare metal restore	367
Differences in bare metal restore for Windows and Linux machines	368
Prerequisites for performing a bare metal restore for Windows or Linux machines	372
Managing a Windows boot image	373
Understanding driver injection in a boot CD	374
Using UltraVNC for remote access	374
Creating a boot CD ISO image	374
Transferring the boot CD ISO image to media	376
Loading the boot CD and starting the target machine	377
Managing a Linux boot image	377
About the boot ISO image for Linux	378
Downloading a boot ISO image for Linux	378
Saving the Live DVD ISO image to media	379
Loading the Live DVD and starting the target machine	379
Connecting to the BMR target from the Rapid Recovery Core	379
Performing a bare metal restore using the Restore Machine Wizard	380
Using the Universal Recovery Console for a BMR	385
About Windows Universal Recovery Console tools	386
Loading drivers using the Universal Recovery Console	387
Loading drivers in the Universal Recovery Console using portable media	387
Loading a driver in the URC using Chromium	388
Selecting a recovery point and initiating a BMR	388
About disk mapping for a bare metal restore	389
Automatically mapping disks for a BMR	390
Manually mapping disks for a BMR	390
Performing a BMR from an archive	391
Loading drivers to the operating system	394
Performing a bare metal restore for Linux machines	395
Managing Linux partitions	396

Creating partitions on the destination drive	397
Formatting partitions on the destination drive	397
Mounting partitions from the command line	397
Launching a bare metal restore for Linux	397
Starting the Screen utility	398
Launching a bare metal restore for a Linux machine using the command line	398
Restoring volumes for a Linux machine using the command line	400
Verifying a bare metal restore	402
Viewing the recovery progress	402
Starting a restored target server	402
Troubleshooting connections to the Universal Recovery Console	403
Repairing boot problems	403
Verifying the bare metal restore from the command line	403
Performing a file system check on the restored volume	404
Using the command line to make a restored Linux machine bootable	404
Managing aging data	407
Data retention and archiving	407
Managing retention policies	407
Configuring Core default retention policy settings	408
Customizing retention policy settings for a protected machine	411
Forcing rollover for a protected machine	414
Archiving	415
Understanding archives	415
Archive creation and storage options	416
Amazon storage options and archiving	416
Recovery point chain options for archives	417
Methods to access an archive	417
Uses for archives	418
Creating an archive	418
Archiving to a cloud	422
Editing a scheduled archive	422
Pausing or resuming a scheduled archive	425
Forcing an archive job	425
Checking an archive	425
Attaching an archive	427
Detaching an archive	428
Importing an archive	428
Cloud accounts	431
About cloud accounts	431
Considering cloud storage options	432

Secure cloud accounts for US Government	432
Balancing access time, cost and convenience for archiving to cloud accounts	432
Adding a cloud account	433
Editing a cloud account	437
Removing a cloud account	438
Core Console references	439
Viewing the Core Console user interface	439
Button bar	441
Icon bar	443
Left navigation menu	445
Viewing protected machines	448
Viewing events for a protected machine	449
Viewing the More menu for a protected machine	451
REST APIs	452
Intended audience	452
Working with Rapid Recovery REST APIs	452
Downloading and viewing Core and Agent APIs	453
Recommended additional reading	454
Glossary	1
About us	463
Technical support resources	463

Introduction to Rapid Recovery

Rapid Recovery is a backup, replication, and recovery solution that offers near-zero recovery time objectives and recovery point objectives. Rapid Recovery offers data protection, disaster recovery, data migration and data management. You have the flexibility of performing bare-metal restore (to similar or dissimilar hardware), and you can restore backups to physical or virtual machines (VMs), regardless of origin. Rapid Recovery lets you create backup archives to a wide range of supported systems including archiving to the cloud. With Rapid Recovery, you can replicate to one or more targets for added redundancy and security.

Rapid Recovery offers:

- **Flexibility.** You can perform universal recovery to multiple platforms, including restoring from physical to virtual, virtual to physical, virtual to virtual, and physical to physical.
- **Cloud integration.** You can export a VM, archive and replicate to the cloud, and perform bare metal restore from archives in the cloud. Compatible cloud services include Microsoft Azure, Amazon Web Services (AWS), any OpenStack-based provider (including Rackspace), and Google Cloud. US government-specific platforms include AWS GovCloud (US) and Azure Government.
- **Intelligent deduplication.** You can reduce storage requirements by storing data once, and referencing it thereafter (once per repository or encryption domain).
- **Live Recovery.** Using the Live Recovery feature of Rapid Recovery Agent, you have instant access to critical data first, while remaining restore operations complete in parallel. You can use Live Recovery to restore data from a recovery point of any non-system volume of a Windows machine, physical or virtual. Live Recovery is not supported for agentlessly protected machines, Linux machines, or cluster-shared volumes.
- **File-level recovery.** You can recover data at the file level on-premises, from a remote location, or from the cloud.
- **File-level search.** Using criteria you specify, you can search a range of recovery points for one or more files. From the search results, you can then select and restore the files you want to the local Core machine directly from the Rapid Recovery Core Console.
- **Virtual machine export.** Rapid Recovery supports one-time virtual export, letting you generate a VM from a recovery point; and virtual standby, in which the VM you generate is continually updated after each backup. Compatible VM hypervisors include VMware vCenter/ESXi, VMware Workstation, Microsoft Hyper-V, Oracle VM VirtualBox, and Microsoft Azure. You can even perform virtual export to Hyper-V cluster-shared volumes.
- **Rapid Snap for Virtual support.** Enhanced support for virtualization includes agentless protection for vCenter/ESXi VMs and for Hyper-V VMs. Rapid Snap for Virtual includes protection and autodiscovery for VMware ESXi 6.0 and higher with no software agent installed. Host-based protection supports installing Rapid Recovery Agent on a Microsoft Hyper-V host only, letting you agentlessly protect all its guest VMs.

- **Application support.** Rapid Recovery is built with application support. When you protect SQL Server or Microsoft Exchange machines (whether using Rapid Recovery Agent or agentless protection), the backup snapshots captured are automatically application-aware; open transactions and rolling transaction logs are completed and caches are flushed to disk before creating snapshots. Specific application features are supported, including SQL attachability checks (for SQL Server) and database checksum and mountability checks (for Exchange Server). If you protect Oracle 12c or 18c servers with Rapid Recovery Agent, you can also perform DBVERIFY database integrity checks.

See the following resources for more information about Rapid Recovery.

- The Rapid Recovery product support website at <https://support.quest.com/rapid-recovery/>.
- The documentation website at <https://support.quest.com/rapid-recovery/technical-documents/>.

The Core Console

The Rapid Recovery Core Console is a web-based user interface (UI) that lets you fully manage your Rapid Recovery Core, including its protected machines and clusters, replicated machines, and so on.

This section describes the different elements of the Rapid Recovery Core Console UI.

i | **NOTE:** For more detailed information about the Core Console, see [Core Console references](#).

Topics include:

- [Accessing the Rapid Recovery Core Console](#)
- [Understanding the Quick Start Guide](#)
- [Navigating the Rapid Recovery Core Console](#)
- [Viewing the Protected Machines menu](#)
- [Viewing replicated machines from the navigation menu](#)
- [Viewing the Recovery Points Only menu](#)
- [Viewing the Custom Groups menu](#)
- [Using the Error dialog box](#)

Accessing the Rapid Recovery Core Console

To access the Rapid Recovery Core Console, perform one of the following actions:

- Log in locally to your Rapid Recovery Core server, and then double click the **Core Console** icon on your desktop.
- Or, type the following URL in your web browser...
 - `https://<yourCoreServerName>:<port>/apprecovery/admin/` or
 - `https://<yourCoreServerIPAddress>:<port>/apprecovery/admin/`

... where you replace the server name or IP address with yours, and insert the appropriate port number for the Rapid Recovery service.

For example, if your company domain name is **companyabc.com**, and you use the default port of **8006**, the appropriate URL to access your Core is `https://companyabc.com:8006/apprecovery/admin`.

i | **NOTE:** Since the Rapid Recovery Core Console UI depends on JavaScript, the web browser you use to access the Core Console must have JavaScript enabled.

Understanding the Quick Start Guide


The Quick Start Guide is a feature that provides you with a guided flow of suggested tasks for configuring and using Rapid Recovery Core. As you navigate through the guide, animated pages identify where in the Rapid Recovery Core Console you can perform important tasks. You are not required to perform the steps suggested by the guide. You can simply view the suggested tasks, navigating through them using the **Skip Step** and **Back** options.

The Quick Start Guide appears automatically the first time you upgrade or install the Rapid Recovery Core software and navigate to the Core Console. When it first appears, the Quick Start Guide begins with the *Welcome* page. From the *Welcome* page, click **Start Guide** to run through the guide to get an understanding of key tasks you can perform. Once started, navigate through the guide to see the various suggested configuration tasks. When you have seen the last suggested task, click **Finish** to close the guide.


Optionally, if you don't want to view this feature now, you can click **Close** from the *Welcome* page to exit the guide. You can also click **Exit Guide** from any page in the guide to close this feature.

The Quick Start Guide, with the *Welcome* page showing, reappears each time you return to the *Home* page until you do any one of the following:

- Navigate through at least one page of the guide.
- Select **Don't show again** on the *Welcome* page.





Thereafter, the *Welcome* page no longer appears. However, you can launch the Quick Start Guide again at any time by selecting  **Quick Start Guide** from the **Help** menu in the Core Console.











If you choose to perform any configuration tasks suggested by the Quick Start Guide, follow the prompts indicated in any step of the guide, and the appropriate wizard or relevant area in the user interface appears. Procedures to complete each task suggested by the guide are described in this document, as indicated in the table below.

 **NOTE:** Not all configuration tasks suggested by the Quick Start Guide are required for all users. You must understand which tasks you want to accomplish for your specific needs.


The Quick Start Guide addresses the following configuration tasks:



Table 1: Quick Start Guide configuration tasks

Function	Short Description	Result of Selecting Task, Link to Procedure
Protection or Deployment	Protecting a single machine, protecting a server cluster, or protecting multiple machines simultaneously. You can also deploy Rapid Recovery Agent to one or more machines on your network.	Click Protect or select  Protect Machine from the button bar drop-down menu to open the Protect Machine Wizard. For information about completing the Protect Machine Wizard, see Protecting a machine .
		Select  Protect Cluster from the button bar drop-down menu to open the Protect Cluster Wizard. For more information about protecting a cluster, see Protecting a cluster .
		Select  Protect Multiple Machines from the button bar drop-down menu to open the Protect Multiple Machines Wizard. For information about completing the Protect Multiple Machines Wizard, see About protecting multiple machines .
		Select  Deploy Agent Software from the button bar drop-down menu to deploy the latest available version of Rapid Recovery Agent to machines on your network, active directory domain, hypervisor host or cluster. Unless the machine is already protected in your Core, deployment only makes the












Function	Short Description	Result of Selecting Task, Link to Procedure
		software available to that machine; you must still explicitly add it to protection on your Core.
Replication	Setting up replication from a primary (source) Core to a secondary (target) Core	Click  Replication from the icon bar to open the <i>Replication</i> page. Prompts you to add a target Core using the Replication Wizard. For information about using the Replication Wizard to set up replication on a self-managed Core, see Replicating to a self-managed target Core . For general information about replication, see Configuring replication .
Virtual Export	Performing a one-time export or establishing continual export from a protected machine to a virtual machine	Click  VM Export from the Restore button bar drop-down menu to perform an export of data from your protected machine to a virtual machine. You can either perform a one-time export, or set up virtual standby for continual export to a VM. For information about virtual exports, see Exporting to virtual machines using Rapid Recovery .
Manage and Configure	Allows you to set up additional configuration for the Rapid Recovery Core	Click  More from the icon bar to display a drop-down menu of additional functions you can manage or configure. Functions include seeing system information or virtual environment data; managing archives; mounts; boot CDs; repositories; encryption keys; cloud accounts; file search; retention policies; the credentials vault; notifications; mail restore settings; available downloads; reports; and log files.
Configure Encryption	Adding or importing encryption keys that you can use for one or more protected machines	Select  Encryption keys from the  (More) drop-down menu to manage security for protected data by adding or importing encryption keys. You can apply encryption keys to one or more protected machines. Encryption is described in the topic Encryption .
Configure Notifications	Setting up notifications for events, warnings and alerts	Select  Notifications from the  (More) drop-down menu to specify notification groups for events, warnings, and alerts. To send these by email, you must also establish SMTP server settings. For more information about managing events, see the topic Events , including the topics Configuring notification groups and Configuring an email server .
Manage Retention	Viewing or changing the default retention policy for the Core	Select  Retention Policy from the  (More) drop-down menu to open the <i>Retention Policy</i> page for the Core. From here you can define how long to keep a recovery point before rolling it up. For conceptual information about retention policies, see the topic Managing aging data . For procedural information, see Managing retention policies .
Restore	Restoring data from a recovery point on the Core	Click  Restore from the button bar to open the Restore Machine Wizard. For information about restoring data, see the topic About restoring volumes from a recovery point .

Navigating the Rapid Recovery Core Console


When you log into the Core Console, and any time you click the **Home**  icon, the *Home* page appears. The *Home* page gives you a view of your Rapid Recovery Core, with two view options for displaying content about your Core. You can toggle between these two viewing options by clicking the appropriate icon at the top right side of the page title area. The content viewing options are described in the following table:

Icon	Name	Page Title Shows...	Description
	Summary Tables view	The display name of your Rapid Recovery Core appears as the page title in this view.	<p>Summary tables describe information about your Core. Each of the following is summarized:</p> <ul style="list-style-type: none">• Machines protected on your Core• DVM repositories associated with your Core.• Recent alerts about activity in your Core. <p>For more information, see Understanding the Home page (summary tables view).</p>
	Dashboard view	"Dashboard" appears as the page title in this view.	<p>The Dashboard displays a set of real-time reports on your system. Default dashboard reports include a trouble monitor showing notifications; transfer job status for the last 24 hours and transfer jobs per machine; a repository overview, and connectivity state for protected, replicated, and recovery points-only machines.</p> <p>For more information, see Understanding Core dashboard reports.</p>

On the *Home* page (and on every page in the Core Console), the left navigation area shows the items that are protected on your Core. You can navigate to other pages in the UI by doing one of the following:

- Clicking the corresponding icon from the icon bar in the left navigation area. The options accessible from the icon bar include  Home,  Replication,  Virtual Standby,  Events,  Settings, and  More.
- Expanding the  (More) menu on the icon bar, and then selecting a destination.
- Clicking a button or menu option from the button bar. Buttons include  Protect,  Restore,  Archive, and  Replicate.

When you select an item from the left navigation area, the focus of the Core Console changes to display summary information about that item. For example, if you click the name of a protected machine, the Core console displays information about that machine only, rather than the Core. In this example, the display name of the protected machine appears as the title of the page. A submenu appears to the right, letting you view specific information about the protected machine. The menu options include: Summary, Recovery Points, Events, Settings, Reports, and More.

To return to viewing information about the Core, including dashboard reports, or a summary view of multiple protected or replicated machines, click on the **Home**  icon on the top left of the UI.

You can use the title at the top of the Core Console to provide context for the information you are viewing in the Core. For example:

- Any time you see the display name or IP address of the Core as the page title, you are viewing summary information about the Core.
- If the title is "Dashboard," you are viewing the Core dashboard.
- If you see the display name or IP address of a protected machine, or a Summary pane at the top of a page, you are viewing information about a single machine protected by or replicated in the Core.
- If you see the title "Protected Machines," you are viewing information about all of the machines protected in the Rapid Recovery Core.
- If you see the title "Machines replicated from...," you are viewing information about all of the machines replicated in the Rapid Recovery Core.
- If you see the page title "Recovery Points Only," you are viewing information about all of the recover points-only machines on this Core.

For information about the features and functions available from each page, see the appropriate section below.

For more information about viewing protected machines, see [Viewing the Protected Machines menu](#). For more information on managing protected machines, see [Managing protected machines](#).

For more information about viewing replicated machines, see [Viewing incoming and outgoing replication](#).

For more information about viewing recovery-points only machines, see [Viewing the Recovery Points Only menu](#).

Understanding the left navigation area

The left navigation area of the Core Console appears on the left side of that user interface. The contents of this navigation area may differ based on the type of objects protected in your Rapid Recovery Core.

The left navigation area always contains the following:

- **Icon bar.** For navigation among the main pages of the Core Console.
- **Text filter.** The text filter is a text field that lets you filter the items displayed in the various menus that appear below it. Clicking the arrow to the right of the text filter expands and collapses each of the menus appearing below it.

Following these elements, the left navigation area typically displays menus to help you navigate, filter, and view the objects protected on your Core. This includes protected machines, replicated machines, and so on.


Each menu is context-sensitive; that is, each menu only appears in the Core Console if it is relevant. For example, if you protect at least one machine, the Protected Machines menu appears, and so on.

For more information, see the left navigation area tables in [Viewing the Core Console user interface](#).



Related topics include the following:

- [Viewing summary information for a protected machine](#)
- [Viewing summary information for a hypervisor or cluster host](#)
- [Replication](#)
- [Understanding custom groups](#)

Viewing the Rapid Recovery Core Console Home page

Each time you log into the Rapid Recovery Core Console, or each time you click on the **Home**  icon in the icon bar, the Home page appears.

The Home page of the Core Console offers a **summary tables** (default) view and a **dashboard** view.

You can toggle between views on the Home page by clicking the  **Summary Tables view** and  **Dashboard view** icons at the top right of the title bar on the Home page.



From the *Home* page, and every other page of the Core Console, you can navigate to the functions you want by using the left navigation area.

For more information, see the following topics:

- [Understanding the left navigation area](#)
- [Understanding Core dashboard reports](#)
- [Understanding the Home page \(summary tables view\)](#)

Understanding the Home page (summary tables view)


The Home page is only applicable to the Core. In the dashboard view, it shows real-time graphical reports. When you switch to the summary tables view, the Home page displays all of the machines the Core protects or replicates, the repositories associated with your Core, and alerts for machines on this Core.

The view of each pane on the Home page can be expanded or contracted. For example, if you click the  (contract view) icon at the top right-hand side of the Protected Machines pane, the view of protected machines contracts, and only the name of the pane is visible. To expand the view to see all protected machines again, click the  (expand view) icon.

The following table describes the various elements on the Home page when in the summary tables view.

Table 2: Home page options

UI Element	Description
Protected Machines	<p>The Protected Machines pane lists the machines that this Core protects. This pane appears regardless of whether any machines have been added to your Core for protection.</p> <p>This section includes the following information for each protected machine:</p> <ul style="list-style-type: none">• Protection icon. An icon shows whether the machine is protected.• Status. Colored circles in the Status column show whether the protected machine is online and protected, paused, or offline and unreachable.• Display Name. The display name or IP address of the protected machine.

UI Element	Description
	<ul style="list-style-type: none"> • Additionally, if virtual standby has been set up for this machine, place your cursor over the  Virtual Standby icon to see a summary of the following: <ul style="list-style-type: none"> • Synchronization status • Last export • Export type • Export destination • Repository Name. The display name of the repository storing the recovery points for that machine. • Last Snapshot. The date and time on which Rapid Recovery took the most recent recovery point snapshot for that machine. • Next Snapshot. The date and time on which Rapid Recovery is scheduled to capture the next recovery point for that machine. • Recovery Points. The number of recovery points stored in the repository and space usage for each protected machine. • Version. The version of the Rapid Recovery Agent software installed on that machine. <p>If you click on a specific machine name shown in this pane, a <i>Summary</i> page appears, showing summary information for the selected machine. For more information on what you can accomplish on the <i>Summary</i> page, see Viewing summary information for a protected machine.</p>
Replicated Machines	<p>The Replicated Machines pane lists any machines that this Core replicates from another Core. This pane does not appear unless your Core replicates machines from another Core. This section includes the following information for each replicated machine:</p> <ul style="list-style-type: none"> • Machine type. An icon shows whether the machine is a physical machine, virtual machine, or a protected cluster. • Status. Colored circles in the Status column show whether the replicated machine is accessible, paused, or offline and unreachable. • Display Name. The display name or IP address of the replicated machine. • Replication Name. The display name of the originating source Core for any machines you replicate on this target Core. You can define this name when setting up replication. • Repository Name. The name of the repository storing the recovery points for that machine. • Last Replicated Snapshot. The date and time on which Rapid Recovery took the most recent replica of the original protected machine. • Recovery Points. The number of recovery points stored in the repository and space usage for each replicated machine. • Version. The version of the Rapid Recovery Agent software installed on that machine. <p>If you click on a specific machine name shown in this pane, the <i>Summary</i> page appears, showing summary information for that replicated machine.</p>



UI Element	Description
Recovery Points Only Machines	<p>The Recovery Points Only Machines pane lists any machines that were removed from protection or replication, if the recovery points have been retained. These machines can be used for file-level recovery, but cannot be used for bare metal restore, for restoring entire volumes, or for adding snapshot data. This pane does not appear unless you have any machines that meet this definition.</p> <p>This section includes the following information for each recovery points only machine:</p> <ul style="list-style-type: none"> • Machine type. An icon shows whether the machine is a physical machine, virtual machine, or a protected cluster. • Status. Colored circles in the Status column show whether the recovery points only machine is accessible, paused, or offline and unreachable. • Display Name. The display name or IP address of the machine for which you kept recovery points. • Repository Name. The name of the repository storing the remaining recovery points for that machine. • Recovery Points. The number of recovery points stored in the repository and space usage for each recovery points-only machine. <p>If you click on a specific machine name shown in this pane, the <i>Summary</i> page appears for this recovery points only machine.</p>
DVM Repositories	<p>This pane does not appear unless your Core has one or more DVM repositories. It includes the following information for each DVM repository:</p> <ul style="list-style-type: none"> • Type. An icon depicts a repository. • Status. Colored circles in the Status column show whether the repository is mounted and can accept recovery point transfers, or is unreachable, or in an error state. • Repository Name. The display name of the repository. • Space Usage. The total amount of space used in the repository, and the size of the storage volume or extent. • Protected Data. The amount of used space in the repository. • Machines. The number of machines for which the repository stores recovery points. • Recovery Points. The number of recovery points stored in the repository. • Compression Ratio. The rate at which the repository compresses the protected data to save space. <p>For more information, see Understanding repositories.</p>
Alerts	<p>This section lists the important alerts for the Core and every machine it protects. The section includes the following information:</p> <ul style="list-style-type: none"> • Icons. The column of icons indicates the nature of the alert. These include informational messages, errors and warnings. • Date. Displays the date and time of when Rapid Recovery issued the alert.

UI Element	Description
	<ul style="list-style-type: none"> • Message. Describes the alert. You can also see these details on the <i>Core Events</i> page. For more information, see Viewing events using tasks, alerts, and journal pages.

Understanding Core dashboard reports

The Core dashboard displays a set of real-time graphical reports of data relevant to your Core, the machines you protect, and the snapshots saved to your repository. Dashboard reports include:

- **Core System reports**
 - **Trouble Monitor.** This report shows job activity, connections with the license portal, and transfer activity to detect trouble early on your system. Click the clock widget to clear all activity tracked and monitor new events.
 - **Repository.** This report shows the repositories associated with your Core. It shows the number of repositories, how many machines are protected in each, the number of recovery points and the percentage of compression or deduplication. This report is refreshed every minute.
- **Machine reports**
 - **Machine Connectivity.** This report shows the connectivity state of machines protected and replicated on your Core. It also shows connectivity for data on a recovery points-only machine.
- **Snapshot reports**
 - **Transfer Job.** This report shows all snapshot data transfers (including base images and incremental snapshots) that completed in the last 24 hours. Snapshots include base images and incremental snapshots. This dashboard report appears as a circle graph.
 - **Transfer Job per Machine.** This job shows, by protected machine, the number of successful and failed transfer jobs in the last 24 hours. This dashboard report appears as a line graph.



At the top of the Dashboard pane, you can click the  (Expand) widget to display a list of reports displayed on the dashboard, in the three categories described above. Each report category and report has a check box. If the box is selected, the report or category appears on the dashboard. If you clear the option, the report does not display. In this way, you can easily change the display of reports on the dashboard. If you click  (Contract), the menu disappears.


You can also collapse or expand the view of any reports on the dashboard by clicking the up or down arrow in the header of the report. Some dashboard reports (machine connectivity and repository) have a plus sign next to the arrow, from which you can add another protected machine or another repository, respectively.

You can also drag and drop to move the location of one of the reports elsewhere on the dashboard, to order the reports in a manner most effective for your use.



Viewing the Protected Machines menu

In the Rapid Recovery Core Console user interface, a Protected Machines menu appears in the left navigation area. As with all menu labels in the navigation area, the label for this menu appears in all upper-case letters. By default, this menu is fully expanded, and shows a list of any machines that are protected by this Core. If you have any server clusters protected, then they are included in this list.

You can collapse or expand the view for protected machines and server clusters in your Core by clicking the  [Contract menu] or  [Expand menu] arrows on the left side of this menu.

The Protected Machines menu includes a drop-down menu on the right side which lists functions that can be performed on all protected machines. Click the  (More) icon to the right of **Protected Machines** to see the menu. Each machine listed under the Protected Machines menu also has a drop-down menu that controls functions only for that machine.


If you are managing server clusters from the Rapid Recovery Core, the cluster also appears in the left navigation menu. From the drop-down menu for any cluster, you can also navigate to the *Protected Nodes* page for the selected cluster.

If you click the  [Contract menu] to the left of the Protected Machines menu, the list of protected machines and server clusters contracts, and no machines are listed. Clicking again on the  [Expand menu] arrow causes the list of machines to expand again.

Clicking any machine name in the Protected Machines menu opens the *Summary* page for that machine. For more information on what you can accomplish on the *Summary* page, see [Viewing summary information for a protected machine](#).

Finally, clicking directly on the **Protected Machines** menu causes the *Protected Machines* page to appear in the main content area, with a single pane showing protected machines on this Core. For more information on what you can accomplish on the Protected Machines pane of the *Protected Machines* page, see [Viewing protected machines](#).



NOTE: From the Protected Machines page, you can return to a view from the Core perspective by clicking the  **Home** icon in the icon bar.

Viewing summary information for a protected machine

When you click the name of a protected machine in the Core Console, the Summary page appears. When displaying information for a protected machine—on the *Summary* page and all other views—there is a menu at the top of the page with functions you can perform. This menu appears immediately below the name of the protected machine.

On the *Summary* page, at minimum, is a [Summary](#) pane, and a [Volumes](#) pane. At the bottom of the Summary pane, to view system information for the protected machine, click **System Information**.

If a machine is added to replication, a [Replication](#) pane also appears.

If you have one or more protected Exchange servers, you will also see an [Exchange Server Information](#) pane that contains information about your protected Exchange server.

If you have one or more protected SQL servers, you will also see a [SQL Server Information](#) pane that contains information about your protected SQL servers.

Related topics:

- [Viewing the Summary pane](#)
- [Understanding system information for a protected machine](#)
- [Viewing Volumes on a protected machine](#)
- [Viewing replication information](#)

- [Viewing the Exchange Server Information pane](#)
- [Viewing the SQL Server Information pane](#)

Viewing the Summary pane

The *Summary* pane contains summary information about the protected machine, including the host name, date and time of the last snapshot, date and time of the next scheduled snapshot, encryption key information, and version information for the Rapid Recovery Agent software. There is also a link to a detailed *System Information* page for the machine.

Viewing Volumes on a protected machine

For any protected machine, from the *Summary* page, in the Volumes pane, you can perform the following actions for any of the volumes listed:

- **Set a protection schedule for a selected volume.** Protection schedules are typically established when you first protect a machine. For more information about modifying a protection schedule, see [Creating custom protection schedules in Simple Mode](#).
- **Force a base image or snapshot.** Snapshots typically occur based on the protection schedule. However, at any time, you can force a base image or an incremental snapshot for selected volumes. For more information, see [Forcing a snapshot](#).

Viewing replication information

The Replication pane contains summary information about the replicated machine, including the replication name, the state of replication, progress, and available space.

Viewing the Exchange Server Information pane

The *Exchange Server Information* pane appears only for protected machines that are Exchange servers.

This pane contains summary information about the protected Exchange server, including the installed version of Microsoft Exchange, the path in which Exchange is installed, and the path defined for Exchange mailbox data.

The Mail Stores grid shows the Exchange Database (EDB) name, the path of the EDB file, the path in which the log files are stored, the log prefix, the system path, the Database Availability Group (DAG), and the mail store type.

Viewing the SQL Server Information pane

The *SQL Server Information* pane appears only for protected machines that are SQL Servers.

This pane contains summary information about the protected SQL Servers. You can expand the database information to see detail for each table in the database. You can also see the database or table name and the database path.

Viewing summary information for a hypervisor or cluster host

When you click the name of a hypervisor or cluster host machine in the Core Console, the *Summary* page displays. On this page, at minimum, is an actions bar of functions, a Summary pane, and a Processors pane. The actions bar displays at the top of all pages for the host. The Summary pane includes information such as the host name and the virtualization software. The Processors pane includes a table which lists the architecture, number of Cores, number of threads, clock speed, and description for each processor.

Depending on the type of host the machine is, the following other panes may also display:

- **SMB Servers.** If an agentlessly protected Hyper-V host uses one or more Server Message Block (SMB) servers, this section lists the host name for each SMB server.
- **Nodes.** The *Summary* page of a cluster host includes the section for listing the host name and Rapid Recovery version number for each node.
- **Volumes.** When using agentless protection, the *Summary* page for a Hyper-V host includes this section to list the name, file system, space usage, current schedule, and next snapshot for each shared or protected volume.
- **Shared Disks.** If a Hyper-V host using agentless protection has one or more shared VHDXs, this section lists the name and path for each virtual hard disk.

For information about CSV clusters, see "Support for Cluster Shared Volumes" in the *Rapid Recovery System Requirements Guide*. For information about other related topics, see the following links:

- [Viewing summary information for a protected machine](#)
- [Support for Hyper-V guest cluster](#)
- [Understanding Rapid Snap for Virtual](#)

Viewing recovery points for a machine

The *Recovery Points* page shows a list of the recovery points collected for that protected machine as well as pertinent machine and repository data. On this page, you can archive, mount, export, and restore specific recovery points, as well as delete recovery points.

The page is divided into two panes: Recovery Points Summary and Recovery Points, which contains detailed information. The Summary pane does not include any actionable links.

For a more detailed explanation of the summary and detailed recovery points information on this page, see [Viewing protected machines](#).


Viewing events for a protected machine

On the *Events* page, you can view the jobs that occurred or are in progress for the protected machine you selected. Buttons at the top of the page let you navigate to lists of jobs in each of the three categories of activities:

- **Tasks.** A job that the Rapid Recovery Core must perform to operate successfully.
- **Alerts.** A notification related to a task or event that includes errors and warning.
- **Journal.** A composite of all protected machine tasks and alerts.


The following table includes descriptions of each element on the *Events* page.

Table 3: *Events* page elements

UI Element	Description
Search keyword	Lets you search for a specific item within each category. Available for tasks only.
From	To narrow your results, you can enter a date at which to begin searching. Available for tasks only.
To	To narrow your results, you can enter a date at which to stop searching. Available for tasks only.
Status icons	<p>Each icon represents a different job status. For alerts and tasks, clicking one of the icons lets you filter the list by that status, essentially generating a report. Clicking the icon a second time removes the filter for that status. You can filter by more than one status. Statuses include:</p> <ul style="list-style-type: none">• Active. A job that is in progress.• Queued. A job that is waiting for another job to complete before it can initiate.• Waiting. A job waiting for your approval or completion, such as a seed drive. (For more information about seed drives, see Replication.)• Complete. A job that completed successfully.• Failed. A job that failed and did not complete.
Service icon	This button adds services jobs to the list of jobs. When you click this icon, a smaller service icon appears on each status icon, which lets you filter by service jobs that have those statuses (if any exist). Examples of services jobs include deleting index files or removing a machine from protection.
Export type drop-down list	<p>The drop-down list includes the formats to which you can export the event report. Available for tasks only. It includes the following formats:</p> <ul style="list-style-type: none">• PDF• HTML• CSV• XLS• XLSX
 (Export icon)	Converts the event report to the format you selected. Available for tasks only.
Page selection	Event reports can include several jobs across multiple pages. The numbers and arrows at the bottom of the <i>Events</i> page let you navigate the additional pages of the report.

The *Events* page displays all events in a table. The following table lists the information shown for each item.

Table 4: Detailed information for the Event summary table

UI Element	Description
Status	Shows the status for the task, alert, or journal item. Available for alerts or journal items, click the header to filter the results by status.
Name	Name is available for tasks only. This text field lists the task type that completed for this protected machine. Examples include transfer of volumes, maintaining repository, rolling up, performing mountability checks, performing checksum checks, and so on.
Start Time	Available for tasks, alerts, and journal items. Shows the date and time when the job or task began.
End Time	Available for tasks only. Shows the date and time when the task completed.
 Job Details	Available for tasks only. Opens the <i>Monitor Active Task</i> dialog box, so you can view details of the specific job or task. These details include an ID for the job, rate at which the Core transferred data (if relevant), elapsed time for the job to complete, total work in amount of gigabytes, and any child tasks associated with the job.
Message	Available for alerts and journal items. This text field provides a descriptive message of the alert or journal item.

Viewing reports for a protected machine

The Reports ▼ drop-down menu lets you generate reports on demand for the selected protected machine.

- The Job report provides a report on the status of successful jobs and failed jobs for the selected machine. Failed jobs can be further viewed in a Failure report. For more information on this report type, see [Understanding the Job report](#).
- The Failure report provides information on failed and canceled Core jobs for the specified machine. For more information on this report type, see [Understanding the Failure report](#).

For more information about generating these reports, see [Generating a Core report on demand](#).

Viewing replicated machines from the navigation menu

If your Core replicates machines from another Rapid Recovery Core, the display name of the source Core appears as a collapsible menu in the left navigation of the Core Console. As with all menu labels in the navigation area, this replicated machines menu name appears in all upper-case letters, below the Protected Machines menu. By default, the replicated machines menu is fully expanded, and lists all machines originating from that source Core that are replicated on your target Core.

You can collapse or expand the view of replicated machines from that source Core by clicking the arrow on the left side of this menu.

Each replicated machines menu includes a drop-down menu on the right side, which includes functions you can perform simultaneously on all of the replicated machines originating from that Core. Click the arrow to the right of replicated machines menu to see a drop-down list of functions you can perform. These actions include the following:

- **Pause replication.** If replication is currently active, it stops the action until you resume it. For more information, see [Pausing and resuming replication](#).
- **Resume replication.** If replication has been paused, it begins replicating again. For more information, see [Pausing and resuming replication](#).
- **Force replication.** Replicates on demand, rather than at a scheduled time. For more information, see [Forcing replication](#).
- **Remove replication.** Removes the replication relationship between the source Core and your target Core. Optionally, you can delete the recovery points stored in this Core. For more information, see topics [Removing outgoing replication from the source Core](#) and [Removing incoming replication from the target Core](#).

Clicking directly on name of the source Core in the navigation menu causes the Machines replicated from [Source Core Name] page to appear in the main content area. For more information on what you can accomplish on that page, see [Viewing incoming and outgoing replication](#).

Viewing the Recovery Points Only menu

The Recovery Points Only menu appears in the left navigation area if one of the following is true:

- if your Rapid Recovery Core Console retains some recovery points from a machine that was previously protected.
- if you removed replication but retained the recovery points.

As with all menu labels in the navigation area, the label for this menu appears in all upper-case letters.

You can collapse or expand the view of recovery points-only machines by clicking the ▼ [Contract menu] or > [Expand menu] arrows on the left side of this menu.

The menu includes a drop-down menu on the right side which lists functions that can be performed on all recovery points-only machines simultaneously. In this case, the only function you can perform is to remove recovery points from the Core.

! CAUTION: This action removes all of the recovery points-only machines in your Rapid Recovery Core, permanently deleting them and precluding you from restoring information from those recovery points from this Core.

Viewing the Custom Groups menu

The custom groups menu appears in the left navigation area only if you have defined one or more custom groups.

As with all menu labels in the navigation area, the label for this menu appears in all upper-case letters.

You can collapse or expand the view of items in this menu by clicking the arrow on its left side.

The custom groups menu includes a drop-down menu on the right side which lists functions that can be performed simultaneously on all of the like items in that group.

For more information, see [Understanding custom groups](#).

Using the Error dialog box

When an error occurs in the Rapid Recovery Core Console user interface, such as trying to enter an invalid parameter, an *Error* dialog box appears. The dialog box typically indicates the cause of the error, includes some links to provide more information about the error, and includes a **Close** button. You must close the *Error* dialog box before you continue, but you may want to view more information about the error.

In the *Error* dialog box, choose from the following options:

User interface errors that cause the *Error* dialog box to appear are not tracked in the Rapid Recovery *Events* page, since they are simply validation or data entry errors. However, when you click the **Search Knowledge Base** option for any error, then the URL link provided for that error is recorded to the Core AppRecovery.log file. You can search the log for the text string “KB article url generated” to see the URL for each error that was viewed in a browser. For more information on downloading or viewing Core error logs, see topics [Downloading and viewing the Core log file](#) or [Accessing Core logs](#), respectively.

Repositories

This section describes how to work with repositories. It discusses the features and attributes of the repository technology supported by Rapid Recovery release 6.7, called Deduplication Volume Manager (DVM). It briefly describes deduplication used in Rapid Recovery. Then this section describes how to manage DVM repositories, including creating a repository, viewing and editing its details, and deleting a repository. You can learn how to open a repository from one Core on another Core.

Topics include:

- [Understanding repositories](#)
- [Deduplication in Rapid Recovery](#)
- [Managing a DVM repository](#)
- [Managing an Azure repository](#)
- [About repository optimization](#)
- [Connecting to an existing repository](#)
- [Viewing or modifying repository details](#)
- [Checking a repository](#)
- [Deleting a repository](#)

Understanding repositories

A repository is a data structure used to store and manage Rapid Recovery data. Backup snapshots are saved to a repository in the form of recovery points. Before you can protect machines, replicate, or restore data in Rapid Recovery, you need at least one repository.

There are two types of repositories you can use for storing Rapid Recovery recovery points. Deduplication Volume Manager (DVM) is a Rapid Recovery technology and can be installed on the Rapid Recovery Core using the Core Console, the Command Line utility, or PowerShell. You can use it to store original recovery points or recovery points replicated from another Core. Azure repositories originate from Microsoft Azure. You can connect one Azure repository at a time to a single Rapid Recovery Core and use it to store replicated recovery points.

i **NOTE:** Quest discourages using an Azure repository for storing original recovery points and recommends using it only for storing replicated recovery points. For more information, see [Replication with Rapid Recovery](#).

On the *Repositories* page, you can create a new repository, or connect your Core to an existing repository (currently used by another Core). In general, you can view details for a repository, view repository settings, check a repository, or delete a repository, as well as add a storage location or optimize a repository. For more information about managing a repository, including how to create one, see [Managing a DVM repository](#).

DVM repositories

A Rapid Recovery administrator can create a DVM repository can reside on different storage technologies, including Direct Attached Storage (DAS), Storage Area Network (SAN), or Network Attached Storage (NAS).

i NOTE: When designating a location for a Rapid Recovery repository, speed for the storage volume is the most critical factor. Archival storage devices such as Data Domain are not supported due to performance limitations. Do not store repositories on NAS filers that tier to the cloud, as these devices tend to have performance limitations.

DAS offers the highest data bandwidth and fastest access rate, and is easy to implement. For optimum results, use DAS with Redundant Array of Independent Disks (RAID) 6 storage.

If installing on a NAS, Quest recommends limiting the repository size to 6TB when using the CIFS protocol, since CIFS is not designed as a high-I/O storage protocol. Any storage device must meet the minimum input/output requirements. For these requirements, and for additional guidance for sizing your hardware, software, memory, storage, and network requirements, see the *Rapid Recovery Sizing Guide*.

The storage location for a DVM repository should always be in a subdirectory that you specify (for example, `E:\Repository`), never in the root of a volume (for example, `E:\`).

The following list describes the features of the DVM repository type:

- DVM repositories are created and managed from the Rapid Recovery Core Console.
- DVM repositories support multiple volumes, up to 255 repositories on a single Core.
- You can specify the size of a DVM repository upon creation, and can add extents later.
- Once you specify the size of a repository, its size cannot be reduced later.
- You can create DVM repositories on machines with Windows operating systems only.
- DVM repository volume can be local (on storage attached to the Core server), or on a storage location on a Common Internet File System (CIFS) shared location.
- You can use this repository type when upgrading existing AppAssure installations, and when using new Rapid Recovery installations.
- DVM repositories can span across different storage technologies.
- Supported storage types include Storage Area Network (SAN), Direct Attached Storage (DAS), or Network Attached Storage (NAS)
- Requires 8GB RAM, preferably Error Checking and Correction (ECC) memory Rapid Recovery Core requirement)
- Requires quad core processor on Core machine
- Supports multiple DVM repositories for each host
- No additional services required; DVM repository uses system-provided Core services for communication with Core and for tracking events
- Each DVM repository supports up to 4096 repository extents (also called storage locations)
- Fixed size; DVM repository requires you to specify the repository size on a volume. The size that you specify cannot exceed the size of the volume. Each volume you define as a storage location must have a minimum of 1GB of free space available on it.
- Repository storage location can be a simple or dynamic disk, with speed the most important factor
- Can use standard encryption keys created and managed in the Core Console (Core-based encryption)

- Deduplicates data across the entire repository (or across encryption domains within each repository, if encryption keys are used)
- Uses a dedicated, resizable DVM deduplication cache, with a configurable storage location in Core settings
- Optimized for writing data, storing snapshot data in a repository local to the Core, with all data processed through the Core
- Cannot be renamed after creation
- New repositories of this type can be created using REST APIs, the Rapid Recovery Command Line Management Utility (cmdutil.exe), or Windows PowerShell cmdlet

When you create a DVM repository, the Rapid Recovery Core pre-allocates the storage space required for the data and metadata in the specified location. The minimum DVM repository size is 1GB, which for practical purposes is too small except for testing.

Since DVM deduplication requires a primary and secondary cache, ensure the storage space you reserve is twice the size of your deduplication cache. For example, if you have 1.5GB reserved in the DVM deduplication cache settings on the Core, reserve 3GB on the cache volume. The default installation path for the cache is on the C drive. For more information, see [Understanding deduplication cache and storage locations](#).

Azure repositories

A Rapid Recovery administrator can add one Azure repository to each Rapid Recovery Core. This repository best serves as a secondary destination on a replication target core. The contents of the repository can come only from another source Core, not the same Core to which it is attached.

i **NOTE:** Currently, the Rapid Recovery Core Console lets you connect to an Azure repository only from the source Core that replicates to it, and does not support the Connect to Existing feature. For more information, see <https://support.quest.com/kb/332647>.

Azure stores this type of repository under Blob Containers. The contents of the Azure repository are only accessible from a Rapid Recovery Core. While the files may be visible from the Azure user interface, only Rapid Recovery can open them.

The following list describes features of an Azure repository:

- Used as secondary storage on a replication target Core
- Azure repositories are created and managed from the Rapid Recovery Core Console
- You can specify the size of an Azure repository, and then reduce and extend the size of the repository at any time
- You can create an Azure repository on a machine with Windows operating systems only
- Requires 8GB RAM, preferably Error Checking and Correction (ECC) memory (Rapid Recovery Core requirement)
- Requires quad core processor on Core machine
- Supports one Azure repository for each Core
- No additional services required; an Azure repository uses system-provided Core services for communication with Core and for tracking events
- Can use standard encryption keys created and managed in the Core Console (Core-based encryption)

- Deduplicates data across the entire repository (or across encryption domains within each repository, if encryption keys are used)
- Uses a dedicated, resizable deduplication cache, with a configurable storage location in Core settings
- Cannot be renamed after creation

As with DVM, the recovery points in the Azure repository are compressed and deduplicated. For Azure, this process requires a location for deduplication cache and a location for metadata cache. You have the option to store the cache locally or in Azure, but you must store the metadata on your local machine.

Unlike with DVM, which has pre-allocated storage space that you can increase but not decrease, you can change the size of an Azure repository to be larger or smaller at any time. When you configure the Azure repository in Rapid Recovery, you set a space limitation that tells Rapid Recovery not to send more data after the contents exceed that amount, but you can reduce or increase the limitation at any time. For example, in Repository Settings, if you set the Azure repository to store no more than 1TB of data but you discover you need less than half that amount, you can reduce the maximum size of the Azure repository to 500GB.

Deduplication in Rapid Recovery

Deduplication is a data compression technique that reduces both storage requirements and network load. The process involves physically storing unique blocks of data only once on disk. In Rapid Recovery, when any unique data block occurs a second time within a repository, instead of storing the data again, the Core stores a reference to the first occurrence of the data in the repository. When the information is needed (for example, when restoring), the Core retrieves the data by following the references and reconstructing the original data stream.

Deduplication occurs in backup snapshots captured by Rapid Recovery Core.

- Backup information is deduplicated within a single repository. It cannot be deduplicated across multiple repositories.
- DVM repositories use target-based encryption, in which deduplication is further limited to the data protected with a single encryption key. For security purposes, each key serves as a separate encryption domain.

For maximum gains, Rapid Recovery uses different types of deduplication, as described in the following sections.

When deduplication occurs

DVM technology uses target-based deduplication.

Target-based deduplication can take place inline (during the transfer of backup information), or as post-processing (occurring on the repository). Post-processing is sometimes called pass-through deduplication.

As for when deduplication occurs, standard deduplication occurs inline.

Rapid Recovery also uses post-processing in one instance: when performing a repository optimization job. This feature is also called duplicate block reclamation.

For more information about the repository optimization job, see [About repository optimization](#). For more information about performing this task, see [Optimizing a repository](#).

Thus, Rapid Recovery takes advantage of target-based deduplication, inline deduplication, and post-processing deduplication.

For more information about where the references to unique blocks are stored for DVM repositories, see [Understanding deduplication cache and storage locations](#). For information about adjusting DVM deduplication cache settings, see [Configuring DVM deduplication cache settings](#).

Managing a DVM repository

Managing a DVM repository involves the following operations:



1. **Creating a DVM repository.** Before creating a repository, consider the appropriate technology type.
 - For information about the different repository types, see [Understanding repositories](#)
 - For information about creating a DVM repository, see [Creating a DVM repository](#).
2. **Connecting to a repository.** For more information about connecting to an existing repository currently managed by another Core, see [Connecting to an existing repository](#).
3. **Checking a repository.** For more information about checking a DVM repository, see [Checking a repository](#).
4. **Modifying repository details.** For more information about viewing repository details or modifying the details that display, see [Viewing or modifying repository details](#).
5. **Changing repository settings.** For more information about changing the settings for a DVM repository, see [Changing DVM repository settings](#).
6. **Performing DVM repository optimization.** For more information about the repository optimization job, see [About repository optimization](#). For steps to optimize an existing DVM repository, see [Optimizing a repository](#).
7. **Deleting a repository.** For more information about deleting a repository, see [Deleting a repository](#).

Creating a DVM repository

This process describes how to create a repository on your Core using the Deduplication Volume Manager (DVM) repository technology.

You must have administrative access to the machine on which you want to create a DVM repository. This process requires you to associate at least one volume on which to store data and metadata for the repository.

Complete the following steps to create a DVM repository.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More), and then select  **Repositories**.
The *Repositories* page displays.
3. At the top of the page, under the *Repositories* page title, click **+ Create**.
The *Create Repository Wizard* appears.
4. On the *Repository type* page, select **DVM**.

5. On the *Configuration* page, enter the information as described in the following table.

Table 5: Repository configuration settings

Text Box	Description
Name	<p>Enter the display name of the repository.</p> <p>By default, this text box consists of the word Repository and a number, which corresponds to the number of repositories for this Core. For example, if this is the first repository, the default name is Repository 1. Change the name as needed.</p> <p>Repository names must contain between 1 and 40 alphanumeric characters, including spaces. Do not use prohibited characters or prohibited phrases.</p>
Concurrent operations	<p>Enter the number of operations you want the repository to perform at the same time. The default number is 64.</p>
Threshold	<p>Enter the amount of space you want to reserve so that the repository does not become full. A full repository would prevent you from deleting any recovery points to regain space. The default value is 50GB.</p>
Comments	<p>Optionally, enter a descriptive note about this repository. You can enter up to 254 characters. Examples might include a message such as DVM Repository 2 or This repository contains protected SQL Server data only.</p> <p>This information can later be viewed and edited by accessing the Settings for this repository.</p>

6. Click **Next**.
The *Storage Location* page appears.

7. On the *Storage Location* page, specify the path to a storage location on a local disk or on a CIFS shared network storage location.

- To add a storage location on a local disk, enter the following information:

i **NOTE:** If installing on a NAS, Quest recommends limiting the repository size to 6TB when using the CIFS protocol, since CIFS is not designed as a high-I/O storage protocol. Any storage device must meet the minimum input/output requirements. For these requirements, and for additional guidance for sizing your hardware, software, memory, storage, and network requirements, see the *Rapid Recovery Sizing Guide*.

Table 6: Local disk settings

Text Box	Description
Location	<p>Specify the path for a local drive on which to store repository data. For example, type <code>E:\Repository\Data</code>.</p> <p>Use only alphanumeric characters, hyphen, or period, with no spaces or special characters.</p>
Metadata path	<p>Specify the path for a local drive on which to store repository metadata. It can be the same local volume (for example, <code>E:\Repository\Metadata</code>), or for more efficiency, you can specify a separate local volume (for example, <code>F:\Repository\Metadata</code>).</p> <p>When specifying the path, use only alphanumeric characters, the hyphen, and the period (but only to separate host names and domains). You can use the backslash character only to define levels in the path. Do not use spaces. No other symbols or punctuation characters are permitted.</p> <p>i NOTE: If you do <i>not</i> specify a path for metadata, then repository data and metadata are both stored in the Location path. Ten percent of the volume space allocated for repository data is then reserved specifically for metadata.</p>

- To specify a storage location on a CIFS network shared location, enter the following information:

Table 7: CIFS share credentials

Text Box	Description
Location	<p>Specify the path for a CIFS network share location on which to store repository data and metadata.</p> <p>The path must begin with \\. When specifying the path, use only alphanumeric characters, the hyphen, and the period (but only to separate host names and domains). The letters a to z are case-insensitive. Do not use spaces. No other symbols or punctuation characters are permitted.</p> <p>If this location is at the root, define a dedicated folder name (for example: \\Server12\Repository).</p>
User name	<p>Enter the user name associated with an administrative account with access to the network share location, or to use a set of credentials saved to Credentials Vault, use the drop-down list and select a user name.</p>
Password	<p>Specify a password for accessing the network share location.</p>

8. Click **Next**.
The *Space Allocation* page appears.

9. On the *Space Allocation* page, perform the following tasks:

- a. Specify the amount of storage space to leave unused on the repository volume. The remainder of space on the volume is dedicated to repository file storage.

i **NOTE:** Industry best practice suggests that drive speed and performance is enhanced when 10 to 20 percent of a storage volume is left unused. While Quest makes no recommendations in this area, the default setting for reserved free space is 20 percent. You can change this amount explicitly by moving the **Percentage of available space** slider, or implicitly by changing the amount in GB of space in the **Data** setting.

- To increase the amount of free space to leave on the repository volume, move the slider from its current position to the left.
- To increase the amount of storage space on the volume to be used for repository storage, move the slider from its current position to the right.

i **NOTE:** Quest does not recommend filling the entire volume with repository files. Consider leaving at least 10 percent of the volume free.

If the value shown on the **Percentage of available space** slider is set as intended, then data and metadata values are informational only. If you change the amount in GB of space in the **Data** setting, the slider and metadata values change accordingly.


- b. Optionally, you can specify in GB the amount of space on the volume to be used to store data.

Before you change this setting, note the following:

- If you change the amount in GB of space in the **Data** setting, the slider and metadata values change accordingly.
- The metadata value always represents between 8 and 10 percent of the total space reserved to back up data in the repository.
- If you designated two separate volumes for storing data and metadata on the *Storage Location* page of the wizard, then the metadata value shows in GB the amount of space reserved for metadata on the specified volume. In this case, if you add the amount of free space reserved as shown in the slider with the amount of space in the Data text field, the sum equals the total amount of storage space available on the volume.
- If storing repository data and metadata on the same storage volume, the 10 percent allocation for metadata is reserved on the single repository volume. The amount shown in the Data text field represents the remaining 90 percent of the storage space allocated for the repository. In this case, if you add the amount of free space reserved as shown in the slider, the amount of space in the Data text field, and the amount in the Metadata field, the sum equals the total amount of storage space available on the volume

- c. Optionally, on the *Space Allocation* page, if you want to enter advanced options, select **Show advanced options** and enter the details for the advanced settings as described in the following table.

Table 8: Space allocation details

Text Box	Description
Bytes per sector	<p>Specify the number of bytes you want each sector to include. The default value is 512.</p> <div> NOTE: Quest recommends setting the bytes per sector to match the physical sector size of the storage location in which the repository resides. For example, if the disk on the intended storage location has a 4096 byte sector size, change the bytes per sector setting to 4096. If using multiple storage locations with different sector sizes, Quest recommends retaining the default setting of 512 bytes per sector.</div>
Average bytes per record	<p>Specify the average number of bytes per record. The default value is 8192.</p>
Write caching policy	<p>The write caching policy controls how the Windows Cache Manager is used in the repository and helps to tune the repository for optimal performance on different configurations. As of release 6.4, the default for this setting is Off. Set the value to one of the following:</p> <ul style="list-style-type: none">• On. If set to On, Windows controls the caching.• Off. If set to Off, which is the default, Rapid Recovery controls the caching.• Sync. If set to Sync, Windows controls the caching as well as the synchronous input/output.

- d. When you have completed your repository configuration, click **Finish**. The *Create Repository Wizard* closes, and Rapid Recovery applies the settings to your Core. If Toast alerts are enabled, you see messages indicating that repository creation has started, and the repository is mounted. Alternatively, you can monitor the progress of the repository creation by viewing alerts on the *Events* page.
- After a brief pause, on the *Repositories* page, in the DVM Repositories summary table, your new repository is listed.

Changing DVM repository settings

This procedure assumes that your Core is already using at least one DVM repository.

In the settings for a DVM repository, you can change such settings as number of concurrent operations, and enabling or disabling deduplication or compression.

Complete the following task to change the available settings for a DVM repository.

To change DVM repository settings




1. From the Rapid Recovery Core Console, in the icon bar, click  (More), and then select  **Repositories**.
2. On the *Repositories* page, from the row representing the repository you want to update, click  (More options) and then select **Settings**.
The *Repository Settings* dialog displays.
3. On the *Repository Settings* dialog, you can change the settings described in the following table.

Table 9: DVM settings

Option	Description
Maximum concurrent operations	The number of jobs that the repository can perform at one time. The default is 64.
Concurrent operations	Enter the number of operations you want the repository to perform at the same time. The default number is 64.
Threshold	Enter the amount of data you may want to delete at one time in the event that the repository becomes full. This threshold lets you delete incremental recovery point chains up to this size.
Description	Can contain and display notes or a description that you want to associate with this repository.
Enable deduplication	When this option is selected, Rapid Recovery Core deduplicates data so that only unique blocks are saved to the repository. This setting is enabled by default. Clear this option and save to disable deduplication.
Enable compression	When this option is selected, Rapid Recovery Core compresses data to reduce space used. This setting is enabled by default. Clear this option and save to disable compression.

4. Click **Save**.
The changes are applied to the repository.

Managing an Azure repository

Managing an Azure repository involves the following operations:

1. **Creating an Azure repository.** Before creating a repository, consider the appropriate technology type.
 - For information about the different repository types, see [Understanding repositories](#)
 - For information about creating an Azure repository, see [Creating a DVM repository](#).
2. **Checking a repository.** For more information about checking an Azure repository, see [Checking a repository](#).
3. **Modifying repository details.** For more information about viewing repository details or modifying the details that display, see [Viewing or modifying repository details](#).

4. **Changing repository settings.** For more information about changing the settings for an Azure repository, see [Changing Azure repository settings](#).
5. **Deleting a repository.** For more information about deleting a repository, see [Deleting a repository](#).

Creating an Azure repository

To perform this procedure, you must have the appropriate credentials for the Azure account.

To create an Azure repository

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **...** (More), and then select **Repositories**.
The *Repositories* page displays.
3. At the top of the page, under the *Repositories* page title, click **+ Create**.
The *Create Repository Wizard* appears.
4. On the *Repository type* page, select **Azure**.
5. On the *Configuration* page, enter the information as described in the following table.

Table 10: Repository configuration settings

Text Box	Description
Name	Enter the display name of the repository. By default, this text box consists of the word Repository and a number, which corresponds to the number of repositories for this Core. For example, if this is the first repository, the default name is Repository 1. Change the name as needed. Repository names must contain between 1 and 40 alphanumeric characters, including spaces. Do not use prohibited characters or prohibited phrases .
Comments	Optionally, enter a descriptive note about this repository. You can enter up to 254 characters. Examples might include a message such as Azure Replication Repository or This repository contains protected SQL Server data only . This information can later be viewed and edited by accessing the Settings for this repository.
Show advanced options	Optionally, select this option and reveal the Concurrent operations setting. The default value is 64.

6. Click **Next**.
The *Storage Location* page appears.

7. On the *Storage Location* page, enter the details described in the following table.

Table 11: Azure repository settings

Text Box	Description
Metadata location	<p>Specify the path for a local drive on which to store repository metadata. It can be the same local volume (for example, E:\Repository\Metadata), or for more efficiency, you can specify a separate local volume (for example, F:\Repository\Metadata).</p> <p>When specifying the path, use only alphanumeric characters, the hyphen, and the period (but only to separate host names and domains). You can use the backslash character only to define levels in the path. Do not use spaces. No other symbols or punctuation characters are permitted.</p> <p>i NOTE: If you do <i>not</i> specify a path for metadata, then repository data and metadata are both stored in the Location path. Ten percent of the volume space allocated for repository data is then reserved specifically for metadata.</p>
Dedupe cache location	Select whether to store dedupe cache in an Azure Blob or your Local file system .
Index files location	Select whether to store index files in an Azure Blob or your Local file system .
Data storage account	Optionally, if your Azure account is pre-configured on this Rapid Recovery Core, to autofill the following details, select the storage account from the drop-down list.
Account name	Enter the name of your Azure data storage account.
Access key	Enter an access key for your Azure account.
Account type	From the drop-down list, select the type of Azure account you want to add.
Use https protocol	Optionally, clear this option if you do not want to use the https secure protocol.

8. Click **Next**.
The *Space Allocation* page appears.
9. On the *Space Allocation* page, perform the following tasks:

Table 12: Azure repository settings

Text Box	Description
Azure repository size	<p>Specify the maximum amount of space you want to reserve for the Azure repository.</p> <p>i NOTE: Rapid Recovery uses only as much space in the repository as it needs for your replicated data. The size that you enter tells Rapid Recovery not to use more space than the specified amount.</p>
Deduplication cache size	Specify the maximum amount of space you want to reserve for the dedupe cache.
Metadata cache size	Specify the maximum amount of space you want to reserve for the metadata cache.

10. Click **Finish**.

Changing Azure repository settings

This procedure assumes that your Core is already replicating to an Azure repository.

In the settings for an Azure repository, you can change such settings as number of concurrent operations, and enabling or disabling deduplication or compression.

Complete the following task to change the available settings for an Azure repository.

To change Azure repository settings

1. From the Rapid Recovery Core Console, in the icon bar, click **...** (More), and then select **Repositories**.
2. On the *Repositories* page, from the row representing the repository you want to update, click **...** (More options) and then select **Settings**.
The *Repository Settings* dialog displays.
3. On the *Repository Settings* dialog, you can change the settings described in the following table.

Table 13: DVM settings

Option	Description
Maximum concurrent operations	The number of jobs that the repository can perform at one time. The default is 64.
Description	Can contain and display notes or a description that you want to associate with this repository.
Azure repository size	The maximum amount of space you want to reserve for the Azure repository. NOTE: Rapid Recovery uses only as much space in the repository as it needs for your replicated data. The size that you enter tells Rapid Recovery not to use more space than the specified amount.
Deduplication cache size	The maximum amount of space you want to reserve for the dedupe cache.
Metadata cache size	The maximum amount of space you want to reserve for the metadata cache.
Enable deduplication	When this option is selected, Rapid Recovery Core deduplicates data so that only unique blocks are saved to the repository. This setting is enabled by default. Clear this option and save to disable deduplication.
Enable compression	When this option is selected, Rapid Recovery Core compresses data to reduce space used. This setting is enabled by default. Clear this option and save to disable compression.

4. Click **Save**.
The changes are applied to the repository.

About repository optimization

When you protect a machine with Rapid Recovery, the data you capture in each snapshot is deduplicated and stored in a DVM repository and then optionally replicated to a DVM repository on another Core or to an Azure repository. This deduplication occurs incrementally, as snapshots are saved to the initial repository. One occurrence of each string of information is saved to the repository. When an information string is duplicated, a reference to the original string in the deduplication cache is used, saving storage space in the repository.

If the deduplication cache is filled, only snapshot data that is already referenced in the cache is deduplicated. As deduplication occurs, the cache continues to update with new unique values, overwriting the oldest values in the cache. This results in less than optimal deduplication.

For more information about deduplication, see [Understanding deduplication cache and storage locations](#).

You can choose to increase your duplication cache before it is full, which ensures continued optimal deduplication of your data in that repository. For more information, see [Configuring DVM deduplication cache settings](#).

You can also increase your deduplication cache after it is full. If you want to reclaim space in the repository after increasing your cache, you can optimize the repository. This action forces a comparison of the data in your snapshots to the information in the deduplication cache. If any repeated strings are found in the repository, that data is replaced with references to the data, which saves storage space in the repository. This is sometimes referred to as off-line deduplication, since this deduplication process occurs upon your request, instead of incrementally as snapshot data is transferred.

The optimization process is processor-intensive. The amount of time it takes to run this job depends on several factors. These factors include the size of your repository; the amount of data in your repository; available network bandwidth; and existing load on the input and output of your system. The more data in your repository, the longer this job runs.

The following actions are superseded or canceled when the Repository Optimization Job is occurring.

- Delete All Recovery Points Job
- Delete Recovery Points Chain Job
- Maintain Repository Job
- Delete Recovery Points Job Base
- Optimize Repository Job

For steps on optimizing an existing repository, see [Optimizing a repository](#).




You can interrupt the Optimize Repository job for a limited time if required. For more information, see [Interrupting or resuming repository optimization](#).

Optimizing a repository

You can perform on-demand deduplication of data saved to an existing repository. This is accomplished by launching the Repository Optimization Job.

i | **NOTE:** Quest recommends performing the Optimize Repository job only after increasing your deduplication cache size. This action lets you reclaim repository space and more effectively use the deduplication cache.

Complete the steps in this procedure to optimize a repository.





1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More), and then select  **Repositories**.
The *Repositories* page appears.
3. From the row representing the repository you want to optimize, click  (More options) and then select **Optimize**.
A warning prompt appears asking you to confirm the optimization.
4. To confirm the optimization, click **Yes**.
The optimization job takes precedence over most other jobs. If necessary, you can interrupt an optimization job in progress. For more information on interrupting or resuming this job, see [Interrupting or resuming repository optimization](#).

Interrupting or resuming repository optimization

When you initiate the Optimize Repository Job, the selected repository is deduplicated. This deduplication optimization is a processor-intensive job intended to save space in the repository. For more information, see [About repository optimization](#).

Once this job has been initiated, you can interrupt the job using the following procedure. This pauses deduplication. If you have already interrupted a optimization, you can resume the process using this procedure.

Complete the steps in this procedure to interrupt or resume a repository optimization job.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More), and then select  **Repositories**.
The *Repositories* page appears.
3. If you want to interrupt an optimization job, do the following:
 - a. In the repositories summary table, from the row representing the appropriate repository, click  (More options) and then select **Interrupt Optimization Job**.
A warning prompt appears asking you to confirm the interruption.
 - b. To confirm the interruption, click **Yes**.
4. If you want to resume an interrupted optimization job, do the following:
 - a. In the repositories summary table, from the row representing the appropriate repository, click  (More options) and then select **Optimize**.
A warning prompt appears asking you to confirm the optimization.
 - b. In the dialog box, select the option **Continue job from the interrupted point**, and then click **Yes**.

The dialog box closes, and the repository optimization job resumes from the point where it was last interrupted.

Connecting to an existing repository

To perform this procedure, you must have the appropriate credentials for the original Core, and you must have the local or network path, IP address or server name.

If connecting to a DVM repository that is currently owned by another functioning Core, you must first prepare for the impending transfer of ownership. The following prerequisites apply:

- If the original Core containing the repository is functioning, temporarily stop protection of machines on that Core.
- Then, stop the Core service of the original Core while you transfer ownership.
- After connecting the repository to the new Core as described in this procedure, return to the original Core. Ensure that all machines you continue to protect have a repository associated with them.
- Then restart the Core services on the original Core.

For more information about pausing protection, see [Pausing and resuming protection](#). For more information about shutting down and restarting the Core service, see [Restarting or shutting down the Core service](#).

From your Rapid Recovery Core Console, you can connect to an existing repository that is currently managed in a different Core. The repository you connect to must be accessible on a shared network location, or on a storage device accessible to the second Core.

This process is useful if your original Core is down and you wish to stand up a replacement Core. If you later wish to change ownership again from the second Core to a third Core (or to the original), you can do so. The same rules apply.

CAUTION: Connecting to a repository from another Core changes ownership of the repository. After you connect to the repository, the information is then accessible only to the second Core. The original Core losing the repository must not be in use. For example, the machine must be turned off, not accessible to the network, or the Core services must be stopped.

Complete the following procedure to connect to an existing repository.




1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More), and then select  **Repositories**. The *Repositories* page displays.
3. To select a repository type, click the drop-down  menu next to **Connect to Existing**.
4. In the *Connect Existing DVM Repository* dialog box, enter the following information for the repository you want to open, and then click **Connect**.

Table 14: Open Existing DVM Repository options

Text Box	Description
Path	The path for the repository (for example, D:\work\machine for a local path, or \\10.10.99.155\repositories by IP address, or \\servername\sharename for a network path).
User name	If the repository has a network path, enter the user name for logging in to the network share.
Password	If the repository has a network path, enter the password for logging in to the network share.

The dialog box closes, and the selected repository is added to your current Core.

5. On the *Cores* page of the wizard, click to select the appropriate Core, and then click **Next**.



6. On the *Details* page of the wizard, to ensure that the selected Core is the one to which you want to connect, review the displayed details, and then click **Finish**.

If you see an error indicating that the selected repository is in use, log into that Core and prepare it for transferring ownership of its repository to this Core. Pause any existing protection; pause any existing replication; wait for queued jobs to complete, or cancel them. Shut down the Core services or gracefully power down the Core server, and then repeat this procedure. For more information on pausing protection, or resuming protection that is paused, see [Pausing and resuming protection](#). For more information on pausing replication, or resuming replication that is paused, see [Pausing and resuming replication](#). For more information about shutting down and restarting the Core service, see [Restarting or shutting down the Core service](#).



Viewing or modifying repository details

To view repository details, your Core must first contain a repository. For information about creating a DVM repository, see [Creating a DVM repository](#). For more information about creating an Azure replication repository, see [Creating an Azure repository](#).

In the pane for each repository type, each  repository added to the Core is displayed.





- If you click  (Expand), child rows display the data and metadata locations for the repository.
- You can also  (Contract) the view.

Repository details differ based on the repository technology type. To view or modify repository details, use the following procedure.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More), and then select  **Repositories**.
The *Repositories* page is displayed.

3. At the top of the page, under the *Repositories* page title, the following general actions are available:

Table 15: Repository actions

Option	Description
 Create	This option opens the <i>Create Repository Wizard</i> , a workflow that helps you define all requirements to create a new DVM repository. For more information, see Creating a DVM repository .
 Connect to Existing	For DVM repositories only, this option opens the <i>Connect Existing DVM Repository</i> dialog box. When you provide the local or CIFS share path and connection credentials, this lets you open a repository from another Core.  NOTE: This process changes ownership of the repository to your Core. For more information, see Connecting to an existing repository .
 Refresh	Refreshes the list of repositories shown on the page.





4. From the  (More options) drop-down menu for any  repository, you can perform the following actions:

Table 16: Repository options

Option	Description
Expand repository	For DVM repositories only, expand the existing repository by adding a storage location.  NOTE: Before extending a DVM repository volume, Quest recommends pausing protection and pausing all jobs (or waiting for them to complete). Then extend the volume. Finally, resume protection and any paused jobs.
Check	Perform a repository check.
Settings	View or modify repository settings. These settings include: <ul style="list-style-type: none">• Viewing the repository name• Viewing or changing the maximum concurrent operations• Viewing or changing a description for the repository• Enabling or disabling deduplication• Enabling or disabling compression for data stored in the repository
Optimize	Perform a repository optimization job. For more information, see About repository optimization .
Delete	Delete a repository.  CAUTION: This option completely deletes the selected DVM repository and all data it contains.

Checking a repository

Rapid Recovery lets you perform a diagnostic check of a repository volume when errors occur on the Core. Errors could be the result of the Core being improperly shut down, a repository failing to mount or unmount, a hardware failure, or other environmental, lower IP stack factors that can be exposed in Rapid Recovery functionality.

For all repository technology types, the check performs the following tasks:

- Check repository
- Mounting repository
- Loading the recovery points from repository

The check also performs the "Recalculate deduplication cache for repository" task.

NOTE: This procedure should only be performed for diagnostic purposes. For example, run this check in the event of hardware failure, improper shutdown of the Core, or failure when importing a repository.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **More** (More), and then select **Repositories**.
The *Repositories* page appears.
3. To check a repository, in any row of the summary table for any repository technology type, click **More** (More options), and then select **Check**.
The *Check Repository* dialog box appears.
4. In the *Check Repository* dialog box, confirm that you understand that all active tasks associated with this repository will be canceled and that you want to proceed.
Active jobs are canceled and the Checking Repository job starts.

Deleting a repository

Complete the steps in this procedure to delete a repository.


1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **More** (More), and then select **Repositories**.
The *Repositories* page is displayed.
3. In the appropriate repositories summary table, from the row representing the repository you want to delete, click **More** (More options) to expand the drop-down menu, and then select **Delete**.
A warning message appears to confirm deletion.

CAUTION: When a repository is deleted, the data contained in the repository is discarded and cannot be recovered.

4. Click **Yes** to confirm the deletion of the repository.

Core settings

This section describes how to manage and change the settings for your Rapid Recovery Core, and describes key function buttons and tools available from the Core console.

The Rapid Recovery Core has adjustable settings that are configured by default for optimum performance for most users. These settings affect display information in the Core Console or performance of the Rapid Recovery Core. From the icon bar, click  (Settings) to access Core settings.

A set of key functions are shown as buttons on the top of the page horizontally. To access one of these functions, click the corresponding button. For more information about these buttons, see [Core settings key functions](#).

Below the key function buttons are the configurable Core settings. To see all configuration options for any of setting, click a shortcut link on the left side of the Settings pane, or scroll down the right side of the page. For more information about the full set of Core settings, see [Rapid Recovery Core settings](#).


You can also access Core tools such as viewing a summary of system information, or downloading Core log files. For more information, see [Core-level tools](#).

Topics include:

[Core settings key functions](#)




[Rapid Recovery Core settings](#)


Core settings key functions

A set of key functions are shown as buttons arranged horizontally at the top of the  (Settings) page. To access one of these functions, click the corresponding button.

The function buttons accessible on the Settings page are described in the following table.

Table 17: Core settings key functions

Icon	Key Function Button	Description
	Back Up Settings	Backs up Core configuration settings to an XML file you name. Specify a fully qualified path local to the Core server.
	Restore Settings	If you have a backup XML file of your Core settings, this option lets you specify the name and local path of the file, from which you restore Core settings. Use this function to restore Core settings, or to migrate from another Core. Optionally, you can also restore repositories.
	Restart Core Service	This option gracefully shuts down and then starts the Core service.

Icon	Key Function Button	Description
	Shut Down Core Service	This option gracefully shuts down the Core service.

For more information about backing up and restoring Core settings, see [Backing up and restoring Core settings](#).

For more information about shutting down and restarting the Core service, see [Restarting or shutting down the Core service](#).

Backing up and restoring Core settings




You can back up Core setting information to a file, and later restore these settings if you have problems with the Core machine or if you want to migrate those settings to a different machine. Information that gets backed up includes repository metadata (such as the repository name, data path, and metadata path); machines protected in the Core; replication relationships (targets and sources); which machines are configured for virtual standby; and information about encryption keys.

This process restores the configuration settings only, not the data. Security information (such as authentication credentials) is not stored in the configuration file. There is no security risk to saving a Core configuration file.




NOTE: You must first back up Core setting information before you can use this process to restore Core settings.

Use this procedure to back up and restore Core settings.

1. Navigate to the Rapid Recovery Core.
2. On the icon bar, click  (Settings).
The Settings page appears. At the top of the Settings pane, above the categories of settings, you see options to  **Back Up Settings** or  **Restore Settings**.
3. If you want to back up Core settings, proceed to step 4. If you want to restore Core settings, proceed to step 6.
4. To back up the current settings in an XML file, from the top of the Settings page, click **Back Up Settings**.
The Back Up Core Configuration dialog box appears.
5. In the Local path text box, type a directory path accessible locally to the Core machine where you want to store Core settings as an XML file, and then click **Back Up**.
For example, type `C:\Users\Your_User_Name\Documents\RRCoreSettings` and then click **Back Up**.
A file named AppRecoveryCoreConfigurationBackup.xml is saved to the local destination you specified.

6. To restore Core settings from a backup XML file saved previously using this method, perform the following steps.

i | **NOTE:** When you restore the Core configuration settings, the Rapid Recovery Core service restarts.

- a. From the top of the Settings page, click  **Restore Settings**.
The Restore Core Configuration dialog box appears.
- b. In the **local path** text box, enter the local path of the location where you stored the Core configuration settings.
For example, type `C:\Users\Your_User_Name\Documents\RRCoreSettings`.
- c. If you do not want to restore repository information, proceed to step g.
- d. Optionally, if you want to restore repository information as configured in the backup file, select **Restore Repositories** and then click **Restore**.
The Restore Repositories dialog box appears.
If you choose to restore repository information from the backed-up configuration data, then any repositories configured when the Core settings were saved appear for verification. By default, each existing repository is selected.
- e. Verify the repository information you want to restore. If multiple repositories appear in the lists for verification, and you only wish to restore information for some of them, then clear the selection for each repository you do not want.
- f. When you are satisfied with the selection of repositories you want to restore, click **Save**.
The Restore Repositories dialog box closes.
- g. In the Restore Repositories dialog box, click **Restore**.
The Restore Repositories dialog box closes, and the restore process begins. An alert appears indicating that the repository service configuration has changed.
- h. If any configuration settings could not be restored you will see an error message. Review the details of the error to see if any action is required on your part. For more information, see [Viewing events using tasks, alerts, and journal pages](#). To continue, click **Close** to clear the error dialog box.

- i. After restoring the configuration, verify the following:
- Unlock all encryption keys. For more information, see [Unlocking an encryption key](#).
 - If virtual standby is configured to continually update a VM to a network destination, you must specify the network credentials in the virtual standby settings before a successful synchronization. For more information, see [VM export](#).
 - If scheduled archive is configured to archive to a cloud account, you must specify credentials so the Core can connect to the cloud account. For more information on linking the Core with a cloud account, see [Adding a cloud account](#).
 - If replication is set up and you want to restore to a target Core, verify the target Core settings (particularly the host) on the source Core. For more information, if managing your own Core, see [Replicating to a self-managed target Core](#). If replicating to a Core managed by a third party, see [Replicating to a third-party target Core](#).
 - If the SQL attachability check is configured, and if the SQL Server instance performing the check is on the Core machine, then specify the SQL credentials in Attachability settings. For more information, see [Managing Core SQL attachability settings](#).
- Verify that the Replay Engine configuration was restored, and update the settings if they were not to ensure effective communication. For more information, see [Configuring Replay engine settings](#).




Restarting or shutting down the Core service

Invariably, a machine on which Rapid Recovery Core is running shuts down or must be rebooted. In release 6.2, Rapid Recovery Core is enhanced to improve its ability to gracefully shut down and restart Core services.

The Core UI now offers UI features to either restart or shut down the Core service with one click. This feature is useful when planned maintenance of the Core server (including rebooting or restarting) is required. Users receive notification when corresponding services have finished shutting down. These features are accessed from the top of the Core *Settings* page.

i | **NOTE:** Another useful feature supporting graceful shutdown is the ability to suspend the Core from scheduling future tasks. For more information, see the topic [Suspending or resuming scheduled tasks](#).

Use this procedure to restart or shut down the Core service.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings).
The Settings page is displayed. At the top of the Settings pane, above the categories of settings, you see options to **Restart Core Service** or **Shut Down Core Service**.
3. If you want to restart the Core service, at the top of the Settings pane, click  **Restart Core Service**.
A list of each task that must be completed before restarting is displayed, including a display of task progress.
4. Optionally, you can do one of the following:
 - a. To force restart (a less graceful process in which tasks are explicitly stopped before completion), click **Force Restart**.
 - b. To cancel restarting of the Core service, click **Cancel**.
5. If you want to shut down the Core service without immediately restarting it, at the top of the Settings pane, click  **Shut Down Core Service**.

6. Optionally, you can do one of the following:
 - a. To force shutdown (a less graceful process in which tasks are explicitly stopped before completion), click **Force Shutdown**.
 - b. To cancel restarting of the Core service, click **Cancel**.

Rapid Recovery Core settings

The Settings pane contains a navigation column on the left side, listing each Core setting. Click any link in this list, or scroll down on the right side of the page to see all configuration options for each Core setting.

When you click on a setting you want to change, that setting becomes an editable control. Do one of the following:

- When the control is a drop-down menu, click the downward arrow to list the options, and select the desired option from the menu.
- When the control is a text field, enter a value.
- When the option displays **Yes** or **No**, click the value, which is replaced by a check box. For a setting of **Yes**, select the check box. For a value of **No**, clear the check box.
- When the option displays a time value (for example, showing hours, minutes, and seconds), you can click on each component and type a new value or use the up and down arrows to select new values.

For each setting, when satisfied with your changes, click ✓ to confirm and save the change and exit edit mode, or click ✕ to exit edit mode without saving.

The Rapid Recovery Core settings that you can configure are described in the following table. Each setting has a link to a relevant topic with more information.

Table 18: Rapid Recovery Core configurable settings

Configuration Setting	Description
General	<p>General settings include configuration options that apply generally to the Rapid Recovery Core, including display options and ports for the web server and for the Rapid Recovery service.</p> <p>For more information about the general settings for Rapid Recovery Core, including how to configure these settings, see Configuring Core general settings.</p>
Updates	<p>Update settings control aspects of the automatic update feature, which checks for updated versions of Rapid Recovery software.</p> <p>For more information about settings for updating the Rapid Recovery Core, including how to configure these settings, see Configuring update settings.</p>
Nightly Jobs	<p>Nightly jobs settings are automated tasks which the Core performs daily. You can configure the time the jobs begin and which jobs are performed. Quest recommends scheduling the jobs outside of normal business hours to reduce load on the system when demand for resources is high.</p> <p>For more information, see Understanding nightly jobs, Configuring nightly jobs for the Core, and Customizing nightly jobs for a protected machine.</p>

Configuration Setting	Description
Transfer Queue	<p>Transfer queue settings control the number of times transfer operations are attempted if jobs fail due to unavailability of resources. You can establish the maximum number of concurrent transfers and the maximum number of retries for transferring data.</p> <p>For more information about transfer queue settings, see Modifying transfer queue settings.</p>
Client Timeout	<p>Client timeout settings determine the length of time before that specific connection requests or read and write operations should be attempted before timing out.</p> <p>For more information about client timeout settings, see Adjusting client timeout settings.</p>
DVM Deduplication Cache	<p>Deduplication ensures that unique blocks of information are stored only one time in your repository, creating references to repeated data blocks. The references are stored in a deduplication cache. If encryption keys are used, then deduplication occurs within each encryption domain.</p> <p>DVM deduplication cache settings let you configure the size and specify the locations for the primary and secondary cache, as well as the location for the metadata cache.</p> <p>For more information about deduplication cache, see Understanding deduplication cache and storage locations. For information about adjusting the settings, see Configuring DVM deduplication cache settings.</p>
Replay Engine	<p>Replay engine settings control information regarding the communication channel for the Replay engine, such as IP addresses and timeout settings, to help adjust the performance specific to your network needs.</p> <p>For more information about engine settings for Rapid Recovery, see Configuring Replay engine settings.</p>
Deploy	<p>Deploy settings let you set options for deploying the Rapid Recovery Agent software from your Core to the machines you want to protect.</p> <p>For more information about configuring deployment settings, see Configuring deployment settings.</p>
Database Connection	<p>Rapid Recovery stores transactional information in a MongoDB service database that is installed locally by default on the Core machine. You can configure these settings to change how long information is retained in the database, or to change the connection pool size to allow for more or fewer concurrent connections.</p> <p>For more information about establishing or modifying database connection settings for the service database, see Configuring database connection settings.</p>
Local Database Settings	<p>Rapid Recovery displays information about Core tasks, events, and alerts on the <i>Events</i> page. Rapid Recovery stores this transactional information in a MongoDB service database that is installed locally on the same machine as the Rapid Recovery Core.</p> <p>You can configure credential information (user name and password) for the local Mongo service database using the Local database settings. For more information about adjusting local database settings, see Modifying local database connection settings.</p>
SMTP Server	<p>Configure Simple Mail Transfer Protocol (SMTP) server settings for the Core to send Core event information by email.</p> <p>For more information about configuring an SMTP email server, see Configuring an email server.</p>


Configuration Setting	Description
	<p>i NOTE: To send event information by email, you must also configure notification group settings. For more information about specifying events to receive email alerts, see Configuring notification groups.</p>
Cloud Accounts	<p>The Cloud Accounts settings let you specify configuration settings for supported cloud accounts. These settings do not create cloud accounts. Instead, they associate existing external cloud storage or cloud service provider accounts with your Rapid Recovery Core to facilitate actions such as archiving Rapid Recovery information. For information about setting timeout settings for cloud accounts, see Configuring cloud account connection settings.</p> <p>For more information about managing cloud accounts in the Rapid Recovery Core Console, see Cloud accounts.</p>
Reports	<p>Report settings include configuration parameters that allows you to select the font used when a report is generated from the Rapid Recovery Core. You can also set the paper size and page orientation for reports.</p> <p>For more information about changing report settings, see Managing report settings.</p>
Attachability	<p>Attachability settings let you specify whether to perform SQL attachability checks on the protected machine, or whether to use the SQL Server instance on the Core. If specifying SQL on the Core, you must provide credential information.</p> <p>For more information about managing SQL attachability settings for the Core, see Managing Core SQL attachability settings.</p>
Jobs	<p>Core jobs are automatically created whenever you initiate operations such as replication. You can specify settings for each job using the Jobs settings for the Core.</p> <p>You can configure the number of jobs to run at one time. In case network or other communication errors prevent any job from succeeding the first time, you can set how many times a job should be attempted using the Try Count setting.</p> <p>For more information about Core jobs, which jobs are available, and how to configure them, see Core job settings.</p>
Licensing	<p>From the Core console, Rapid Recovery lets you change the license associated with your Core, limit the number of daily snapshots, view license pool information, and contact the license server.</p> <p>For more information about managing licenses from the Core, see "Understanding Rapid Recovery licenses" in the <i>Rapid Recovery Installation and Upgrade Guide</i>.</p> <p>For more information about managing licenses, see the <i>Rapid Recovery License Portal User Guide</i>.</p> <p>i NOTE: The Rapid Recovery License Portal has a different release cycle than Rapid Recovery software. For the latest product documentation, see the Quest Technical Documentation website.</p>
SNMP Configuration	<p>Simple Network Management Protocol (SNMP) is a protocol for managing devices on an IP network. You can configure the Rapid Recovery Core as an SNMP agent. The Core then can report information such as alerts, repository status, and protected machines.</p> <p>For more information about using SNMP with Rapid Recovery, see Understanding SNMP settings.</p>

Configuration Setting	Description
vSphere	vSphere Core settings apply only for users of the agentless protection of virtual machines. If using a vSphere host, these settings include connection settings that apply to the VMs. For more information about vSphere settings for VMware or ESXi agentless protection, see Configuring vSphere settings .
VMware Proxy	A VMware proxy service installed with the Core lets users set service timeouts associated with VMware disk storage. For more information about these settings, see Managing VMware proxy settings .
QorePortal	If managing two or more Cores, you can integrate your Core server with the QorePortal. This feature, particularly useful for managed service providers, lets you manage multiple Cores; access a dashboard where you can monitor tasks and events, view repository status, and check system health; generate reports; and perform a growing list of other functions from a single web-based user interface. To enable or disable access to the portal, use this setting.
vFoglight	Quest Foglight for Virtualization (vFoglight) helps administrators monitor, analyze, and optimize hypervisors across VMware, Hyper-V, and OpenStack environments. For customers managing VMs using vFoglight and protecting them on a Rapid Recovery Console, this Core setting lets you integrate navigation for the two products. After successfully entering vFoglight settings, by clicking the vFoglight URL on the Summary page for VM in the Core Console, users navigate to the corresponding page for that virtual machine in vFoglight. For more information about the vFoglight Core settings, see Configuring vFoglight settings . For more information about vFoglight, see the vFoglight product page on the Quest website.
Synthetic incremental	When enabled, this feature compares the data on a protected machine with the data that is in the previous backup stored in the repository and only sends the difference to the Core server, which significantly reduces the amount of repository space needed to store the backup and thereby reduces the amount of time needed to re-replicate or re-export that recovery point. For more information, see https://support.quest.com/rapid-recovery/kb/331728/rapid-recovery-synthetic-incremental-feature-overview .

You can also access Core tools such as viewing a summary of system information, or downloading Core log files. For more information, see [Core-level tools](#).






Configuring Core general settings

General settings for the Rapid Recovery Core include the Core ID, display name, the web server port, service port, locale (the Core console display language), and the display color theme.



1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **General**.
 - Scroll down on the right side of the *Settings* page until you can see the **General** heading.
3. Click on the general setting you want to change.
The setting you selected becomes editable, as a text field or a drop-down menu.

4. Enter the configuration information as described in the following table.

Table 19: General Settings information

Text Box	Description
Core ID	Each Core has a unique Core ID. This ID is used, for example, to integrate your Core with the QorePortal to provide reporting or management of two or more Cores. The Core ID is now listed in General settings.  NOTE: This field is not configurable.
Display name	Enter a new display name for the Core. This is the name that will display in the Rapid Recovery Core Console and (if enabled) in the QorePortal. You can enter up to 64 characters.
Web server port	Enter a port number for the Web server. The default port is 8006.  NOTE: Quest recommends using the default port.
Service port	Enter a port number for the Rapid Recovery Core service. The default port is 8006.  NOTE: Quest recommends using the default port.
Locale	From the Locale drop-down list, select the language you want to display.  NOTE: If changing the languages, confirm the message indicating that the Rapid Recovery Core service must restart before the updated language can display in the Core Console. You can restart this service using the  Restart Core Service button from the top of the <i>Core Settings</i> page.
Theme	From the Theme drop-down list, select the style you want to apply to the Core Console. Three themes are available: <ul style="list-style-type: none">• Dark. This theme features a solid dark gray background throughout the interface (left navigation menu, top button bar, and primary pane). Text elements and text buttons not in focus appear in off-white, or white when in focus. Clickable links appear in a medium blue. Occasional buttons have white text over a medium blue background.• Hybrid. This theme features the familiar dark gray background for the left navigation menu and button bar at the top of the Core Console. Text elements in these areas are white. The primary pane has a white background with off-white highlights, with text elements and text buttons in black. Clickable links appear in a dark blue. Occasional buttons have white text over a blue background.• Light. This theme features a clean white background throughout the interface (left navigation menu, top button bar, and primary pane). The Quest logo and some design elements are orange. Text elements are dark gray, with titles in black. Clickable links appear medium blue on hover. Occasional buttons have white text over a medium blue background.
Agree to use of personal	To change the setting that allows the application to use personal information, from the Agree to use of personal data drop-down list, select Yes or No , as appropriate. In the resulting dialog box, select and register the appropriate license file.

Text Box	Description
data	<p>When you upgrade or install Rapid Recovery Core, you have the option to set sharing of personal information. If you agree to share information with Quest, you can use features such as automatic update and the QorePortal (which is then enabled by default).</p> <p>If you decline to share information with Quest when installing, you are prompted to register a non-phone-home license. You must have access to the non-phone-home license to save to confirm the change.</p> <p>i NOTE: Regardless of the option you selected during installation, you can change the Agree to use of personal data setting in Core General settings. Ensure you have access to the non-phone-home license, since this action prompts you to upload the non-phone home license.</p> <p>In release 6.7, when you change this setting from "Yes" to "No," the following applies:</p> <ol style="list-style-type: none"> You are prompted to upload the non-phone-home license file. After confirming the non-phone-home license, Updates settings for your Core automatically adjust to never check for or install Core updates. The QorePortal setting Enable connection to QorePortal is set to "No." <p>However, the reverse is not true. When you change this setting from "No" to "Yes," you give the Core permission to share your information, but no information is shared until you explicitly change the license to phone-home mode and update the appropriate Core settings. For example:</p> <ol style="list-style-type: none"> The non-phone-home key remains registered until you explicitly upload a standard phone-home key (which you can obtain from the license portal). To use automatic update, change "Check for new updates" from "Never" to "Daily," "Weekly," or "Monthly." Optionally, change "Install updates" to notify you or to automatically install updates. To share information with the QorePortal, set Enable connection to QorePortal to "Yes." <p>i NOTE: If you configure this first, you are prompted to upload a phone-home license.</p> <p>To understand the effect of sharing personal information, see Managing privacy, including the topic How Rapid Recovery uses personal information.</p> <p>For more information about the functions you cannot perform when using non-phone-home mode, see the topic Non-phone-home license restrictions.</p> <p>To request a license for non-phone-home mode, see Obtaining and using non-phone-home licenses.</p>

- For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Configuring update settings

Rapid Recovery includes the automatic update feature. When installing the Rapid Recovery Core, you can choose whether to automatically update the Rapid Recovery Core software when new updates are available, and how

frequently the system should check for updates.

NOTE: The automatic update feature requires a license using the standard phone-home mode. If using a software license in non-phone home mode, your Core does not have permission to communicate with the Rapid Recovery License Portal and cannot update the Core or notify you of available updates. For more information, see [Managing privacy](#).

Rapid Recovery release numbers typically include four chunks of information, separated by decimal points: the major release number, minor release number, revision, and build number. For example, the first rebranded release of Rapid Recovery was 6.0.1.609. The next release was 6.0.2.142.


The automatic update feature compares all digits in a release number. If you enable automatic update, the Core software is only updated without intervention when the major and minor release numbers are identical. For example, automatic update would occur from Core version 6.0.1.609 to 6.0.2.142 (both start with 6.0). On the same machine, the Core would not update automatically from 6.0.2.142 to 6.1.1.XXX, because the digits after the first decimal are not equal. Instead, you are notified (by a banner at the top of the Core Console) that an update to the Core software is available. This notification gives you an opportunity to review release notes, and determine if updating to the latest Core version is appropriate for your needs.

NOTE: For information on installing Rapid Recovery Core software, see the *Rapid Recovery Installation and Upgrade Guide*.

You can view and change the settings the system uses to check for updates at any time.

CAUTION: When using replication, configuring your system to install updates automatically could result in upgrading the source Core before the target Core, which may result in replication failure or the inability to set up new replication between Cores. For replication users, Quest recommends administrators apply automatic upgrades only to the target Core, and then manually upgrade the source Core, and lastly upgrade the protected machines.



Complete the steps in this procedure to configure update settings.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the Settings page, click **Updates**.
 - Scroll down on the right side of the Settings page until you can see the Updates heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.

4. Enter the configuration information as described in the following table.

Table 20: Update settings information

Text Box	Description
Check for new updates	Select how frequently Rapid Recovery checks for and installs updates. You can choose from the following options: <ul style="list-style-type: none">• Never• Daily• Weekly• Monthly If you choose automatic updates, after the selected threshold of time passes, if an update is available, it is installed after nightly jobs have completed.
Install updates	Specify the handling of available updates by choosing one of the following options: <ul style="list-style-type: none">• Never check for updates• Notify me about updates, but do not install them automatically• Automatically install updates
Status	The status indicates whether any new updates are available.
Last check	The Last check field indicates the date and time the system last checked for an update. Click Check Now to immediately verify whether a software update is available. This check occurs regardless of the frequency you have set.


5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.




Understanding nightly jobs



Nightly jobs are daily automated tasks that occur at a predetermined time outside of normal business hours. These jobs are memory-intensive, and include various integrity checks and data consolidation tasks that are best conducted when the Rapid Recovery Core is less active.



All the nightly jobs, and the scope for which they can be applied, are described in the following table. Nightly jobs can be managed at the Core level (which applies to all machines protected on the Core). Those nightly jobs which can also be applied for a specific protected machine list the scope as "Protected machine."

Table 21: Nightly jobs information

Job Name	Scope	Description
 Change	N/A	This control opens the Nightly Jobs dialog box, where you can enable, disable, or change settings for each nightly job.
Nightly jobs time	All	This setting represents the time that nightly jobs are scheduled to start running. Quest recommends configuring your Core to run nightly jobs during a time of low activity.


Job Name	Scope	Description
		The default time is 12:00 AM.
Check attachability of SQL Server databases	Protected machine	Checks the integrity of recovery points containing SQL databases. For more information, see Managing Core SQL attachability settings .
Check checksum of Exchange databases	Protected machine	Checks the integrity of recovery points containing Exchange Database (EDB) files.  NOTE: This option does not appear if you are not protecting an Exchange Server in your Core.
Check integrity of Oracle databases	Core or protected machine	Checks the integrity of Oracle databases using the DBVERIFY utility. Process: <ul style="list-style-type: none"> • Mount the latest recovery point for every protection group. • Enumerate the files and folders for each volume. • Examines the recovery points to ensure that data files are valid and data blocks are not corrupted. • Dismount the recovery point.
Check integrity of recovery points	Core or protected machine	Checks the integrity of recovery points for each protected machine.  NOTE: By default, the <code>Check integrity of recovery points</code> option is not enabled. Process: <ul style="list-style-type: none"> • Mount the latest recovery point for every protection group. • Enumerate the files and folders for each volume. • Examines the recovery points to ensure that they are valid. • Dismount the recovery point.
Clean orphaned registry keys on Hyper-V agents		For Hyper-V hosts using Rapid Recovery release 6.1.x agentless protection, this nightly job cleans orphaned keys made in the Windows registry for each attach and detach operation. The registry entries are harmless, but over time can accumulate, leading to slower performance.  NOTE: As of Rapid Recovery release 6.2, an improved approach to obtaining storage metadata for Hyper-V agentless protection precludes creating registry entries.
Consolidate VMware snapshots for protected virtual machines	Core or protected machine	This nightly job is relevant if you use native VMware APIs to protect machines without the Rapid Recovery Agent software. You should periodically consolidate VMware snapshots. Enabling this nightly job lets you perform these consolidations on a daily basis. This nightly job contains one parameter, Maximum simultaneous consolidations, which must be set to a number between 1 and 100.

Job Name	Scope	Description
Deferred delete	Core	<p>This setting lets you defer removal of recovery points from the repository until the time specified in your Core to perform nightly jobs. When enabled, then after other nightly jobs run, Core processing is dedicated to running the "Deleting records previously flagged for deletion" job. That job removes marked recovery points from the repository until they are all removed, or until four hours have passed from the nightly jobs execution time. Nightly jobs then end, and other queued jobs resume. Any remaining deletions occur in the background, concurrent with other tasks, until the next day's nightly jobs run.</p> <p> NOTE: By default, the <code>Deferred Delete</code> option is not enabled.</p> <p>Quest recommends leaving this nightly job disabled unless you are encountering transfer performance issues related to backed-up recovery point deletions.</p> <p>If you enable this option, Quest recommends reviewing your Core jobs to ensure most recovery points marked for deletion are removed from the repository within a one-week period. This approach helps to balance maximum transfer performance with maximum reclamation of repository space.</p>
Delete old events and jobs	Core	Maintains the scale of the events database by removing old events. The number of days is configurable, defaulting to 30 days.
Log truncation for Exchange	Protected machine	<p>Maintains the size of Exchange logs by truncating the exchange database transaction log to match the last recovery point.</p> <p> NOTE: This option does not appear if you are not protecting an Exchange server in your Core.</p>
Log truncation for Oracle	Protected machine	<p>Controls truncation for Oracle logs. When enabled, truncation occurs when nightly jobs run, according to the deletion policy selected.</p> <ul style="list-style-type: none"> You can select the Core Automatic deletion policy. When this policy is in effect, then when nightly jobs run, all of the locally stored Oracle archivelogs included in the last recovery point are truncated. Business logic prevents archivelogs not included in a snapshot from truncation. New archive logs that are captured in a subsequent snapshot become eligible for automatic truncation when the next nightly job is run. You can select a custom deletion policy for a specific protected Oracle server. The Keep newest policy lets you specify the duration of time before which Oracle logs are truncated, and the Keep specified number policy lets you keep a specified number of log files before truncating the older ones. When this nightly job is disabled, substantial accumulation of log files occurs. In such cases, users can also truncate log files manually, as described in the topic Manually truncating Oracle database logs.
Log truncation for SQL Server	Protected machine	Maintains the size of SQL Server logs by truncating the database transaction log to match the last recovery point.



Job Name	Scope	Description
		 NOTE: This option does not appear if you are not protecting a SQL Server in your Core.
Rollup	Core or protected machine	<p>Applies the retention policy to your backed-up data by combining or "rolling up" recovery points on the schedule dictated in the policy. You can customize the policy on the Core, which applies by default to all protected machines. By default, the rollup job is run for the whole Core; or click  [Expand] to expand the view of protected machines. You can then define the set of protected machines you want to roll up using the Core policy.</p> <p>For more information about using a retention policy on a protected machine that differs from the default policy set in the Core, see Customizing retention policy settings for a protected machine.</p>

Configuring nightly jobs for the Core

When any nightly job option is enabled on the Rapid Recovery Core, the selected job executes once daily at the time specified for all machines that are protected by the Core. Conversely, if you disable any nightly job at the Core level, the specified job no longer executes for all machines protected by the Core.

 **NOTE:** If the scope of a nightly job, as described in the topic [Understanding nightly jobs](#), includes protected machines, you can configure that nightly job to apply only for one or more specific protected machines individually. For more information about applying nightly job settings specific to a protected machine, see [Customizing nightly jobs for a protected machine](#).

Because nightly jobs are memory-intensive, Quest recommends configuring your Core to execute them during a time of low activity. The default schedule to run nightly jobs is 12:00 am. If another time is more suitable, change this setting in the Nightly Jobs Time field using this procedure.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the Settings page, click **Nightly Jobs**.
 - Scroll down on the right side of the Settings page until you can see the Nightly Jobs heading.
3. To change any nightly job, or to change the time that nightly jobs begin to execute, click  **Change**. The Nightly Jobs dialog box displays.
4. If you want to change the time nightly jobs execute, enter a new time in the **Nightly job times** text box.
5. In the first column, click to select each nightly jobs option you want to set for the Core. Click any selected option to clear it.
6. Click **OK**.
The Nightly Jobs dialog box closes and your nightly jobs settings for the Core are saved.

Modifying transfer queue settings

Transfer queue settings are Core-level settings that establish the maximum number of concurrent transfers and the maximum number of retries for transferring data.

Complete the steps in this procedure to modify transfer queue settings.




1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Transfer Queue**.
 - Scroll down on the right side of the *Settings* page until you can see the Transfer Queue heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.
4. Enter the configuration information as described in the following table.


Table 22: Transfer queue settings information

Text Box	Description
Maximum concurrent transfers	Enter a value to update the number of concurrent transfers. Set a number from 1 to 60. The smaller the number, the lesser the load on network and other system resources. As the number of agents that are processed increases, so does the load on the system.
Maximum retries	Enter a value to set the maximum number of attempts before canceling the transfer operation. Set a number from 1 to 60.


5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Adjusting client timeout settings

Client timeout settings control the length of time that various operations are attempted before the operation times out.





 **NOTE:** Quest recommends leaving default timeout settings unless you experience specific issues in your environment, and you are advised by a Quest Data Protection Support representative to modify the settings.



Complete the steps in this procedure to adjust client timeout settings.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Client Timeout**.
 - Scroll down on the right side of the *Settings* page until you can see the Client Timeout heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.

4. Enter the configuration information as described in the following table.

Table 23: Client timeout settings information

Setting	Description
Connection timeout	<p>Controls the timeout for the connection between the Core and protected machines when sending data across the hypertext transfer protocol (http).</p> <p>Enter the amount of time you want to lapse before a connection timeout occurs. Uses HH:MM:SS format.</p> <p> NOTE: The default setting is 0:05:00 or five minutes.</p>
Read/Write timeout	<p>Controls the timeout for the connection between the Core and protected machines when reading or writing stream data across http. An example is receiving changed data blocks from a protected machine to the Core for an incremental snapshot.</p> <p>Enter the amount of time you want to lapse before a timeout occurs during a read/write event. Uses HH:MM:SS format.</p> <p> NOTE: The default setting is 0:05:00 or five minutes.</p>
Connection UI timeout	<p>Controls the timeout for the connection between the graphic user interface and the Rapid Recovery Core service across http.</p> <p>Enter the amount of time you want to lapse before a connection UI timeout occurs. Uses HH:MM:SS format.</p> <p> NOTE: The default setting is 0:05:00 or five minutes.</p>
Read/Write UI timeout	<p>Controls the timeout for the connection for reading and writing data streams between the graphic user interface and the Rapid Recovery Core service across http.</p> <p>Enter the amount of time you want to lapse before a timeout occurs during read or write events. Uses HH:MM:SS format.</p> <p> NOTE: The default setting is 0:05:00 or five minutes.</p>

5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Understanding deduplication cache and storage locations

Global deduplication reduces the amount of disk storage space required for data your Core backs up. Each repository is deduplicated, storing each unique block once physically on disk, and using virtual references or pointers to those blocks in subsequent backups. To identify duplicate blocks, Rapid Recovery includes a deduplication cache for deduplication volume manager (DVM) repositories. The cache holds references to unique blocks.

By default, for DVM repositories, this deduplication cache is 1.5GB. This size is sufficient for many repositories. Until this cache is exceeded, your data is deduplicated across the repository. Once the amount of redundant information

is so great that the deduplication cache is full, your repository can no longer take full advantage of further deduplication for newly added data. The amount of data saved in your repository before the deduplication cache fills varies by the type of data being backed up, and is different for every user.

You can increase or decrease the size of the deduplication cache by adjusting the **Deduplication cache size** value in the DVM Deduplication Cache Core settings. For more information on how to increase the cache size, see the topic [Configuring DVM deduplication cache settings](#).

When you increase the DVM deduplication cache size, there are two factors to consider: disk space and RAM usage.

Disk space. Two copies of the DVM deduplication cache are stored on disk: a primary cache, and a secondary cache which is a parallel copy. Thus, if using the default cache size of 1.5GB for a DVM repository, 3GB of disk storage is used in your system. As you increase the cache size, the amount of disk space used remains proportionally twice the size of the cache. To ensure proper and fault-resistant performance, the Core dynamically changes the priority of these caches. Both are required, the only difference being that the cache designated as primary is saved first.

RAM usage. When the Rapid Recovery Core starts, it loads the deduplication cache to RAM. The size of the cache therefore affects memory usage for your system. The total amount of RAM the Core uses depends on many factors. These factors include which operations are running, the number of users, the number of protected machines, and the size of the deduplication cache. Each operation the Core performs (transfer, replication, rollup, and so on) consumes more RAM. Once an operation is finished, memory consumption decreases accordingly. However, administrators should consider the highest RAM load requirement for efficient operations.

Default settings for the Rapid Recovery Core place the primary cache, secondary cache, and the metadata cache for DVM repositories in the AppRecovery directory. This folder is installed on the Core machine.

NOTE: Depending on your settings, the AppRecovery directory may not be visible on the Rapid Recovery Core. To see this directory, you may need to change the Folder Options control panel to show hidden files, folders, and drives.

Assuming the Rapid Recovery Core is installed on the C drive, these locations are typically as follows:

Table 24: Default storage locations for DVM deduplication cache settings

Setting	Default Storage Location
Primary cache location	C:\ProgramData\AppRecovery\RepositoryMetaData\PrimaryCache
Secondary cache location	C:\ProgramData\AppRecovery\RepositoryMetaData\SecondaryCache
Metadata cache location	C:\ProgramData\AppRecovery\RepositoryMetaData\CacheMetadata

Optimizing deduplication

When you install Rapid Recovery Core, the installer defaults to the C : \ drive of the Core machine to store the deduplication cache used for DVM repositories.

To optimize deduplication performance, Quest recommends installing the deduplication cache to a separate fast physical storage location. For maximum performance, install the deduplication cache on a solid-state drive (SSD). You can change the storage location of these caches. For example, for increased fault tolerance, you can change location of your secondary cache to a different physical drive than the primary cache, assuming the Rapid Recovery Core has access to the location.



For more information on how to change storage locations for any of these settings, see the topic [Configuring DVM deduplication cache settings](#).

For conceptual information about deduplication, see [Deduplication in Rapid Recovery](#).

Configuring DVM deduplication cache settings



For conceptual information about deduplication, see [Deduplication in Rapid Recovery](#). For recommendations about settings to use, see [Understanding deduplication cache and storage locations](#).



Complete the steps in this procedure to configure deduplication cache settings for DVM repositories.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the Settings page, click **DVM Deduplication Cache**.
 - Scroll down on the right side of the *Settings* page until you can see the DVM Deduplication Cache heading.
3. If you want to restore default DVM deduplication cache settings at any time, do the following:
 - a. At the top of the deduplication cache settings area, click  **Restore Default**.
The *Restore Default* dialog box appears
 - b. Click **Yes** to confirm the restore.
4. Click the setting you want to change.
The setting you selected becomes editable.

5. To change individual deduplication cache settings, enter the configuration information as described in the following table.

Table 25: DVM deduplication cache settings information

Setting	Description
 Restore Default	This control resets DVM cache locations to system default locations, which are described for each setting.
Primary cache location	<p>If you want to change the primary cache location for DVM repositories, then in the Primary cache location text box, type the path for a storage location accessible to the Core.</p> <p>The default location is:</p> <p><code>C:\ProgramData\AppRecovery\RepositoryMetaData\PrimaryCache</code></p> <p>Since the primary and secondary caches are the same size, collective storage for these two caches requires twice the amount of space as the amount allocated for the deduplication cache size. For example, if you specify the default amount of 1.5GB for the deduplication cache size, ensure that each of the two storage locations have at least 1.5GB. In particular, if both locations belong to the same drive (for example, the C drive), there must be at least 3.0GB of free disk space.</p>
Secondary cache location	<p>If you want to change the secondary cache location for DVM repositories, then in the Secondary cache location text box, type the path for a storage location accessible to the Core.</p> <p>The default location is:</p> <p><code>C:\ProgramData\AppRecovery\RepositoryMetaData\SecondaryCache</code></p>
Cache metadata location	<p>If you want to change the cache metadata location for DVM repositories, then in the Cache metadata location text box, type the path for a storage location accessible to the Core.</p> <p>The default location is:</p> <p><code>C:\ProgramData\AppRecovery\RepositoryMetaData\CacheMetadata</code></p>
Deduplication cache size (GB)	<p>If you want to change the deduplication cache size for DVM repositories, then in the Deduplication cache size (GB) text box, enter a new amount in gigabytes.</p> <p> NOTE: If you decrease the size of the deduplication cache, the existing contents of the cache are flushed, and the cache is recreated.</p> <p>The default location is:</p> <p><code>C:\ProgramData\AppRecovery\RepositoryMetaData\SecondaryCache</code></p> <p>The minimum cache size setting is 1.5GB. Additionally, the cache size cannot exceed 50 percent of the installed RAM.</p>

6. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Configuring Replay engine settings

You can configure information regarding the Replay engine, which is the communication channel for Rapid Recovery. These settings determine Core settings to provide effective communication.

In general, Quest recommends using default settings. In some cases, you may be directed by Quest Data Protection Support to modify these settings to help adjust the performance specific to your network needs. Complete the steps in this procedure to configure Replay engine settings.




1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Replay Engine**.
 - Scroll down on the right side of the *Settings* page until you can see the Replay Engine heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.
4. Enter the configuration information as described in the following table.

Table 26: Replay engine settings information

Text Box	Description
IP address	<p>The Core uses this IP address when performing mount and restore for a recovery point, to allow feedback between protected machines and the Core.</p> <p>The IP address for the Replay engine automatically populates with the IP address of the Core machine. If you manually enter the server IP address, then this value is used in cases where the protected machine cannot resolve the automatically provided IP address.</p> <p>You do not need to set this value manually unless you are having issues with protected machines being able to communicate with the Core.</p>
Preferred port	<p>Enter a port number or accept the default setting. The default port is 8007.</p> <p>The port is used to specify the communication channel for the Replay engine.</p>
Port in use	Represents the port that is in use for the Replay engine configuration.
Allow port auto-assigning	Click for allow for automatic TCP port assignment.
Admin group	Enter a new name for the administration group. The default name is BUILTIN\Administrators.
Minimum asynchronous I/O length	<p>Enter a value or choose the default setting. It describes the minimum asynchronous input/output length.</p> <p>The default setting is 65536.</p>
Read timeout	Enter a read timeout value or choose the default setting. The default setting is 00:05:00.
Write timeout	Enter a write timeout value or choose the default setting. The default setting is 00:05:00.
Receive buffer size	Enter an inbound buffer size or accept the default setting. The default setting is 8192.
Send buffer size	Enter an outbound buffer size or accept the default setting. The default setting is 8192.
No delay	It is recommended that you leave this check box unchecked as doing otherwise will impact network efficiency. If you determine that you need to modify this setting, contact Quest Data Protection Support for guidance in doing so.

5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Configuring deployment settings

Rapid Recovery lets you download installers from the Rapid Recovery Core to machines you want to protect.



NOTE: You can also download Rapid Recovery software from the QorePortal.

You can configure settings related to the deployment of the Rapid Recovery Agent software from your Core to the machines you want to protect.

Complete the steps in this procedure to configure deployment settings.




1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the Settings page, click **Deploy**.
 - Scroll down on the right side of the *Settings* page until you can see the **Deploy** heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.
4. Enter the configuration information as described in the following table.

Table 27: Deployment settings information

Text Box	Description
Agent installer name	The default filename is Agent-Web.exe. If you wish to change this file name for any reason, you can use this setting to specify a new name of the Core Web Installer executable file. This file streams a download of the latest version of the Rapid Recovery Core installer, which runs directly from the web and lets you pause and resume the process as needed.
Core address	Enter the address of your Core server. This typically consists of the protocol, the name of your Core server and port, and the directory where the Core files reside. For example, if your server is <i>Sample</i> , this setting is <code>https://sample:8006/apprecovery/admin/Core.</code>
Failed receive timeout	The amount of time deployment of the Agent software should be attempted before timing out. The default setting is 00:25:00 or twenty-five minutes. If you wish to change this setting, enter the length of time you want the system to attempt to deploy the Agent software before a timeout occurs during read or write events. Uses HH:MM:SS format.
Maximum parallel installs	This setting controls the maximum number of deployments of the Agent software for the Core to attempt at one time. The default setting is 100.

5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Core-level tools

In addition to configuring Core settings, you can also use the Core-level tools described in the following table.

Table 28: Other Core-level tools



UI Element	Description
System information	<p>Rapid Recovery lets you view information about the Rapid Recovery Core that includes system information, local and mounted volumes, and Replay engine connections.</p> <p>For more information on the information displayed on the System information page, see Understanding system information for the Core.</p> <p>For more information on how to view System information, see Viewing system information for the Core.</p>
Downloading Core log files	<p>Information about various activities for the Rapid Recovery Core are saved to the Core log file. To diagnose possible issues, you can download and view logs for your Rapid Recovery Core. For more information on accessing and viewing the Core logs, see Accessing Core logs.</p> <p>Each protected machine also saves a log of activity. This log can be uploaded to the Core if you select the nightly job called Downloading the logs from the protected machines. For more information about nightly jobs, see Understanding nightly jobs. For more information about how to configure nightly job settings for the Core, see Configuring nightly jobs for the Core. For more information about configuring nightly jobs for specific protected machines, see Customizing nightly jobs for a protected machine.</p>

Viewing system information for the Core

System information for the Core includes general information, information about Core volumes, microprocessors, and Replay Engine connections. For a detailed description of the information available on this page, see [Understanding system information for the Core](#).

i **NOTE:** You can also see system information for a specific protected machine. For more information, see [Viewing system information for a protected machine](#).

Complete the steps in this procedure to view system information for the Core.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More) and then click  **System Information**.
The System Information page appears.

Understanding system information for the Core

Rapid Recovery lets you view information about the Rapid Recovery Core. You can view [general information](#), information about [local volumes](#) on your Core machine, [processor information](#) for your Core, and [Replay Engine](#) connections.

i **NOTE:** You can also see system information for a specific protected machine. For more information, see [Viewing system information for a protected machine](#).

In the *General* pane, you can see information about your Core as described in the following table.

Table 29: Core system information

UI Element	Description
Host name	The machine name of your Rapid Recovery Core.
OS version	The version of the operating system installed on the Rapid Recovery Core.
OS architecture	Lists the underlying structure and design of the machine hosting your Rapid Recovery Core. Potentially includes chipset and lists 64-bit system. Rapid Recovery Core supports 64-bit systems only.
Memory (physical)	Lists the amount of Random Access Memory installed on the Core machine.
Display name	Shows the display name of the Core, which is configurable (see Configuring Core general settings).
Fully qualified domain name	Shows the fully qualified domain name for the Core machine.
Metadata cache location	Shows the path of the metadata cache location. For more information, see Understanding deduplication cache and storage locations .
Primary cache location	Shows the path of the primary deduplication cache location. For more information, see Understanding deduplication cache and storage locations .
Secondary cache location	Shows the path of the secondary deduplication cache location. For more information, see Understanding deduplication cache and storage locations .

The *Volumes* pane includes the following information about storage volumes for the Core machine: Name, device ID, file system, raw capacity, formatted capacity, used capacity, and mount points.

The *Processors* pane displays detailed information about processors on the Core, including architecture, the number of Cores, the number of threads, clock speed, and description, if any.

The *Replay Engine Connections* pane displays detailed information about currently mounted recovery points. You can view the local end point, remote end point, mounted image agent ID, the authenticated user, number of bytes read, and number of bytes written.

You can dismount recovery points that are mounted locally on a Core from the *Mounts* page. For more information about dismounting recovery points, see [Dismounting recovery points](#).

For more information, see [Viewing system information for the Core](#).

Related topics:

- [Viewing system information for the Core](#)
- [Configuring Core general settings](#)
- [Understanding deduplication cache and storage locations](#)
- [Dismounting recovery points](#)

Accessing Core logs

Information about various activities for the Rapid Recovery Core are saved to the Core log file. This file, `AppRecovery.log`, is stored by default in the path `C:\ProgramData\AppRecovery\Logs`.

NOTE: Depending on your settings, the AppRecovery directory may not be visible on the Rapid Recovery Core. To see this directory, you may need to change the Folder Options control panel to show hidden files, folders, and drives. If these settings include the option to hide extensions for known file types, the Core log file may appear as AppRecovery with no `.log` extension.

The Core log includes information about completed Core jobs, connection failures, results of attempts on the part of the Core to contact the License Portal, and other information. Each statement stored in the Core log file is preceded by one of four qualifiers: INFO, DEBUG, ERROR, and WARN. These qualifiers help categorize the nature of information stored in the log when diagnosing an issue.




NOTE: Similarly, a log file is also stored on each protected machine containing information relating to its attempts at communicating with the Core. For more information about machine logs, see [Downloading and viewing the log file for a protected machine](#).

The ability to access logs can be useful when troubleshooting an issue or working with Quest Data Protection Support. To access logs, see the following procedures:

- [Downloading and viewing the Core log file](#)
- [Downloading and viewing the log file for a protected machine](#)

Downloading and viewing the Core log file

If you encounter any errors or issues with the Core, you can download the Core logs to view them or to share them with your Quest Support representative.

1. From the Rapid Recovery Core Console, on the icon bar, click  (More) and then click  **Core Log**.
2. On the Download Core Log page, click  **Click here to begin the download**.
3. If prompted to open or save the Core `AppRecovery.log` file, click **Save**.
4. If you see the Opening Core AppRecovery.log dialog box, do one of the following:
 - To open the log file, select **Open with**, then select an application (such as Notepad) for viewing the text-based log file, and finally click **OK**.
The Core `AppRecovery.log` file opens in the selected application.
 - To save the file locally, select **Save File** and click **OK**.
The Core `AppRecovery.log` file saves to your Downloads folder. It can be opened using any text editor.

Configuring database connection settings

Rapid Recovery displays information about Core tasks, events, and alerts on the *Events* page. Rapid Recovery stores this transactional information in a MongoDB service database that is installed locally by default on the Core machine. You can configure these settings to change how long information is retained in the database, or to change the connection pool size to allow for more or fewer concurrent connections.




If using a second Rapid Recovery Core, you can configure the database connection settings on the first Core to point to the second Core machine. In this way, the event data for both Cores will be stored in the MongoDB on the second Core.

Alternatively, you can configure the database connection settings on the Core to point to another machine that has a separately installed MongoDB which is accessible over the network to the Rapid Recovery Core. The event transaction data for your Core is then saved to that service database, not locally.

i **NOTE:** For more information about viewing event information from the Rapid Recovery Core, see [Viewing events using tasks, alerts, and journal pages](#).







Customers can choose to specify installation of the MongoDB service database on another machine accessible on the network to the Rapid Recovery Core. If the service database for your Rapid Recovery Core is installed on a machine other than the machine hosting the Rapid Recovery Core, you must provide database credentials (a user name and password) in these settings.



Complete the steps in this procedure to modify the database connection settings for the service database used by the Rapid Recovery Core.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Database Connection**.
 - Scroll down on the right side of the *Settings* page until you can see the Database Connection heading.
3. From the top of the Database Connection settings area, you can do the following:
 - Click  **Test Connection** to verify your settings.
Testing the connection is recommended when you change any of the database connection settings.
 - Click  **Restore Default** to restore all default database connection settings.
You are prompted to confirm this action, which results in abandoning any customized database connection settings.
4. Click on the setting you want to change.
The setting you selected becomes editable.

5. Enter the configuration information as described in the following table.

Table 30: Database connection settings information


Text Box	Description
Host name	Enter a host name for the database connection.  NOTE: When localhost is the parameter specified as the host, the MongoDB is installed locally on the machine hosting the Core.
Port	Enter a port number for the database connection.  NOTE: The default setting is 27017.
User name	Enter the name of a user with administrative privileges to the MongoDB service database.  NOTE: If the host name parameter is localhost, this field is not required.
Password	Enter the password associated with the user name you specified.  NOTE: If the host name parameter is localhost, this field is not required.
Retention period (day)	Enter the number of days to retain the event and job history in the service database.
Maximum connection pool size	Sets the maximum number of database connections cached to allow dynamic reuse.  NOTE: The default setting is 100.
Minimum connection pool size	Sets the minimum number of database connections cached to allow dynamic reuse.  NOTE: The default setting is 0.

6. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Modifying local database connection settings

You can view system events related to the Rapid Recovery Core on the *Events* page. The Rapid Recovery Core stores this transactional information in a MongoDB service database. By default, this database is installed locally on the Core machine, and the hostname in the database connection settings defaults to localhost. In this situation, the loopback interface bypasses local network interface hardware, and database credentials are not required.

Optionally, to increase security, you can explicitly specify database credentials (a user name and password) for the MongoDB database used by the Rapid Recovery Core.

 **NOTE:** For more information about viewing event information from the Rapid Recovery Core, see [Viewing events using tasks, alerts, and journal pages](#). For information about database connection settings, see [Configuring database connection settings](#).

Complete the steps in this procedure to modify the local database connection settings to specify database credentials.




1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Local Database Settings**.
 - Scroll down on the right side of the *Settings* page until you can see the Local Database Settings heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.
4. Enter the appropriate credentials for connecting to the service database, as described in the following table.

Table 31: Local database settings information

Text Box	Description
User name	Enter the name of a user with administrative privileges to the MongoDB service database.
Password	Enter the password associated with the user name you specified.

5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Managing SMTP server settings


If you configure simple mail transfer protocol (SMTP) server settings for the Core, you can send task, event, and alert notifications by email.


Information about configuring an SMTP email server is described in the topic [Configuring an email server](#).

i **NOTE:** To send event information by email, you must also configure notification group settings. For more information on specifying events to receive email alerts, see [Configuring notification groups](#).

Configuring cloud account connection settings

Cloud account connection settings let you determine how much time should pass between Rapid Recovery attempts to connect to your cloud account before the operation times out. Complete the steps in the following procedure to configure the connection settings for your cloud account.


1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  **Settings**.
The *Settings* page appears.
3. In the left menu, click  **Cloud Accounts**.

4. In the Cloud Accounts table, click the  drop-down menu next to the cloud account you want to configure, and then complete one of the following actions:
 - To return any custom cloud configuration settings to the following default settings, click **Reset**.
 - **Request timeout** 01:30 (minutes and seconds)
 - **Write buffer size** 8388608 (bytes)
 - **Read buffer size** 8388608 (bytes)
 - To change the cloud account connection settings, click **Edit**, and in the Cloud Configuration dialog box, complete any of the following actions:
 - For **Request timeout**, use the up and down arrows to determine the amount of time in minutes and seconds that Rapid Recovery should spend on a single attempt to connect to the cloud account when there is a delay. Connection attempts will cease after the entered amount of time.
 - For **Write buffer size**, enter the buffer size you want to reserve for writing archived data to the cloud.
 - For **Read buffer size**, enter the block size you want to reserve for reading archived data from the cloud.

Managing report settings



You can generate reports for the Rapid Recovery Core or for protected machines. For information on the reports you can generate, see [Reporting](#).

Complete the steps in this procedure to manage report settings for Core reports.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Reports**.
 - Scroll down on the right side of the *Settings* page until you can see the Reports heading.

The Reports Core settings appear. Report settings are described in the following table.

Option	Description
Restore Default	This option restores all report settings to the default settings. Defaults are listed below for each setting.
Font	<p>This option controls the default font used for reports. The default typeface is Trebuchet MS. You can change this font to any typeface available to your system.</p> <p>This option controls the default paper size for printing reports. The default is letter. You can choose from the following paper sizes:</p> <ul style="list-style-type: none"> • A3 • A4
Paper size	<ul style="list-style-type: none"> • B4 • Executive • Ledger • Legal • Letter • Tabloid
Page orientation	<p>This option controls the page orientation for exported reports. The default orientation is Landscape. You can choose from the following layout options:</p> <ul style="list-style-type: none"> • Landscape • Portrait

3. To change any of the settings for Reports, click in the appropriate setting field. The setting field appears as a configurable drop-down menu.
4. Click the drop-down menu, and select one of the values available. For example, in the Font field, click **Times New Roman**.
5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving. The option you selected now appears as the new setting for the selected Reports parameter.

Managing Core SQL attachability settings

SQL attachability checks occur as part of the Rapid Recovery nightly jobs. To ease licensing costs, Rapid Recovery gives you two options for performing attachability checks: using a licensed instance of SQL Server installed on the Rapid Recovery Core machine or using the instance of SQL Server already installed on your protected machine.

This second option is now the default setting. However, if your protected machine is already exerted during the time when the nightly jobs occur, consider performing the checks with an instance of SQL Server on the Core.

The ability to perform attachability checks using the SQL Server instance on a protected machine is a function of the Rapid Recovery Agent software. This feature does not work if the SQL Server is protected agentlessly. Therefore, if using agentless protection on the SQL machine, configure this Core setting to **Use SQL Server on the Core**, as shown in [step 3](#) of this procedure.

In summary, the process of managing Core SQL attachability settings involves the following tasks:


- Mount the latest recovery point for protection groups containing databases.
- Connect to the database from SQL Server.
- Open the database.
- Close the database.
- Dismount the recovery point.

To enable this nightly check, specify a SQL Server instance to use to perform attachability checks for SQL Server databases on protected machines.

i | **NOTE:** This option does not appear if you are not protecting a SQL Server in your Core.

To configure the Core to perform SQL attachability checks as part of the nightly jobs, complete the following steps.

i | **NOTE:** If you select the default option to use the instance of SQL Server installed on the protected machine, that SQL Server instance will manage SQL attachability for all protected SQL machines. If you do not want this setting to apply to all protected SQL machines, select **Use SQL Server on the Core**. To perform attachability checks on the Core, you must install or use a licensed version of SQL Server on the Core machine.


1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Attachability**.
 - Scroll down on the right side of the *Settings* page until you can see the Attachability heading.
3. To use the SQL Server instance installed on the protected SQL Server machine, select **Use SQL Server on the protected machine**. This is the default option.

4. To use the SQL Server instance installed on the Rapid Recovery Core, select **Use SQL Server on the Core**, and then enter the authentication information as described in the following table.

Table 32: SQL Server credentials information

Text Box	Description
SQL Server	From the SQL Server drop-down menu, select the appropriate SQL Server instance from the Core server.
Credential Type	Select the appropriate authentication method for your credentials from the following options: <ul style="list-style-type: none">• Windows• SQL
User Name	Specify a user name for accessing the SQL Server on the Core based on your selected credential type.
Password	Specify a password for accessing the SQL Server on the Core based on your selected credential type.

5. Click **Test Connection**.

 **NOTE:** If you entered the credentials incorrectly, a message displays to alert you that the credentials failed. Correct the credential information and test the connection again.

6. After you are satisfied with your changes, click **Apply**.

Understanding Core jobs

Core jobs are processes that the Rapid Recovery Core performs to support its operations, including backing up to recovery points, replicating data, archiving data, exporting data to VMs, maintaining repositories, and so on. Core jobs are initiated automatically for some operations, such as replicating or archiving on an established schedule. You can also invoke some jobs on demand from various elements on the Core Console.

- When viewing or editing Core job settings, each Core job has two parameters: Maximum concurrent jobs and Try count.
 - The Maximum concurrent jobs parameter determines how many jobs of that type can be run at the same time.
 - The Try count parameter determines how many times the job should be tried before abandoning the job, if network or other communication errors prevent the job from succeeding the first time.
- In the Core Jobs table, the Settings column indicates if the job listed is included in Core job settings by default or must be explicitly added.

The following table describes the primary Core jobs available and their functions.

Table 33: Core jobs

Job Name	Description	Maximum Concurrent Jobs	Try Count	Settings
Check attachability of SQL databases in snapshots	<p>Lets the Core check the consistency of SQL databases and ensures that all supporting MDF (data) and LDF (log) files are available in the backup snapshot. Process:</p> <ul style="list-style-type: none"> • Mount the latest recovery point for protection groups containing SQL databases. • Mount the database. If performing attachability from the protected SQL server, mount using UNC path. • Connect to the database from SQL Server. • Perform the attachability check. • Perform cleanup operations. • Close the database. • Dismount the database. • Dismount the recovery point. 	1	0	Default
Check checksum of Exchange databases	<p>Checks the integrity of recovery points containing Exchange databases. Process:</p> <ul style="list-style-type: none"> • Mount the latest recovery point for protection groups containing SQL databases. • Connect to the database from SQL Server. • Open the database. • Close the database. • Dismount the recovery point. 	1	0	Default
Check mountability of Exchange databases	Checks that Exchange databases are mountable.	1	0	Default
Replicate protected machines data from remote source	Transfers a copy of recovery points for a protected machine from a source Core to a target Core. This job runs on the target Core receiving the incoming replicated recovery points.	3	0	Default
Replicate protected machines data to	Transfers a copy of recovery points for a protected machine from a source Core (on which they were originally saved) to a target Core. This job runs on the source Core and controls outgoing replication.	1	3	Default

Job Name	Description	Maximum Concurrent Jobs	Try Count	Settings
remote target				
Roll up recovery points	Applies the retention policy to your backed-up data by combining or "rolling up" recovery points on the schedule defined in the retention policy.	1	0	Default
Check recovery points	Checks the integrity of recovery points.	1	0	Add
Delete all recovery points	Deletes the full set of recovery points on a protected machine.	1	0	Add
Delete chain of recovery points	Deletes a complete recovery point chain on a protected machine.	1	0	Add
Delete range of recovery points	Deletes a set of recovery points on a protected machine, by recovery point identifier or date range.	1	0	Add
Deploy Agent software to machines	Deploys Rapid Recovery Agent software to the specified machine or machines.	1	0	Add
Download Exchange libraries	Downloads Microsoft Exchange libraries from the protected machine to the Core machine at path C:\ProgramData\AppRecovery\ExchangeLibraries.	1	0	Add
Export to archive	Creates backup in the specified path with an archive of the selected recovery points. Process: <ul style="list-style-type: none"> Mount recovery points. Write data to backups. Dismount the recovery point. 	1	0	Add
Export to virtual machine	Exports data from specified recovery point of protected machine to destination path as a virtual machine. Process: <ul style="list-style-type: none"> Mount recovery point. Create virtual machine from the recovery point data in the destination path. Dismount the recovery point. 	1	0	Add

Job Name	Description	Maximum Concurrent Jobs	Try Count	Settings
Import archives	Imports recovery point from the specified backup on a previously created Core archive.	1	0	Add
Maintain repository	Performs a check of the repository. Process: <ul style="list-style-type: none"> • Check repository file system. • Mount recovery point. • Recalculate deduplication cache for repository. • Load recovery points from repository. 	1	0	Add
Mount recovery point snapshots	Performs mount of recovery point to the specified path.	1	0	Add
Protect ESX virtual machines	Adds all specified virtual machines to agentless protection. Job is performed immediately after adding agentless protection of one or more VMs to the Core using the Protect Multiple Machines Wizard. Job sets ID number for each specified VM, writes information about the Core to a configuration file, and retrieves metadata from the file.	1	0	Add
Restore from recovery point	Performs a restore from a recovery point to a specified target machine. Process: <ul style="list-style-type: none"> • Mount recovery point. • Write all data from the recovery point to the specified machine. • Dismount the recovery point. 	1	0	Add
Uploading logs	Uploads logs to specified server.	1	0	Add


Some Core jobs are included in Settings. The Jobs settings let you specify how many concurrent jobs of the same type the Core can run, and how many retries should be attempted if the first job attempt fails.

For more information about these Settings, see [Core job settings](#).

For information on adding jobs to Core Settings, see [Adding Core jobs to settings](#).

For information on editing settings for jobs in the Settings list, see [Editing Core job settings](#).

Core job settings

When you select  (Settings) from the icon bar, you can access settings for some Core jobs. The *Jobs* area on the Core settings page lets you determine two settings for each job type listed:

1. The maximum number of jobs of this type for the Core to attempt at one time. This must be set to a value between 1 to 50.
2. The number of times a job should be attempted if a network or other communication error prevents the job from succeeding the first time. This must be set to a value between 0 to 10.

Several jobs are automatically included in Core settings. These jobs include a value of "Default" in the Settings column (as shown in the topic [Understanding Core jobs](#)).

You can add some other jobs to settings if you want to configure those settings to control the maximum number of jobs or retries for those functions. These jobs include a value of "Add" in the Settings column. For information on how to add these jobs to the Settings table, see [Adding Core jobs to settings](#).

Core jobs not available in Settings do not provide the ability to set these two parameters.

For jobs that are listed in settings, you can edit existing settings. This lets you customize the two parameters, delete a job type from the job settings list, or restore default settings. For detailed information, see the topic [Editing Core job settings](#).


Adding Core jobs to settings

Core job settings let you define, for each job type, the maximum number of jobs for the Core to attempt at one time, and how many times that job should be retried if the first attempt failed.

Each Core job type has default values for these two parameters, as described in the topic [Core job settings](#). This list also indicates which of the job types are included in the Core settings by default.

Adding a Core job to settings lets you change these parameters for the job type you added.

Complete the steps in the following procedure to add a job to Core settings.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the **Settings** page, click **Jobs**.
 - Scroll down on the right side of the Settings page until you can see the **Jobs** heading. The Jobs settings for the Core appear.
3. On the Core **Settings** page, under Jobs, click **+ Add**.
The **Job Settings** dialog box appears.
4. In the **Job Settings** dialog box, from the **Jobs** field, select the name of a job you want to add to the Core settings.
These jobs are described in the topic [Core job settings](#).
5. To set the maximum number of jobs for the Core to attempt at one time, in the **Maximum concurrent jobs** text box, enter a new value between 1 to 50.
6. To set the number of attempts the Core should make before abandoning the job, in the **Try count** text box, enter a new value between 0 and 10.
7. Click **Save**.
The Job Settings dialog box closes, and your new job settings are applied.

Editing Core job settings

Core job settings let you define, for each job type, the maximum number of jobs for the Core to attempt at one time, and how many times that job should be retried if the first attempt failed.

Each Core job type has default values for these two parameters, as described in the topic [Understanding Core jobs](#). This list also indicates which of the job types are included in the Core settings by default. When you edit Core job settings, you can accomplish the following:





- You can customize the settings for each Core job type.
- You can delete a job type from the list of Core settings. This feature is not available if the job type is included in settings by default.

i **NOTE:** Deleting a job from Core settings simply removes the job type from this list. To edit Core settings for that job type again in the future, you can add it to the list as described in the topic [Adding Core jobs to settings](#).

- You can restore the settings for any job type to the default settings.

i **NOTE:** Although you can only use this feature for the job types included in Core settings by default, you can set other job types to defaults by removing them from the list and adding them again.

Complete the steps in the following procedure to edit the settings of a job.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the Settings page, click **Jobs**.
 - Scroll down on the right side of the Settings page until you can see the **Jobs** heading.The Jobs settings for the Core appear.
3. From the Job grid, select a job you want to remove from the list. From the drop-down  menu for that job, select **Delete**.
The job is removed from the list.
4. From the Job grid, select a job from the list for which you want to reset settings. From the drop-down  menu for that job, select **Reset to defaults**.
The job settings for this job are reset to default settings.
5. From the Job grid, select a job you want to change. From the drop-down  menu for that job, select **Edit**.
The *Job Settings: [JobName]* dialog box opens.
6. To change the maximum number of jobs for the Core to attempt at one time, in the **Maximum concurrent jobs** text box, enter a new value between 1 to 50.
7. To change the setting for the number of additional attempts the Core should make before abandoning the job, in the **Try count** text box, enter a new value between 0 and 10.
8. Click **Save**.
The *Job Settings* dialog box closes, and your new job settings are applied.

Understanding SNMP settings

Simple Network Management Protocol (SNMP) is a protocol for managing devices on an IP network. SNMP is used primarily to monitor devices on a network for conditions that require attention. This protocol uses software components (agents) to report information to administrative computers (managers). An SNMP agent handles the manager's requests to get or set certain parameters. The SNMP agent can send traps (notifications about specific events) to the manager.

Data objects that the SNMP agents manage are organized into a Management Information Base (MIB) file that contains Object Identifiers (OIDs). Each OID identifies a variable that can be read or set using SNMP.

Rapid Recovery includes support for SNMP version 1.0.

You can configure the Rapid Recovery Core as an SNMP agent. The Core then can report information such as alerts, repository status, and protected machines. An SNMP host can read this information using a standalone application called an SNMP browser. You can install the SNMP browser on any machine accessible over the network to the Rapid Recovery Core.

To ensure the Core SNMP event notifications can be received by the SNMP browser, verify that the notification options for a notification group are properly configured to notify by SNMP trap.

i | **NOTE:** You can use the default group, or create a custom notification group. The process is identical.

Open the notification group, select the **Notification Options** tab, and ensure the **Notify by SNMP trap** option is enabled. The notification group specifies trap number 1 by default. If necessary, you can change the trap number to ensure that it matches the setting that the SNMP browser expects.

For more information and specific details about configuring notification options, see [Configuring notification groups](#).

Alternatively, you can download a MIB file from the Rapid Recovery Core. This file is readable using an SNMP browser in a more user-friendly fashion than data it receives directly from the Core.

This section includes the following topics:


- [Configuring SNMP settings](#)
- [Downloading the SNMP MIB file](#)

Configuring SNMP settings

Use the SNMP settings to control communication, such as alerts, between the Rapid Recovery Core and an SNMP browser. Available settings include the incoming and outgoing SNMP ports, trap receiver port, and the host name for the trap receiver.

i | **NOTE:** Rapid Recovery builds prior to release 6.1 do not include the ability to change the Community string setting. Release 6.4 and later have the option to specify an outgoing SNMP port.






Use this procedure to configure SNMP settings for the Core.



1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **SNMP Configuration**.
 - Scroll down on the right side of the *Settings* page until you can see the SNMP Configuration heading.

The SNMP Configuration settings are displayed.

3. Modify the SNMP settings as described in the following table.

Table 34: SNMP connection settings information

Text Box	Description
Handle incoming request	To let the Core recognize incoming SNMP protocols, select this option. To block incoming SNMP protocols, clear the option.  NOTE: Selecting the option to handle incoming requests lets you edit the community string setting.
Community string	Enter a name for the community.  NOTE: You can only change this setting if the Handle incoming request setting is set to Yes .
Incoming port	Enter a port number for the SNMP connection.  NOTE: The default setting is 8161.
Outgoing port	Optionally, enter a port number for the outgoing SNMP connection.  NOTE: If not explicitly set, Rapid Recovery Core chooses the most appropriate port.
Send traps	To allow alerts (traps) to be sent using the SNMP protocol, select this option. To block alerts, clear the option.
Trap receiver port	Enter a port number for the incoming alert. The default setting is 162.
Trap receiver host name	Enter a host name for the SNMP connection.  NOTE: The default host name is localhost.

4. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.


Downloading the SNMP MIB file

The Simple Network Management Protocol is used to monitor devices on a network for conditions that require attention. When the Rapid Recovery Core is set as an SNMP agent, the Core report information such as alerts, repository status, and protected machines. This information can be read by an SNMP host using a standalone application called an SNMP browser.

Data objects managed by SNMP agents are organized into a Management Information Base (MIB) file that contains Object Identifiers (OIDs). Each OID identifies a variable that can be read or set using SNMP.

You can download a MIB file from the Rapid Recovery Core. This file, named `quest-rapid-recovery-core.mib`, can then be read by an SNMP browser in a more user-friendly fashion than data it receives directly from the Core.


Use this procedure to download the SNMP MIB file from the Rapid Recovery Core.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More) and then click **Downloads**.
The Downloads page appears.
3. Scroll down to the *Other Files* pane.
4. To download the MIB file, click the **SNMP MIB file** download link.
The SNMP Configuration settings appear.
5. In the Opening quest-rapid-recovery-core.mib dialog box, do one of the following:
 - To open the log file, select **Open with**, then select an SNMP browser application for viewing the text-based MIB file, and finally click **OK**.
The MIB file opens in the selected application.
 - To save the file locally, select **Save File** and click **OK**.

Configuring vSphere settings





VMware vSphere is a suite of virtualization software, from which you can manage ESXi or vCenter Server virtual machines. If using vSphere, you no longer need to load the Rapid Recovery Agent software onto individual VMs to protect them. This is called the agentless protection feature, which applies only to virtual machines.



Use this procedure to configure vSphere settings for the Core.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **vSphere**.
 - Scroll down on the right side of the *Settings* page until you can see the vSphere heading.

3. Modify the vSphere settings as described in the following table.


Table 35: vSphere Core settings information

UI Element	UI Type	Description
Connection lifetime	Spin box	Establishes duration of time before a timeout for the connection with the ESXi server. Uses HH:MM:SS format.  NOTE: The default setting is 00:10:00 or ten minutes.
Maximum simultaneous consolidations	Text field	Sets the maximum number of simultaneous consolidations for protected virtual machines.  NOTE: The default setting is 0.
Maximum retries	Text field	Sets the maximum number of attempts for connection to a virtual disk or read and write operations before a timeout.  NOTE: The default setting is 10.
Allow parallel restore	Boolean (check box)	When this option is checked, enables parallel restore for an agentless virtual machine. When this option is cleared, this function is disabled.  NOTE: The default setting is No (cleared).

4. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.





Managing VMware proxy settings

The VMware proxy settings are intended for protected VMware ESXi machines that require Virtual Disk Development Kit (VDDK) APIs to access the VMware disk storage. Rapid Recovery addresses this access and possible timeouts associated with it by using a process called the VMware proxy. This service is automatically installed with the Rapid Recovery Core and only runs when it is needed. The Core Settings page lets you adjust the service timeout settings as you see appropriate.

1. On the icon bar of the Rapid Recovery Core Console, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **VMware Proxy**.
 - Scroll down on the right side of the *Settings* page until you can see the VMware Proxy heading.

2. Under VMware Proxy, modify the timeout settings described in the following table.

Table 36: VMware proxy settings information


UI Element	Description
Connection timeout	The maximum amount of time that should pass before the VMware proxy should stop trying to connect to the VMware disk storage, designated by hh:mm:ss.  NOTE: The default setting is 5 minutes (00:05:00).
Read/Write timeout	The maximum amount of time that should pass before the VMware proxy should stop trying to read or write to the VMware disk storage, designated by hh:mm:ss.  NOTE: The default setting is 5 minutes (00:05:00).
Start service timeout	The maximum amount of time that should pass before Rapid Recovery should stop trying to start the VMware proxy service, designated by hh:mm:ss.  NOTE: The default setting is 1 minute (00:01:00).
Stop service timeout	The maximum amount of time that should pass before Rapid Recovery should stop trying to stop the VMware proxy service, designated by hh:mm:ss.  NOTE: The default setting is 1 minute (00:01:00).

3. For each setting, when satisfied with your changes, click the check mark to save the change and exit edit mode, or click X to exit edit mode without saving.

Configuring vFoglight settings

If you provide vFoglight information in the Core settings, the Core Console displays a vFoglight URL for each protected VM on its *Summary* page. Clicking this URL opens information about the VM in vFoglight.



Complete the following steps to integrate your vFoglight server with your Rapid Recovery Core server.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **vFoglight**.
 - Scroll down on the right side of the *Settings* page until you can see the **vFoglight** heading. The vFoglight configuration settings are displayed.

3. Modify the vFoglight settings as described in the following table.

Table 37: vFoglight connection settings information

Text Box	Description
Use https	Enables or disables secure hypertext transfer protocol. Secure HTTPS is the default.
Host	Enter a host name or IP address for your vFoglight server. Consult your vFoglight server administrator for details.
Port	Specify the appropriate port. The default port is 32896. Consult your vFoglight server administrator for details.
Authentication token	Provide the proper authentication to let the Core communicate with your vFoglight server. Consult your vFoglight server administrator for details.


4. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

Configuring SAML settings

To integrate your SAML single sign-on authentication identity provider (IdP) server with your Rapid Recovery Core server, complete the following steps.

- NOTE:** Before you configure SAML settings in Rapid Recovery, you must enable SAML with a compatible IdP. For more information, see [Understanding SAML single sign-on](#).
- NOTE:** Rapid Recovery Core supports two types of authentication: Windows-based and SAML-based. The Core uses Windows-based authentication by default. After you enable SAML, the Core begins to use SAML-based authentication.
- CAUTION:** After you configure the SAML settings, you must restart the Rapid Recovery Core Service for the changes to take effect. Before you restart the Core Service, you must complete all of the steps in the configuration procedure. If you restart the Core service before completing the SAML configuration, the Core Console becomes inaccessible.

To configure SAML settings

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **SAML**.
 - Scroll down on the right side of the *Settings* page until you can see the **SAML** heading. The SAML configuration settings are displayed.
3. Next to **Enable configuration**, select **Yes**.
4. Copy or take note of the following information provided by Rapid Recovery Core:
 - Sign on URL
 - Reply URL (Assertion Consumer Service URL)
 - Logout URL
5. Go to your IdP.






6. On the Single Sign-on or SAML page, paste or enter the Rapid Recovery information in the corresponding areas.
7. On the same page of your IdP, copy or take note of the following information:
 - Metadata URL (Issuer URL)
 - Audience (Entity ID)
8. If you have token encryption enabled, upload the certificate file and provide the password.
9. Return to the Rapid Recovery Core Settings page.
10. In the SAML section, paste or enter the **Metadata URL** and **Entity ID** that you copied from the IdP.
 For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.
11. Modify the SAML settings as described in the following table. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.
 Find the following details on the SAML page of your IdP.

Table 38: SAML connection settings information

Text Box	Description
Enable configuration	Select Yes .
Metadata URL	Enter the metadata URL from your IdP. This URL provides an endpoint on the Rapid Recovery Core that Rapid Recovery uses to furnish keys and additional SAML endpoints to the IdP.
Entity ID	Enter the entity ID from the IdP. This URL serves as the identifier that represents the Rapid Recovery Core server, and should be the same as the Entity ID on the IdP.
Signature algorithm	If using a certificate, select the signature algorithm from the following options: <ul style="list-style-type: none"> • RSA-SHA1 • RSA-SHA256 • RSA-SHA384 • RSA-SHA512
Certificate validation	Select the validation you want to use from the following options: <ul style="list-style-type: none"> • None • ChainTrust • PeerOrChainTrust • PeerTrust
Enabled token encryption	Select whether to enable token encryption.  NOTE: Not all IdPs support token encryption.

12. If you enabled token encryption, the **Certificate** option appears.

13. Next to **Certificate**, click **Upload** and complete the following steps:
 - a. In the Upload certificate window, upload the certificate file by clicking **Choose File**.
 - b. Enter the password for the certificate.
 - c. Click **Continue**.
14. To run a validation of the current settings and confirm the login with the IdP, click **Check SAML**.
15. To return the settings to their original state, click **Reset**.
16. To apply the SAML settings, restart the Core service.

Protecting machines

This section describes how to protect, configure, and manage the protected machines in your Rapid Recovery environment.

Topics include:

- [About protecting machines with Rapid Recovery](#)
- [Understanding the Rapid Recovery Agent software installer](#)
- [Deploying Agent to multiple machines simultaneously from the Core Console](#)
- [Using the Deploy Agent Software Wizard to deploy to one or more machines](#)
- [Modifying deploy settings](#)
- [Understanding protection schedules](#)
- [Protecting a machine](#)
- [About protecting multiple machines](#)
- [Enabling application support](#)
- [Settings and functions for protected Exchange servers](#)
- [Settings and functions for protected SQL servers](#)

About protecting machines with Rapid Recovery

To protect your data using Rapid Recovery, you need to add the workstations, servers, desktop, and laptop machines you want to protect to your Rapid Recovery Core.

In the Rapid Recovery Core Console, using one of the Protect Machine wizards, you can identify the machines you want to protect. You can do the following:

- Protect a single machine using the Protect Machine wizard, which connects to the machine using network hostname or IP address. For more information about how to protect a single machine, see [Protecting a machine](#).
- Protect a network cluster using the Protect Cluster wizard, which connects to the cluster and its nodes using network hostname or IP address. For more information about how to protect a cluster, see [Protecting a cluster](#).
- Protect multiple machines simultaneously using the Protect Multiple Machines wizard. This wizard lets you connect to the machines associated with a Microsoft Active Directory server; machines on a vCenter or ESXi host; or to machines on a Hyper-V host or a Hyper-V cluster. You can also manually enter connection information (network hostname or IP address, username and password) for multiple machines. For more information about how to protect multiple machines, see [About protecting multiple machines](#).



NOTE: Quest recommends limiting the number of machines you protect simultaneously to 50 or fewer, to preclude experiencing resource constraints that may cause the protect operation to fail.

When identifying your protection requirements for a single machine in the wizard, you can specify which volumes to protect. When you protect multiple machines, all volumes are protected by default. (You can change this later on an individual machine basis).

When protecting a virtual machine on a vCenter/ESXi or Hyper-V host, you must define whether to protect the machine using the Rapid Snap for Virtual feature or by installing Rapid Recovery Agent. For more information, see [Factors when choosing agent-based or agentless protection](#).

The wizard also lets you define a customized schedule for protection (or re-use an existing schedule).

Using advanced options, you can add additional security measures by specifying or applying an encryption key to backups for the machines you want to protect.

Finally, if one does not already exist, you can define a repository using the wizard.

After installing the Agent software, each machine must be restarted after installation.

For more information on how to protect workstations and servers, see [Protecting a machine](#).

Factors when choosing agent-based or agentless protection

The Rapid Snap for Virtual feature of Rapid Recovery is supported on vCenter/ESXi or on Hyper-V hypervisors. This feature, also known as agentless protection, lets you protect VMs running on your protected hypervisor in your Core without installing the Rapid Recovery Agent software on each guest machine.

General recommendations

Rapid Snap for Virtual has nearly achieved parity with protection provided by installing the Rapid Recovery Agent software. As a general rule, Quest recommends using agentless protection on ESXi or Hyper-V virtual machines. If the Agent software is installed on ESXi or Hyper-V VMs, unless there is a compelling reason to explicitly protect your VM using Rapid Recovery Agent, Quest recommends removing the Agent software, and protecting your VMs agentlessly.

There are some advantages to protecting agentlessly, and some limitations. These are clearly described in the topic [Understanding Rapid Snap for Virtual](#).

Exceptions to the recommendation to use agentless protection are as follows:

- Gathering metadata for agentless machines is slower than for machines protected by the Rapid Recovery Agent software. If you experience performance issues related to metadata (specifically for agentlessly protected Exchange Server or SQL Server machines), Quest Data Protection Support may suggest installing the software-based Agent on specific application servers for troubleshooting purposes.
- If protecting only one or two VMs on a hypervisor with multiple sockets, you may consume fewer licenses by installing Agent directly on the VMs instead of the host.
- If you require features exclusive to Rapid Recovery Agent, install the Agent software on relevant VMs.

Some features are unique to protection by installing the Rapid Recovery Agent software. The following examples apply:

- Performing a SQL attachability check is a capability of the Rapid Recovery Agent software. If protecting your SQL Server machine agentlessly, you must perform SQL Attachability checks using an instance of SQL Server installed on the Core server. To perform this check, you must adjust your Core Attachability setting on the Core to Use SQL Server on the Core.
- Dynamic volumes protected agentlessly are protected at the disk level, not the volume level.

- Live Recovery is a feature of the Rapid Recovery Agent software. You cannot use this feature when restoring volumes protected using Rapid Snap for Virtual (nor for Linux machines or when restoring CSVs).

NOTE: Rapid Recovery supports Windows Server 2012 and 2012 R2 for agentless protection only.

If you require any of the features described in the previous list for a specific VM, Quest recommends installing Agent instead of protecting the VM agentlessly.

For more information, see the topic [Understanding Rapid Snap for Virtual](#).

Release 6.7 license consumption concepts

As described in the *Rapid Recovery Installation and Upgrade Guide* topic "Understanding Rapid Recovery licenses," Rapid Recovery 6.2 and later uses only two license pools: Capacity, and Enterprise. If licensing for your Core is set up to use a capacity-based pool, you cannot use another pool type.

NOTE: In the future, Quest may add license pools based on other units of measure. Capacity and Enterprise pools continue to be supported.

DL series backup appliances use back-end capacity-based licensing, and are not affected by license pool restrictions. Software-based Rapid Recovery environments using front-end capacity licensing likewise receive no license benefits from using agentless protection. Other benefits for using agentless protection are relevant even when Capacity license pools are in use.

If your Rapid Recovery release 6.2 or later environment uses an Enterprise license pool, then the following rules apply:

- Hyper-V or vCenter/ESXi hypervisor hosts protected with Rapid Recovery Agent consume one license from the pool *for each processor socket*. If your hypervisor host has six CPU sockets, it consumes 6 licenses from the Enterprise pool.
- Any other machine (physical or virtual) protected in your Core with Rapid Recovery Agent consumes one license from that pool. This is true even for application servers (such as Exchange Server, SQL Server, or Oracle Database 12c) with multiple CPU sockets.

Licensing benefits of using agentless protection

You can protect guest VMs on a vCenter/ESXi hypervisor host by running the Protect Multiple Machines Wizard. On the Connection page of this wizard, if you specify **Protect selected VMs agentlessly**, the guest VMs on that host are protected agentlessly. For those VMs, no licenses are consumed from your license pool. While Rapid Recovery Agent is not installed on the host, adding that host to your Core consumes one license for each CPU socket.

When you protect a Hyper-V Server, Rapid Recovery Agent is installed on the host. For each CPU socket on that hypervisor host, one license from your Enterprise pool is consumed. If you specify protecting the Hyper-V server agentlessly, guest VMs are protected agentlessly, and for those VMs, no licenses are consumed from your available license pool.

When you protect a Hyper-V cluster, Rapid Recovery Agent is installed on each node in the cluster. Only a single license is consumed from your license pool. The total number of CPU sockets in the cluster are consumed. If you specify protecting the Hyper-V cluster agentlessly, guest VMs are protected agentlessly, and for those VMs, no licenses are consumed from your available license pool on the cluster.

The chief licensing benefit to using Rapid Snap for Virtual is a reduction in consumption of licenses from your Enterprise license pool for the VMs you protect. If you specify agentless protection for an ESXi hypervisor host, or a Hyper-V server or cluster, all new VMs created on the host are automatically protected agentlessly, and do not consume licenses from your Enterprise license pool.

If some of the VMs on that hypervisor host previously had Rapid Recovery Agent installed, and your Core is running Rapid Recovery release 6.2 or later, you should do one of the following:

- Remove the Agent software and protect the VM agentlessly. No licenses from your pool are consumed.
- If you require the machine to be protected by Agent, and the host is added to the Core, associate the VM with its parent host. You get the benefit of Agent-based protection, and no license is consumed.
- Make no changes. The VM is protected using the APIs in Rapid Recovery Agent, and a single license is consumed.

Each virtual machine on a hypervisor added to your Core is protected agentlessly without consuming a license. To obtain this benefit, you must do the following:

The chief licensing benefit to using Rapid Snap for Virtual is a reduction in consumption of licenses from your Enterprise license pool for the VMs you protect. Each virtual machine on a hypervisor added to your Core is protected agentlessly without consuming a license. To obtain this benefit, you must do the following:

- **Protect VMs agentlessly.** You can explicitly protect VMs by using the Protect Multiple Machines wizard. When protecting a hypervisor host, you can also select the option to **Auto protect new virtual machines**, which implicitly protects new VMs when they are created.
- **Associate the guest VM with its protected hypervisor host.** If Rapid Recovery Agent is installed, its APIs (not those native to the hypervisor) are used to protect the VM. However, you can reduce licenses consumed by associating the VM with the host that has been added to the Core. This association is performed at the machine level for each virtual machine. The process of linking the guest VM with its parent hypervisor host is described in [step 3](#) of the procedure [Viewing and modifying protected machine settings](#).
- **Uninstall Agent.** Unless otherwise recommended, remove any copies of the Agent software from the virtual machine.

For a discussion of benefits and limitations regarding agentless protection, additional software recommended, minimum requirements for the host, and so on, see the topic [Understanding Rapid Snap for Virtual](#).

About protecting Linux machines with Rapid Recovery

The Rapid Recovery Agent software is compatible with multiple Linux-based operating systems (for details, see the system requirements defined in the *Rapid Recovery System Requirements Guide*). The Rapid Recovery Core is compatible only with Windows machines. While you can manage protected Linux machines from the Rapid Recovery Core Console, several procedures for Linux machines have steps that differ from their Windows counterparts. Additionally, you can perform some actions directly on a protected Linux machine by using the `local_mount` command line utility.

If you want to protect a single Linux machine, you can now use the Protect Machines Wizard. See the topic [Protecting a machine](#). To protect multiple Linux machines simultaneously using the wizard from the Core Console, see the topic [Protecting multiple machines manually](#).

To deploy or install the Agent software to a Linux machine from the Core Console, you must have the following:

- The user account must have SUDO privileges.
- The Linux machine you want to protect must have access to an SSH server.

If a Linux machine you want to protect does not meet these prerequisites, consult with a Linux administrator. Comply with these requirements and then you can complete the relevant wizard to deploy and install the Agent software.

About protecting Oracle database servers

Rapid Recovery includes application support of Agent-based protection of Oracle 12c and Oracle 18c relational database management systems (RDBMS). You can protect an Oracle database server and all of its databases, and perform related tasks.

In this release, the following restrictions apply:

- Oracle 12c and Oracle 18c are the only tested and supported versions for protection on the Rapid Recovery Core. Use any other Oracle versions at your own risk.
- Protected Oracle database servers must run 64-bit versions of Windows Server 2012 R2 or Windows Server 2016.
- You must install the Rapid Recovery Agent software (release 6.2 or later) on your Oracle server. Agentless protection is planned for a future release.
- Protection of Oracle 12c and Oracle 18c databases is limited to using Volume Snapshot Service (VSS) in the ARCHIVELOG mode.

i NOTE: Support for NOARCHIVELOG (the default Oracle database log mode) is planned for a future release. Oracle databases with ARCHIVELOG mode enabled must also be set to archive all online redo logs using the ARCH (archive) process. A database administrator (DBA) can change the mode and set archiving of redo logs using Oracle SQL*Plus or Oracle Enterprise Manager. For more information about enabling ARCHIVELOG mode and archiving, see Oracle documentation or consult a qualified Oracle 12c or 18c DBA.

To fully protect Oracle servers, perform the following tasks:

- Install the Rapid Recovery Agent software (release 6.2 or later) on your Oracle server and begin protection. Use the Protect Machine Wizard to locate the Oracle server on your network, deploy the Agent software, and establish a protection schedule. For more information, see [About protecting machines with Rapid Recovery](#).
- Enter credentials for each database in the Rapid Recovery Core Console. The Core securely caches your credentials and lets you access the metadata from the UI. Before you enter credentials, you cannot view details for databases on your protected Oracle server. For more information, see [Entering or editing credentials for Oracle databases](#).
- Enable ARCHIVELOG mode for the protected Oracle server, and verify the Oracle VSS writer, from the Core Console. For more information, see [Enabling archive log mode and adding VSS writer for protected Oracle databases](#).
- Review log truncation deletion policies offered in Rapid Recovery Core and implement a log truncation regime. When an Oracle database is set to ARCHIVELOG mode, the logs accumulate quickly and consume substantial disk space. You can truncate Oracle logs manually on demand, or configure automatic log truncation using the "Log truncation for Oracle" nightly job. Both methods offer the same [Deletion policies for Oracle log truncation](#). To avoid eventual errors and snapshot failures when storage for the logs is full, implement a consistent approach for managing Oracle logs. For more information, see [About truncating Oracle logs](#).

You can also truncate Oracle database logs manually on demand. For more information about this procedure, see [Manually truncating Oracle database logs](#).

Once you install the Agent, protect the machine in your Core, and configure settings properly, you can do the following:

- **View metadata.** From the protected machine Summary page, you can view metadata about each database on your Oracle server, including connection and status for each log file, control file and data file.
- **Check database integrity.** You can perform integrity checks from the Core Console using the DBVERIFY utility.
- **Truncate archive logs,** using one of three deletion policies.
- **Restore databases.** Restore entire volumes, or volumes that contain selected databases. Once you enable Archive log mode, snapshots of the Oracle database are crash-consistent from the point of view of the Oracle service.
- **Perform virtual export.** You can make a one-time export, or set up a virtual standby VM that continually updates a VM with new information as backups on your protected database are captured. If you boot up a VM of an Oracle database, you may have to manually start the database services, and manually disable backup mode for database data files.

Entering or editing credentials for Oracle databases

Before performing this procedure, you must first add an Oracle database server to protection on your Core.

To enter or edit Oracle database credentials:

- The Windows user account of the Rapid Recovery user performing this procedure must have SYSDBA privileges on the protected database server.
- The database must be accessible to the Rapid Recovery Core server, and a connection must be successfully established.


i **NOTE:** SYSDBA is an Oracle systems database administrative privilege required to perform high-level administrative operations. These functions include creating, starting up, shutting down, backing up, or recovering an Oracle database.

After protecting your Oracle database server, you cannot access database metadata or view database details until you enter credentials for each database. This one-time step caches your database credentials securely and provides the Core Console with access to status information about all protected transaction log files, control files, and data files that comprise your Oracle databases.

For example, on the *Summary* page for the protected Oracle machine, before entering credentials, you cannot expand details about any of the protected databases in the *Oracle Server Information* pane.


i **NOTE:** This is a one-time step required for each new protected Oracle database.

Complete the steps in this procedure to provide the Rapid Recovery Core Console with access to the required metadata for your protected Oracle databases.

1. Navigate to the protected Oracle machine in the Rapid Recovery Core Console. The *Summary* page displays for the protected machine.
2. On the *Summary* page, scroll down to the Oracle Server Information pane.
3. For the first database in the table, click the  (More options) drop-down list and select **Edit Credentials**. The *Edit Instance Credentials* dialog box displays.

4. Two connection types are supported: basic, and Transparent Network Substrate (TNS, a proprietary Oracle networking technology). Do one of the following:

- To connect using a basic connection, enter the information in the following table:


Option	Description
Connection Type	Basic
Host Name	Enter the host name or IP address.
Port	Enter the appropriate port. The default port open for this purpose is 1521.
SID or Service Name	Select the appropriate connection method. You can use one of the following: <ul style="list-style-type: none">• SID. The Oracle System Identifier (SID) is a unique ID that uniquely identifies your database instance.• Service Name. The service name is the TNS alias used to remotely connect to your database. <div> NOTE: The service name can be found in the TNSNAMES.ORA file.</div>
Service Name	The service name is the TNS alias that you give when you remotely connect to your database and this Service

- To connect using TNS, enter the information in the following table:


Option	Description
Connection Type	TNS
Network Alias	Select this drop-down menu to view database aliases available to the network, and select the appropriate alias.

5. Two credential types are supported: Oracle, and Operating System. Do one of the following:

- To connect with Oracle credentials, enter the user name and password for the Oracle database in the relevant text fields.

 **NOTE:** Your Windows user account must have SYSDBA privileges.

- To connect using cached credentials from the operating system, select **Operating System**.

 **NOTE:** Your Windows user account must be a member of the ORA_DBA local group, which ensures the user has SYSDBA privileges.

6. To verify credentials, click **Verify**.

A dialog box displays, indicating if the test connection was successful.

7. Do one of the following:

- If the verification succeeded, click **OK** to close the message dialog box.
- If not successful, close the dialog box and revise the instance credentials until connection is verified. Consult your system administrator if you have questions about credentials.

8. In the *Edit Instance Credentials* dialog box, after successful verification, click **OK**.
The dialog box closes, and Rapid Recovery Core Console immediately applies and caches the credentials. Very soon afterward, metadata is available in the Core Console, and the status indicator for the selected database displays a green (online) status.
9. Repeat [step 3](#) through [step 8](#) for each database listed in the Oracle Server Information pane.

After entering and caching credentials for all databases on this protected Oracle machine, perform the procedures described in the topic [Enabling archive log mode and adding VSS writer for protected Oracle databases](#).

Enabling archive log mode and adding VSS writer for protected Oracle databases

Before performing this procedure, you must first add an Oracle database server to protection on your Core, and enter credentials for each database into the Core Console.

Database applications require the presence of a combination of specific files (such as configuration, log, and control files), each set to a specific state, to be able application-consistent. In Rapid Recovery Core release 6.2 and later, Oracle databases require archive log mode to be enabled. Until you perform this procedure, your snapshots of the database server will be crash-consistent, but not application-consistent.

Oracle VSS writer captures snapshots using Volume Snapshot Service. This writer must be enabled.

i | **NOTE:** These are one-time steps required for each new protected Oracle database.

Complete the steps in this procedure to verify if archive log mode is enabled; to enable that mode if required; and to add VSS writer to your Core.

1. Navigate to the protected Oracle machine in the Rapid Recovery Core Console.
The *Summary* page displays for the protected machine.
2. On the *Summary* page, scroll down to the *Oracle Server Information* pane.
3. In the *Oracle Server Information* pane, if you see a warning notification that archive log mode is disabled, click **Enable archive log mode for these databases** and then click to confirm restart of the database instances.

i | **NOTE:** Enabling archive mode causes the relevant database instances to restart. This could take a few minutes.

4. In the Oracle Server Information pane, if you see a warning notification that the Oracle VSS writer is excluded from snapshots, click **Include Oracle VSS Writer** and then click to confirm.
The warning message closes, and the VSS writer is added to your Core.
5. Optionally, you can track the progress towards enabling archive log mode or adding a VSS writer. For more information, see [Viewing tasks](#).

About truncating Oracle logs

Currently, protection of Oracle 12c databases in Rapid Recovery Core release 6.2 and later is limited to using Volume Snapshot Service (VSS) in the ARCHIVELOG mode. In this mode, the database is configured to archive all online redo logs using the ARCH (archive) process. The Oracle VSS writer calls the Oracle service to archive the existing re-do logs to Oracle's `fast_recovery_area` folder, and creates a new archive log file.

ARCHIVELOG mode causes a substantial number of log files to accumulate on the database server, using valuable storage space. Each time a recovery point of the Oracle server is captured, the new Oracle logs generated since the

last snapshot are included. This renders the local copies of the logs superfluous. For this reason, Rapid Recovery Core includes several methods for truncating Oracle logs.

1. **Enable Oracle log truncation on the Core as a nightly job.** Rapid Recovery Core includes a nightly job setting called **Log truncation for Oracle**. This nightly job is disabled by default. If you enable this nightly job, Oracle ARCHIVELOGS are truncated according to the deletion policy specified in the settings for this job.
2. **Enable and customize Oracle log truncation as a nightly job.** For each specific Oracle server protected on your Core, you can customize the Oracle log truncation nightly job configuration. As with the Core nightly job, you can select from any of the three deletion policies, described in the following section.
3. **Manually truncate logs on demand.** Regardless of whether you set a log truncation schedule using nightly jobs, you can manually truncate Oracle ARCHIVELOGS on demand at any time.

Truncation of Oracle logs, by nightly job or on demand manually, occurs without requiring a transfer job.

When you specify truncation of Oracle logs by nightly jobs, truncation occurs according to the deletion policy at the time nightly jobs are scheduled to occur. Conversely, when you elect to manually truncate Oracle logs, a log truncation job queues. If the system is not busy, the job executes immediately and the logs are truncated.

Deletion policies for Oracle log truncation

When configuring log truncation for Oracle database servers, you can choose from one of three deletion policies:

- The **Automatic** deletion policy truncates all Oracle ARCHIVELOGS stored in the `fast_recovery_area` folder one time daily when nightly jobs run.
- The **Keep newest** deletion policy lets you specify the duration of time (n days, weeks, months, or years) to retain the Oracle logs before truncating. When the time period expires, log files past the threshold are then truncated one time daily when nightly jobs run.
- The **Keep specified number** deletion policy lets you specify a specific number of log files to retain. After that threshold is reached, newer logs are retained, and the older logs are then truncated one time daily when nightly jobs run.

If you need a truncated ARCHIVELOG file, you can obtain it from the latest recovery point captured prior to truncating the logs.


For more information about truncating jobs as a nightly job, see the topic [Understanding nightly jobs](#).

For more information about manually truncating Oracle logs on demand, see [Manually truncating Oracle database logs](#).

Manually truncating Oracle database logs


This procedure is only appropriate for Oracle database servers protected on your Core.

When protecting an Oracle database in ARCHIVELOG mode, substantial log files accumulate on the local database server. To support the protection of Oracle databases in Rapid Recovery, you can enable a nightly job to truncate Oracle logs.

 **NOTE:** Log truncation for Oracle is disabled by default.

For any protected Oracle server, you can also manually truncate Oracle database logs on demand at any time, as described in this topic. Truncating will delete the logs from the local server. In each recovery point saved to your repository, the database logs persist to reflect the state of the database at the time the backup snapshot was captured.

Complete the steps in this procedure to manually truncate Oracle log files.

1. Navigate to the protected Oracle machine in the Rapid Recovery Core Console. The Summary page displays for the protected machine.
2. On the Summary page, scroll down to the Oracle Server Information pane.
3. In the table row representing the Oracle database instance for which you want to truncate logs, click the  (More options) drop-down list and select **Force Log Truncation**. The Force Log Truncation dialog box displays.
4. If you want to delete all of the Oracle logs from the local database server, from the **Deletion policy** drop-down menu, select **Automatic**, and then click **Force**.
A log truncation job queues. If the system is not busy, the job executes immediately and the logs are truncated.
5. If you want to delete all of the locally stored logs from the Oracle server (copies of which are stored in the most recent recovery point), do the following:
 - a. From the **Deletion policy** drop-down menu, select **Keep newest**.
 - b. In the **Keep logs for** text field, enter a number, and then from the time period drop-down menu, select the relevant a period of time (**days**, **weeks**, **months**, or **years**).
 - c. Click **Force**.
A log truncation job queues. If the system is not busy, the job executes immediately and the logs are truncated.
6. If you want to keep a specific number of Oracle log files, and truncate the remaining logs, then do the following:
 - a. From the **Deletion policy** drop-down menu, select **Keep specified number**.
 - b. In the **Number of archive files** text field, enter a number representing the amount of the newest database logs to retain.
 - c. Click **Force**.
A log truncation job queues. If the system is not busy, the job executes immediately and the logs are truncated.
7. If you want to truncate logs for other database instances on this server, repeat [step 3](#) through [step 6](#) for each relevant database listed in the Oracle Server Information pane.

About managing Exchange and SQL servers in Rapid Recovery Core

Options specific to Exchange Server and SQL Server appear in the Rapid Recovery Core Console only when an instance of the software and related files are detected on protected servers. In those cases, additional options are available when you select the protected machine in the Core Console.

For example, if you select a protected Exchange server in the left navigation menu, then the menu options that appear for that protected machine include an **Exchange** drop-down menu option.

If you select a protected SQL server in the left navigation menu, then the menu options that appear for that protected machine include a **SQL** drop-down menu.

While these options may work differently, there is some commonality. Functions you can accomplish for protected Exchange and SQL servers (and for no other protected machines) include:

- **Forcing server log truncation.** Both SQL servers and Exchange servers include server logs. The process of truncating SQL logs identifies available space on the server. When you truncate logs for an Exchange server, in addition to identifying the available space, the process frees up more space on the server.
- **Setting credentials for the relevant server.** Exchange servers allow you to set credentials for the protected machine on the Summary page for the protected server. SQL servers allow you to set credentials for a single protected SQL Server machine, or to set default credentials for all protected SQL servers.
- **Viewing status for checks on recovery points from Exchange Server or SQL Server.** Recovery points captured from a protected SQL or Exchange server have corresponding color status indicators. These colors indicate the success or failure of various checks relevant for SQL servers or Exchange servers.

This section includes the following topics specific to managing protected machines that use Exchange Server or SQL Server:

- [Understanding recovery point status indicators](#)
- [Settings and functions for protected Exchange servers](#)
- [Settings and functions for protected SQL servers](#)

About protecting server clusters

In Rapid Recovery, server cluster protection is associated with the Rapid Recovery protected machines installed on individual cluster nodes (that is, individual machines in the cluster) and the Rapid Recovery Core, which protects those machines, all as if they were one composite machine.

You can easily configure a Rapid Recovery Core to protect and manage a cluster. In the Core Console, a cluster is organized as a separate entity, which acts as a container that includes the related nodes. For example, in the left navigation area, under the *Protected Machines* menu, protected clusters are listed. Directly below each cluster, the associated individual nodes or agent machines appear. Each of these is a protected machine on which the Rapid Recovery Agent software is installed. If you click on the cluster, the *Summary* page for the cluster appears in the Core Console.

At the Core and cluster levels, you can view information about the cluster, such as the list of related nodes and shared volumes. When showing information for a cluster in the Core Console, you can click **Protected Nodes** in the top navigation menu to view a summary table of individual nodes in the cluster. From that summary table, for each node, you can perform functions such as forcing a snapshot; performing a one-time export or setting up virtual standby; mounting or viewing recovery points; restoring from a recovery point; converting the cluster node to its own protected machine; or removing the node from protection. If the node is an Exchange or SQL Server, you will also see the option to truncate logs.

At the cluster level, you can also view corresponding Exchange and SQL cluster metadata for the nodes in the cluster. You can specify settings for the entire cluster and the shared volumes in that cluster.

If you click on any node in the cluster from the left navigation menu, the information displayed in the Core Console is specific to that node of the cluster. Here you can view information specific to that node, or configure settings just for that node.

For more information about CSV clusters, see topics "Supported applications and cluster types" and "Support for Cluster-Shared Volumes" in the *Rapid Recovery System Requirements Guide*.

Understanding Rapid Snap for Virtual



The Rapid Snap for Virtual feature of Rapid Recovery is also known as agentless protection, because you can protect virtual machines (VMs) in your Core without installing the Rapid Recovery Agent on every VM.

CAUTION: Quest recommends that you limit initiating agentless protection to no more than 200 VMs at once. Do not select more than 200 VMs when using the Protect Multiple Machines Wizard. Attempting to start protection of more than 200 VMs results in slow UI performance. There is no limit to how many VMs a Core can agentlessly protect over time. For example, you could protect 200 VMs today and another 200 VMs tomorrow.

Whether adding protection for a single machine agentlessly or multiple machines simultaneously, the tool to access agentless protection is the Protect Multiple Machines wizard.

Protecting VMware vCenter/ESXi VMs

Rapid Recovery lets you protect vCenter/ESXi VMs without installing the Rapid Recovery Agent on the VM or ESXi host, resulting in fully agentless protection. Rapid Recovery uses the ESXi client and the application program interface (API) native to VMware to detect and protect selected VMs on a single host. The Rapid Recovery Core then communicates with the virtual machine disk (VMDK) to determine the necessary details of the protected volumes. Because Rapid Recovery creates recovery points based on volumes, not VMDKs, each volume can be separately mounted, restored, and exported.

Rapid Recovery does not technically protect a VMware vCenter/ESXi hypervisor host. When you select a vCenter/ESXi host using the Protect Multiple Machines wizard, you add the host as a parent entity on your Rapid Recovery Core. No data files or data from the actual host are included in snapshots on the Core. However, the VM guests on the host can be protected. The protected VMs are represented in the Core GUI as children under the parent host. If you select agentless protection for ESXi VMs, the icon for the protected VM  appears differently than the icon on an ESXi VM protected using Rapid Recovery Agent .

NOTE: Quest recommends that VMware Tools be installed on virtual machines (VMs) you want to protect on vSphere or ESXi hosts. When VMware Tools are installed on a VM using a Windows operating system (OS), the backups that the Rapid Recovery Core captures use Microsoft Volume Shadow Copy Services (VSS). This provides the capacity for application-consistent backups. For information on the behavior of agentless VMs with or without VMware Tools, see [Benefits of installing hypervisor tools for agentless protection](#) and also [Understanding crash-consistent and application-consistent backups](#).

When VMware Tools are installed, agentless protection uses VMware Changed Block Tracking (CBT) to reduce the time needed for incremental snapshots. CBT determines which blocks changed in the VMDK file, letting Rapid Recovery back up only the portions of the disk that have changed since the last snapshot. This backup method often results in shorter backup operations and reduced resource consumption on network and storage elements.

There are multiple benefits to using agentless protection. Some of the most useful attributes include the following characteristics:

- No additional software is required on the host machine.
- Agentless protection lets you opt to automatically protect new VMs added to the ESXi host or Hyper-V host.
- A restart is not required during the protection process.
- Credentials are not required for each individual VM.
- Agentless protection lets you protect a VM even if it is powered off.
- Agentless protection lets you restore to disks.
- Agentless protection supports all guest operating systems.
- If you associate a guest VM with its protected parent hypervisor host, an Enterprise license is not consumed when protecting the VM in your Core.
- You can optionally protect and collect metadata for SQL Server and Exchange.

- The Core (not the Agent) can perform attachability checks, log truncation, and mountability checks on recovery points captured from protected Oracle, SQL and Exchange servers.
- Agentless protection lets you export dynamic disks or volumes.

i **NOTE:** If dynamic volumes are complex (striped, mirrored, spanned, or RAID), they export as disk images and parse into volumes after the export operation completes on the exported VM.

While there are many reasons to use agentless protection for ESXi VMs, opt for the protection method that best suits your environments and business needs. Along with the previously mentioned benefits, there are also the following considerations to keep in mind when choosing agentless protection:

- Agentless protection of Oracle database servers does not gather Oracle-related metadata. While files and the operating system are backed up in snapshots, no Oracle-related features are supported. For Oracle database application support, protect your server using Rapid Recovery Agent.
- Agentless protection does not support protection of dynamic volumes (for example, spanned, striped, mirrored, or RAID volumes) at the volume level. It protects them at the disk level.
- Agentless protection does not support Live Recovery. For more information about this feature, see [Understanding Live Recovery](#).
- During the restore process of a single volume to the protected VM, the VM is automatically restarted.
- Agentless protection does not display the actual amount of space used on a VM if the virtual disk type is thick provision eager zeroed.

If you choose to use agentless protection for your ESXi VMs, the host must meet the following minimum requirements for agentless protection to be successful.

- The host machine must be running ESXi version 5.0.0 build 623860 or later.
- The host machine must meet the minimum system requirements stated in the *Rapid Recovery System Requirements Guide*.
- For volume-level protection, VMDKs must include either Master Boot Record (MBR) partition tables or GUID partition tables (GPTs). VMDKs without these partition tables are protected as whole disks rather than as individual volumes.
- Each VMware virtual machine must have VMware Tools installed to ensure snapshot consistency.

Protecting VMs on Hyper-V servers and clusters

To protect Hyper-V VMs agentlessly, you do not need to install the Rapid Recovery Agent on each VM. You need only install it on the host machine or cluster nodes. The Agent protects the virtual hard disk on the host and converts any changes to the hard disk files to a volume image or disk image, depending on the file system. A new driver provides file-level support for VMs on hosts and on cluster-shared volumes (CSVs).

Agentless support for Hyper-V is determined by the operating system on the host. A complete list of operating systems, and Rapid Recovery components supported for each, is maintained in the *Rapid Recovery System Requirements Guide*. For more information, see the topic "Rapid Recovery release 6.7 operating system installation and compatibility matrix" in that document.

Quest recommends that Hyper-V Integration Services be installed on virtual machines (VMs) you want to protect on Hyper-V hosts. When Hyper-V Integration Services are installed on a VM using a Windows OS, the backups that the Rapid Recovery Core captures use Microsoft VSS. This provides the capacity for application-consistent backups. For information on the behavior of agentless VMs with or without Hyper-V Integration Services, see [Benefits of](#)

installing hypervisor tools for agentless protection and also [Understanding crash-consistent and application-consistent backups](#).

i | **NOTE:** Rapid Recovery supports the VHDx disk file format. It does not support the VHD format.

For protecting VMs on a CSV, the Rapid Recovery Agent and driver must be installed on each cluster node using the auto deployment feature in the Protect Multiple Machines Wizard. From the nodes, the Agent can protect all VMs operating on CSVs by creating two types of changes for every file. The first type of change is saved only before or after a snapshot or clean system restart. The second type of change resides on the disk, which makes an incremental snapshot available even if there is a power failure or dirty shutdown. The Agent installed on the node merges all of the changes into one before transferring the data.

When a host or node is running, Rapid Recovery creates a backup. If the host is not running, no backup can be created; however, if one of the nodes is not running, then Rapid Recovery can continue taking snapshots of the VMs on the cluster.

i | **NOTE:** For best performance, it is recommended that the maximum concurrent transfers for the Hyper-V host or node be set to 1, which is the default setting.

Agentless Hyper-V protection has many of the same capabilities as traditional protection where the Agent is installed on every VM, including:

- Archiving
- Recovery point integrity checks
- Mounting recovery points
- Auto discovery of new VMs (unique to agentless protection)
- Protecting SQL and Exchange servers and collecting their metadata
- Performing Exchange mountability checks
- Performing SQL attachability checks
- Replication
- Restoring VMs, including restoring to CSVs, or to CIFS shared folders
- Restoring files in a guest VHDX format
- Rollup
- Virtual export to Hyper-V VMs and other hypervisors, including ESXi, VMware Workstation, and VirtualBox


However, there are limitations to consider when choosing agentless Hyper-V protection. Capabilities that are not performed include:


- Live Recovery
- Restoring VMs on CIFS using VHD format
- Restoring files in a guest VHD (.vhd) format
- Restoring files in a guest VHD Set (.vhds) format

i | **NOTE:** For an application-consistent snapshot, you must have the SCSI Controller installed on each VM. Without this controller, the result is always a crash-consistent snapshot.

Application support

Rapid Snap for Virtual lets you enable agentless protection of SQL Server and Exchange applications running on Hyper-V and ESXi VMs. This optional capability is available for VMs running Windows operating systems.

 **NOTE:** Application support does not apply to applications installed on Linux VMs.

After you enable application support, application metadata displays on the Summary page for the VM, as does an icon  beside the VM name on the Machines page. If there is an error preventing healthy application support, the icon changes from green to red.

Before you opt to agentlessly protect an SQL Server or Exchange Server, take note of the following considerations:

- To protect the application, the VM must be powered on. The Core does not retrieve metadata from machines that are off.
- VMware Tools or Hyper-V Integration Services utilities must be installed on VMs you want to protect.
- On a VM that has both Exchange Server and SQL Server installed, there is no ability to truncate the logs separately. If both applications are installed, then the logs truncate together.
- SQL Server attachability checks are available only on the Core and cannot be conducted on the protected machine.
- To run log truncation on an ESXi VM, the host must use ESXi version 6.5 or later.

Benefits of installing hypervisor tools for agentless protection

When protecting virtual machines (VMs) without using the Rapid Recovery Agent Agent, Quest recommends installing VMware Tools on protected VMs on vSphere or ESXi hosts. In the same way, Quest recommends installing Hyper-V Integration Services on VMs you want to protect on Hyper-V hosts.

Installing these native hypervisor utilities lets Rapid Recovery take full advantage of Microsoft Volume Shadow Copy Services (VSS) functionality.

When these utilities are installed on VMs running Windows operating systems, the backups that the Rapid Recovery Core captures can also use VSS. When these tools are not installed, Rapid Recovery still collects snapshots, but only in a crash-consistent state. For more information, see [Understanding crash-consistent and application-consistent backups](#).

The following conditions apply based on whether VMware Tools or Hyper-V Integration Services are installed and on the powered-on state of the VM:

Table 39: Backup type conditions for VMs

Hypervisor Tool	VM Powered On	Backup Type
Not installed	Yes	Crash-consistent
Not installed	No (dirty shut-down)	Crash-consistent
Not installed	No (clean shut-down)	Application-consistent
Installed	Yes	Application-consistent
Installed	No (dirty shut-down)	Crash-consistent
Installed	No (clean shut-down)	Application-consistent


Understanding crash-consistent and application-consistent backups

When protecting virtual machines agentlessly using the Rapid Snap for Virtual feature, the data in the backup snapshots you capture can be in one of two states:

- **Crash-consistent.** At minimum, all agentless backups captured by the Rapid Recovery Core are crash-consistent. The backup is a snapshot in time of all the data and operating system files on each protected volume, at the time those files were captured. If you restore from a crash-consistent recovery point, the VM OS starts and can read and understand the file system, and all files in it.
If you recover a transactional application from a crash-consistent state, the database returns to the last valid state. That most recent valid state may be from the time of the crash, or it may be from earlier than the crash. If it is from earlier, then the database must roll forward some work to make the data files match the information in the logs. This process takes some time when you first open the database, which causes a delay when starting up the machine.
- **Application-consistent.** Application-consistent backups use Microsoft's Volume Shadow Copy Service (VSS) to ensure the consistency of application data when a shadow copy is created. Using VSS writers, pending input/output operations are completed and log files committed prior to snapshots being captured. As a result, if you restore from an application-consistent recovery point, the VM OS starts and can read and understand the file system. Additionally, files for transactional applications such as SQL Server or Exchange are in a consistent state. For example, SQL Server logs match the data files, and the database opens quickly without needing any repairs.

Support for Hyper-V guest cluster

Rapid Recovery can protect shared virtual hard disks in the VHDX format on a Hyper-V host by using agentless protection.

 **NOTE:** Rapid Recovery does not support protection of shared VHD sets.

A shared VHDX is a virtual hard disk on a Hyper-V host that is shared among a cluster of guests. When the Agent is installed on the Hyper-V host for agentless — or host-based — protection, protects the shared VHDXs at the host level. can also protect a shared VHDX at the guest level when the Agent is installed on the guest.

When you select and protect the virtual machines (VMs) available on the host, protects all virtual hard disks associated with those VMs, including any shared VHDXs. After protection is complete, you can edit the protection schedule and remove or protect shared VHDXs at the host level under Volumes on the Summary page for the host. All recovery points of these shared virtual hard disks are displayed on the Recovery Points tab for the host, not for the individual VMs.

When using agentless protection, automatically detects new VHDXs and offers the option to protect them. Unprotected disks are displayed in a separate Shared Disks section. Metadata is not gathered for these disks until after they are protected.


On the Summary page for a guest virtual machine, shared VHDXs are listed under Shared Volumes. A clickable icon next to the volume name opens a dialog that displays the VMs that are connected to that disk. If a connected VM is not protected, then the option to protect that VM is available.

Understanding the Rapid Recovery Agent software installer

You can download installers from the Rapid Recovery Core. From the *Downloads* page on the Core Console, you can choose to download the Agent Installer, the Local Mount Utility (LMU), or an SNMP MIB file. For more information about SNMP, see [Understanding SNMP settings](#).




NOTE: For access to the Agent Installer, see [Downloading the Rapid Recovery Agent Installer](#). For more information about deploying the Agent Installer, see the *Rapid Recovery Installation and Upgrade Guide*.

The Agent installer is used to install the Rapid Recovery Agent application on machines that are intended to be protected by the Rapid Recovery Core. If you determine that you have a machine that requires the Agent Installer, you can download the web installer from the *Downloads* page of the Rapid Recovery Core Console.

NOTE: You can download the Core, Agent, and other Rapid Recovery software from the QorePortal. Log into the QorePortal at <https://qoreportal.quest.com/> and click **Settings**, and then click  **Downloads**.

Downloading the Rapid Recovery Agent Installer

Download the Rapid Recovery Agent Installer and deploy it to any machine that you want to protect on the Rapid Recovery Core. Complete the steps in this procedure to download the installer.

1. To download the Agent installer directly from the machine you want to protect, do the following:
 - a. In a web browser, open the QorePortal at <https://qoreportal.quest.com/>.
 - b. Click **Settings**, and then click  **Downloads**.
 - c. Under *Applications*, locate **Agent Software**. If you want a full Windows 64-bit installer, click **Full Installer**. To download a smaller package that streams the latest version of the installer over the web, click **Web Installer**.
The installer file, for example `Agent-X64-6.3.xxxx.exe`, saves to the downloads destination folder.
2. To download the web installer from the Core, on the Core Console icon bar, click the  More icon and then select  **Downloads**.
3. On the Downloads page, from the **Agent** pane, click **Download web installer**.
4. If prompted, from the Opening Agent-Web.exe dialog box, click **Save File**.
The installer file, for example `Agent-X64-6.3.xxxx.exe`, saves to the downloads destination folder.
5. Move the installer to the appropriate machine and install the Rapid Recovery Agent software.
For more information about installing the Rapid Recovery Agent software, see the *Rapid Recovery Installation and Upgrade Guide*.

Deploying Agent to multiple machines simultaneously from the Core Console

You can deploy the Rapid Recovery Agent software simultaneously to multiple Windows machines. The machines can be part of an Active Directory domain, a vCenter or ESXi virtual host, or a Hyper-V virtual host; or they can be machines already protected by the local Rapid Recovery Core, as in the case of a Rapid Recovery Agent software upgrade. You also have the option to manually deploy the software to machines that are not necessarily associated with a specific domain or host.

You can also manually deploy the Rapid Recovery Agent software to one or more Linux machines from the Core Console.

CAUTION: If AppAssure Agent was previously installed on a Linux machine, then before installing Rapid Recovery Agent, remove the AppAssure Agent from the machine using a shell script. For information about removing the Agent from a Linux machine, see the topic "Uninstalling the AppAssure Agent software from a Linux machine" in the *Rapid Recovery Installation and Upgrade Guide*. To successfully deploy the Agent software to Linux machines, see the prerequisites in the topic "About installing the Agent software on Linux machines" in the same document.

Deploying the Rapid Recovery Agent software does not protect the machines automatically. After deploying, you must then select the **Protect Multiple Machines** option from the button bar of the Core Console.

NOTE: The feature in which you *deploy* to multiple machines simultaneously was previously referred to as *bulk deploy*. The feature in which you *protect* multiple machines simultaneously was previously referred to as *bulk protect*.

To deploy and protect multiple machines simultaneously, perform the following tasks:

- Deploy Rapid Recovery Agent to multiple machines. See [Deploying Agent to multiple machines simultaneously from the Core Console](#).
 - Monitor the deployment. See [Verifying the deployment to multiple machines](#).
 - Protect multiple machines. See [About protecting multiple machines](#).
- NOTE:** If you selected the *Protect Machine After Install* option during deployment, skip this task.
- Monitor the activity of the bulk protection. See [Monitoring the protection of multiple machines](#).

Using the Deploy Agent Software Wizard to deploy to one or more machines

You can simplify the task of deploying the Rapid Recovery Agent software to one or more machines by using the Deploy Agent Software Wizard.

If deploying to Linux machines, this method is appropriate.

NOTE: In the past, this feature was called "*bulk deploy*."

When you use the Deploy Agent Software Wizard, Rapid Recovery Core can:

1. Detect Windows machines on an Active Directory domain, and push the Agent software to the machines you select.
2. Connect to a VMware vCenter or ESXi host, detect the guests, and push the Agent software to the machines you select.
3. Connect to a local Rapid Recovery Core and deploy the current (newer) Agent software to Windows machines that Core already protects. (For Linux machines, use the option to deploy manually.)
4. Connect to a Hyper-V server or cluster, detect the guests, and push the Agent software to the machines you select.
5. Manually specify Linux or Windows machines, using IP addresses and credentials, and push the Agent software to the machines you select.

From within the Core Console, you can complete any of the following tasks:

- [Deploying to machines on an Active Directory domain](#)
- [Deploying to machines on a VMware vCenter/ESXi virtual host](#)
- [Deploying an upgrade of the Rapid Recovery Agent to protected machines](#)
- [Deploying to machines manually](#)



i **NOTE:** Quest recommends limiting the number of machines to which you deploy simultaneously to 50 or fewer, to preclude experiencing resource constraints that may cause the deploy operation to fail.

i **NOTE:** The target machines must have internet access to download and install bits, because Rapid Recovery uses the web version of the Rapid Recovery Agent Installer to deploy the installation components. If internet access is unavailable, use the Core Console to download the installer to a storage medium such as a USB drive. Then, physically install the software on the machines that you want to protect. For more information, see [Downloading the Rapid Recovery Agent Installer](#).

Deploying to machines on an Active Directory domain

Before you begin this procedure, have the domain information and login credentials for the Active Directory server on hand.

Use this procedure to simultaneously deploy the Rapid Recovery Agent software to one or more machines on an Active Directory domain.

1. From the Rapid Recovery Core Console button bar, click the **Protect**  drop-down menu, , and then select  **Deploy Agent Software**.
The Deploy Agent Software Wizard opens.
2. On the Connection page of the wizard, from the **Source** drop-down list, select **Active Directory**.

3. Enter the domain information and login credentials as described in the following table.

Table 40: Domain information and credentials

Text Box	Description
Host	The host name or IP address of the Active Directory domain.
User name	The user name used to connect to the domain; for example, Administrator or, if the machine is in a domain, [domain name]\Administrator).
Password	The secure password used to connect to the domain.

4. Click **Next**.
5. On the Machines page, select the machines to which you want to deploy the Rapid Recovery Agent software.
6. Optionally, to automatically restart the protected machines after the Agent is installed, select **After Agent installation, restart the machines automatically (Recommended)**.
7. Click **Finish**.
The system automatically verifies each machine that you selected.
If Rapid Recovery detects any concerns during automatic verification, the wizard progresses to a *Warnings* page, where you can clear machines from selection and manually verify the selected machines. If the machines you added pass the automatic verification, they appear on the Deploy Agent to Machines pane.
8. If the *Warning* page appeared and you are still satisfied with your selections, click **Finish** again.

The Rapid Recovery Agent software deploys to the specified machines.

CAUTION: The machines are not yet protected. Protection begins after you complete the steps in the topic [Protecting multiple machines on a VMware vCenter/ESXi virtual host](#).

Deploying to machines on a VMware vCenter/ESXi virtual host

Use this procedure to simultaneously deploy the Rapid Recovery Agent software to one or more machines on a VMware vCenter/ESXi virtual host.

Before starting this procedure, you must have the following information:

- Logon credentials for the VMware vCenter/ESXi virtual host.
- Host location.
- Logon credentials for each machine you want to protect.


NOTE: All virtual machines must have VMware Tools installed; otherwise, Rapid Recovery cannot detect the host name of the virtual machine to which to deploy. In lieu of the host name, Rapid Recovery uses the virtual machine name, which may cause issues if the host name is different from the virtual machine name

1. From the Rapid Recovery Core Console button bar, click the **Protect** drop-down menu, , and then select **Deploy Agent Software**.
The Deploy Agent Software Wizard opens.
2. On the Connection page of the wizard, from the **Source** drop-down list, select **vCenter / ESXi**.

3. Enter the host information and logon credentials as described in the following table.

Table 41: vCenter/ESXi connection settings

Text Box	Description
Host	The name or IP address of the VMware vCenter Server/ESXi virtual host.
Port	The port used to connect to the virtual host. The default setting is 443.
User name	The user name used to connect to the virtual host; for example, Administrator or, if the machine is in a domain, [domain name]\Administrator.
Password	The secure password used to connect to this virtual host.

4. Click **Next**.
5. On the Machines page of the wizard, select one of the following options from the drop-down menu:
 - Hosts and Clusters
 - VMs and Templates
6. Expand the list of machines, and then select the VMs to which you want to deploy the software. A notification appears if Rapid Recovery detects that a machine is offline or that VMware Tools are not installed.
7. If you want to restart the machines automatically after deployment, select **After Agent installation, restart the machines automatically (Recommended)**.
8. Click **Next**.
Rapid Recovery automatically verifies each machine you selected.
9. On the Adjustments page of the wizard, enter the credentials for each machine in the following format:
hostname::username::password.
 **NOTE:** Enter one machine on each line.
10. Click **Finish**.
The system automatically verifies each machine that you selected.
If Rapid Recovery detects any concerns during automatic verification, the wizard progresses to a *Warnings* page, where you can clear machines from selection and manually verify the selected machines. If the machines you added pass the automatic verification, they appear on the Deploy Agent to Machines pane.
11. If the *Warning* page appeared and you are still satisfied with your selections, click **Finish** again.

The Rapid Recovery Agent software deploys to the specified machines.

 **CAUTION:** The machines are not yet protected. Protection begins after you complete the steps in the topic [Protecting multiple machines on a VMware vCenter/ESXi virtual host](#).

Deploying an upgrade of the Rapid Recovery Agent to protected machines

You can use the Deploy Agent Software Wizard to push an upgrade of the Rapid Recovery Agent software to Windows machines that are already protected by the local Rapid Recovery Core.

NOTE: For Linux users, if the prior version of Agent is branded AppAssure (release 5.4.3 or earlier), you must first remove the Agent software using the shell script appropriate to your specific AppAssure Agent version. Removing AppAssure Agent after installing Rapid Recovery Agent can break the connection between the Linux machine and the Core. For more information on uninstalling AppAssure Agent from a Linux machine, see "Uninstalling the AppAssure Agent software from a Linux machine" in the *Rapid Recovery Installation and Upgrade Guide*.

1. From the Rapid Recovery Core Console button bar, click the **Protect** drop-down menu, and then select **Deploy Agent Software**.
The Deploy Agent Software Wizard opens.
2. On the Connection page of the wizard, from the **Source** drop-down list, select **Local Core**.
3. Click **Next**.
4. On the Machines page of the wizard, select the protected machines to which you want to deploy an upgrade of the Rapid Recovery Agent software.

NOTE: At this time, you cannot use this process to update protected Linux machines.

5. Best practice is to restart each machine after installing or updating the Agent software. If you want to restart the machines after deploying, leave the default option **After Agent installation, restart the machines automatically (Recommended)**. If you do not want to restart upgraded machines immediately, clear this option.
6. Click **Finish**.
The system automatically verifies each machine that you selected.
If Rapid Recovery detects any concerns during automatic verification, the wizard progresses to a *Warnings* page, where you can clear machines from selection and manually verify the selected machines. If the machines you added pass the automatic verification, they appear on the Deploy Agent to Machines pane.
7. If the *Warning* page appeared, and you are still satisfied with your selections, click **Finish** again.

Deploying to machines manually

Use the following procedure to deploy the Rapid Recovery Agent to multiple machines on any type of host other than the local Core, Active Directory, vCenter/ESXi, or Hyper-V.

CAUTION: If AppAssure Agent was previously installed on a Linux machine, then before installing Rapid Recovery Agent, remove the AppAssure Agent from the machine using a shell script. For information about removing the Agent from a Linux machine, see the topic "Uninstalling the AppAssure Agent software from a Linux machine" in the *Rapid Recovery Installation and Upgrade Guide*. To successfully deploy the Agent software to Linux machines, see the prerequisites in the topic "About installing the Agent software on Linux machines" in the same document.

1. From the Rapid Recovery Core Console button bar, click the **Protect** drop-down menu, , and then select **Deploy Agent Software**.
The Deploy Agent Software Wizard opens.
2. On the Connection page of the wizard, from the **Source** drop-down list, select **Manually**.
3. Click **Next**.

4. On the Machines page of the wizard, for each machine to which you want to deploy Agent, enter the machine details in the dialog box. Press *Enter* to separate information for each machine. Use the format `hostname::username::password::port`. For Windows machines, the port setting is optional. For Linux machines, always include the SSH port, which by default is 22. Examples include:

```
10.255.255.255::administrator::&11@yYz90z
```

```
abc-host-00-1::administrator::99!zU$o83r::8006
```

```
Linux-host-00-2::administrator::p@$w0rD::22
```

5. If you want to restart the machines automatically after deployment, select **After Agent installation, restart the machines automatically (Recommended)**.
6. Click **Finish**.
The system automatically verifies each machine that you selected.
If Rapid Recovery detects any concerns during automatic verification, the wizard progresses to a Warnings page, where you can clear machines from selection and manually verify the selected machines. If the machines you added pass the automatic verification, they appear on the Deploy Agent to Machines pane.
7. If the Warning page appeared, and you are still satisfied with your selections, click **Finish** again.


The Rapid Recovery Agent software deploys to the specified machines.

 **CAUTION:** The machines are not yet protected. Protection begins after you complete the steps in the topic [Protecting multiple machines manually](#).

Verifying the deployment to multiple machines

Once you have deployed the Rapid Recovery Agent software to two or more machines simultaneously, you can verify the success by viewing each machine listed under the Protected Machines menu.

You can also view information regarding the bulk deploy process from the *Events* page. Complete the steps in this procedure to verify the deployment.

1. From the Rapid Recovery Core Console, click  (Events), and then click **Alerts**.
Alert events appear in the list, showing the time the event initiated and a message. For each successful deployment of the Agent software, you will see an alert indicating that the protected machine has been added.
2. Optionally, click on any link for a protected machine.
The Summary page for the selected machine appears, showing pertinent information including:
 - the host name of the protected machine
 - the last snapshot, if applicable
 - the time of the next scheduled snapshot, based on the protection schedule for the selected machine
 - the encryption key, if any, used for this protected machine
 - the version of the software.

Modifying deploy settings

Complete the steps in this procedure to modify deploy settings.



1. From the Rapid Recovery Core Console, click  (Settings).
2. On the Settings page, in the left column, click Deploy to navigate to the Deploy section.
3. Modify any of the following options by clicking the setting you want to change to make it editable as a text box or drop-down list, and then click  to save the setting.

Table 42: Deploy options

Option	Description
Agent installer name	Enter the name of the agent executable file. The default is Agent-web.exe.
Core address	Enter the address for the Core.
Failed receive timeout	Enter the number of minutes to wait without activity before timeout.
Maximum parallel installs	Enter a number for the maximum installations you want to install simultaneously. The default and limit is 100.

Understanding protection schedules

A protection schedule defines when backups are transferred from protected agent machines to the Rapid Recovery Core.

The first backup transfer for any machine added to protection on the Core is called a base image snapshot. All data on all specified volumes (including the operating system, applications, and settings), are saved to the repository, which can take a significant amount of time depending on the amount of data being transferred. Thereafter, incremental snapshots (smaller backups, consisting only of data changed on the protected machine since the last backup) are saved to the repository regularly, based on the interval defined (for example, every 60 minutes). This type of backup contains less data than a base image, and therefore takes a shorter amount of time to transfer.

Protection schedules are initially defined using the Protect Machine Wizard or the Protect Multiple Machines Wizard. Using a wizard, you can customize protection schedules (choosing either periods or a daily protection time) to accommodate your business needs. You can then modify the existing schedule or create a new schedule at any time in the Protection Schedule dialog box from the summary page of a specific protected machine.

Rapid Recovery provides a default protection schedule, which includes a single period spanning all days of the week, with a single time period defined (from 12:00 AM to 11:59 PM). The default interval (the time period between snapshots) is 60 minutes. When you first enable protection, you also activate the schedule. Thus, using the default settings, regardless of the current time of day, the first backup will occur every hour, on the hour (12:00 AM, 1:00 AM, 2:00 AM, and so on).

Selecting periods lets you view the default protection schedule and make adjustments accordingly. Selecting a daily protection time causes Rapid Recovery Core to back up the designated protected machines once daily at a time you specify.

You can customize the schedule to define peak and off-peak times using the weekday and weekend periods available. For example, if your protected machines are mostly in use on weekdays, you could decrease the interval for the weekday period to 20 minutes, resulting in three snapshots every hour. Or you can increase the interval for the weekend period from 60 minutes to 180 minutes, resulting in snapshots once every three hours when traffic is low.

Alternatively, you can change the default schedule to define peak and off-peak times daily. To do this, change the default start and end time to a smaller range of time (for example, 12:00 AM to 4:59 PM), and set an appropriate interval (for example, 20 minutes). This represents frequent backups during peak periods. You can then add an

additional weekday time range for the remaining span of time (5:00 pm to 11:59 pm) and set an appropriate (presumably larger) interval (for example, 180 minutes). These settings define an off-peak period that includes 5:00 PM to midnight every day. This customization results in snapshots every three hours from 5:00 PM through 11:59 PM, and snapshots every 20 minutes from 12:00 AM until 4:59 PM.

When you modify or create a protection schedule using the Protection Schedule dialog box, Rapid Recovery gives you the option to save that schedule as a reusable template that you can then apply to other protected machines. Other options in the protection wizards include setting a daily protection time. This results in a single backup daily at the period defined (the default setting is 12:00 PM).

When protecting one or multiple machines using a wizard, you can initially pause protection, which defines the protection schedule without protecting the machines. When you are ready to begin protecting your machines based on the established protection schedule, you must explicitly resume protection. For more information on resuming protection, see [Pausing and resuming protection](#). Optionally, if you want to protect a machine immediately, you can force a snapshot. For more information, see [Forcing a snapshot](#).

For more information, see [Creating multiple protection schedule periods in Advanced Mode](#).

Protecting a machine

The protecting procedure requires the following tasks be completed before you begin:

- The Rapid Recovery Agent has been deployed and installed on the machine you want to protect.
- The machine you want to protect was restarted after the Agent installation.
- A repository has been created and is accessible from the Rapid Recovery Core.

This topic describes how to start protecting the data on a single machine that you specify using the Protect Machine Wizard. To protect multiple machines using one process simultaneously, see [About protecting multiple machines](#).

i NOTE: Unless using agentless protection on a VMware or ESXi host, or a Hyper-V host, the machine you want to protect must have the Rapid Recovery Agent software installed to be protected. You can choose to install the Agent software prior to this procedure, or you can deploy the software to the target machine as a part of completing the Protect Machine Wizard. For more information about agentless protection and its restrictions, see [Understanding Rapid Snap for Virtual](#).

For more information about installing the Agent software, see “Installing the Rapid Recovery Agent software” in the *Rapid Recovery Installation and Upgrade Guide*. If the Agent software is not installed prior to protecting a machine, you will not be able to select specific volumes for protection as part of this wizard. In this case, by default, all volumes on the agentlessly protected machine will be protected.

Rapid Recovery supports the protection and recovery of machines configured with EISA partitions.

When you add protection, you need to define connection information such as the IP address and port, and provide credentials for the machine you want to protect. Optionally, you can provide a display name to appear in the Core Console instead of the IP address. If you change this, you will not see the IP address for the protected machine when you view details in the Core Console. You will also define the protection schedule for the machine.

The protection process includes optional steps you can access if you select an advanced configuration. Advanced options include repository functions and encryption. For example, you can specify which Rapid Recovery repository you want to use. You can also specify an existing encryption key (or add a new encryption key) to apply to the data saved to the Core for this machine. For more information about encryption keys, see [Encryption](#).

The workflow of the protection wizard may differ slightly based on your environment. For example, if the Rapid Recovery Agent software is installed on the machine you want to protect, you will not be prompted to install it from the wizard. Likewise, if a repository already exists on the Core, you will not be prompted to create one.

CAUTION: Rapid Recovery does not support bare metal restores (BMRs) of Linux machines with ext2 boot partitions. Any BMR performed on a machine with this type of partition results in a machine that does not start. If you want to be able to perform a BMR on this machine in the future, you must convert any ext2 partitions to ext3 or ext4 before you begin protecting and backing up the machine.

1. From the Rapid Recovery Core Console button bar, click  **Protect**.

NOTE: If you have not yet create a repository, a dialog box alerts you that no repository was found and prompts you to create one. Clicking **Yes** opens the Create Repository wizard. For more information about repositories, see [Repositories](#).

The *Protect Machine Wizard* appears.

2. On the *Welcome* page, select one of the follow installation options:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you want to define a repository or if you want to establish encryption, select **Advanced (show optional steps)**.Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.
3. When you are satisfied with your choices on the *Welcome* page, then click **Next**.
The *Connection* page appears.
4. On the *Connection* page, enter the information about the machine to which you want to connect, as described in the following table, and then click **Next**.

Table 43: Machine connection settings

Text Box	Description
Host	The host name or IP address of the machine that you want to protect.
Port	The port number on which the Rapid Recovery Core communicates with the Agent on the machine. The default port number for Windows machines is 8006. The default port number for Linux machines is 22 (SSH).
User name	The user name used to connect to this machine; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator). Enter the user name or, to use a set of credentials saved to Credentials Vault, use the drop-down list and select a user name. Optionally, to save your credentials to Credentials Vault, click the plus sign next to the text box. For more information, see Credentials Vault .
Password	Enter the password used to connect to this machine, or select a set of credentials from the drop-down list, which are listed by user name.

If the *Install Agent* page appears next in the Protect Machine Wizard, that means that Rapid Recovery does not detect the Rapid Recovery Agent on the machine and will install the current version of the software. Go to step 6.

If the *Upgrade Agent* page appears next in the wizard, that means that an older version of the Agent software exists on the machine you want to protect.

i **NOTE:** The Agent software must be installed on the machine you want to protect, and that machine must be restarted, before it can back up to the Core. To have the installer reboot the protected machine, select the option **After installation, restart the machine automatically (recommended)** before clicking **Next**.

5. On the *Upgrade Agent* page, do one of the following:
 - To deploy the new version of the Agent software (matching the version for the Rapid Recovery Core), select **Upgrade the Rapid Recovery Agent software to the latest version**.
 - To continue protecting the machine without updating the Agent software version, clear the option **Upgrade the Rapid Recovery Agent software to the latest version**.
6. Click **Next**.
The *Protection* page appears.
7. Optionally, on the *Protection* page, if you want a name other than the IP address to display in the Rapid Recovery Core Console for this protected machine, then in the **Display name** field, type a name in the dialog box.
You can enter up to 64 characters. Do not use the special characters described in the topic [prohibited characters](#). Additionally, do not begin the display name with any of the character combinations described in the topic [prohibited phrases](#).
8. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
Schedule options are added to the wizard workflow.
9. At the bottom of the *Protection* page, if you are re-protecting a machine previously protected in this Core, and want to keep the previous configuration, select **Keep current settings**.
10. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in step 2 and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** to see repository and encryption options.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to the next step to choose which volumes to protect.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** and proceed to step 13 to see repository and encryption options.

The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) will transfer to the repository on the Rapid Recovery following the schedule you defined, unless you specified to initially pause protection.

11. On the *Protection Volumes* page, select which volumes you want to protect. If any volumes are listed that you do not want to include in protection, click in the **Check** column to clear the selection. Then click **Next**.

i | **NOTE:** Typically, it is good practice to protect, at minimum, the System Reserved volume and the volume with the operating system (typically the C drive).

12. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).
If you already have repository information configured, did not indicate that you want to change the repository, and you selected the Advanced option in step 2, then the *Encryption* page appears. Proceed to step 16.
If you already have repository information configured, and you selected the Advanced option in step 2, or if you indicated you want to change the repository, then the *Repository* page appears. Proceed to step 13.
13. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.

14. On the *Encryption* page, do one of the following:
- If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the Select encryption key drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 44: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters and prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

- If you want to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**.

i **NOTE:** This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.

15. Click **Finish** to save and apply your settings.

i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).

Protecting a cluster

This topic describes how to add a cluster for protection in Rapid Recovery. When you add a cluster to protection, you need to specify the host name or IP address of the cluster, the cluster application, or one of the cluster nodes or machines that includes the Rapid Recovery Agent software.

NOTE: A repository is used to store the snapshots of data that are captured from your protected nodes. Before you start protecting data in your cluster, you should have set up at least one repository that is associated with your Rapid Recovery Core.


For information about setting up repositories, see [Understanding repositories](#).


1. From the Rapid Recovery Core Console button bar, click the **Protect** drop-down menu, and then select **Protect Cluster**.
The Protect Cluster Wizard opens.
2. On the Welcome page of the wizard, select one of the following options:
 - Typical
 - Advanced (show optional steps)
3. Optionally, if you want to skip this step the next time you open the Protect Cluster Wizard, select **Skip this Welcome page the next time the wizard opens**.
4. Click **Next**.
5. On the Connection page of the wizard, enter the following information.

Table 45: Connect to Cluster settings

Text Box	Description
Host	The host name or IP address of the cluster, the cluster application, or one of the cluster nodes.
Port	The port number on the machine on which the Rapid Recovery Core communicates with the Agent. The default port is 8006.
User name	The user name of the domain administrator used to connect to this machine: for example, domain_name\administrator. NOTE: The domain name is mandatory. You cannot connect to the cluster using the local administrator user name.
Password	The password used to connect to this machine.

6. Click **Next**.
NOTE: If the cluster nodes already have an older version of the Rapid Recovery Agent installed, an *Upgrade* page appears in the wizard and offers the opportunity to upgrade the Agent.
7. On the Nodes page of the wizard, select the nodes that you want to protect.
The system automatically verifies each machine you selected.
8. Click **Next**.
If the Protection page appears next in the Protect Cluster Wizard, skip to step 11.
If the Agent software is not yet deployed to the machines you want to protect, or if any of the machines you specified cannot be protected for another reason, then the selected machines appear on the *Warnings* page.
9. Optionally, on the Warnings page of the wizard, you can verify any machine by selecting the machine and then clicking **Verify** in the toolbar.

10. Optionally, on the Warnings page, select **After Agent installation, restart the machines automatically**.
 **NOTE:** Quest recommends this option. You must restart agent machines before they can be protected.
11. If the status indicates that the machine is reachable, click **Next** to install the Rapid Recovery Agent software. The Protection page appears.
12. Optionally, on the Protection page, if you want a name other than the IP address to display in the Rapid Recovery Console for this protected machine, then in the **Display name** field, type a name in the dialog box. You can enter up to 64 characters. Do not use the special characters described in the topic [prohibited characters](#). Additionally, do not begin the display name with any of the character combinations described in the topic [prohibited phrases](#).
13. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the Schedule Settings option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the Summary page for the specific protected machine.
 - To define a different protection schedule, in the Schedule Settings option, select **Custom protection**.
14. Optionally, on the Protection page, if you want a name other than the IP address to display in the Rapid Recovery Core for this protected machine, then in the **Display name** field, type a name in the dialog box. You can enter up to 64 characters. Do not use the special characters described in the topic [prohibited characters](#). Additionally, do not begin the display name with any of the character combinations described in the topic [prohibited phrases](#).
15. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the Schedule Settings option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the Summary page for the specific protected machine.
 - To define a different protection schedule, in the Schedule Settings option, select **Custom protection**.

16. Proceed with your configuration as follows:
- If you selected a Typical configuration for the Protect Machine Wizard and specified default protection, and if a repository exists, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified default protection, and if no repository exists, then click **Next** and proceed to step 19 to create a repository.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to the next step to choose which volumes to protect.
 - If you selected Advanced configuration for the Protect Machine Wizard, and default protection, then click **Next** and proceed to step 19 to see repository and encryption options.
 - If you selected Advanced configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to the next step to choose which volumes to protect.
17. On the Protection Volumes page, select which volumes you want to protect. If any volumes are listed that you do not want to include in protection, click in the Check column to clear the selection. Then click **Next**.
-  **NOTE:** Typically, it is good practice to protect, at minimum, the System Reserved volume and the volume with the operating system (typically the C drive).
18. On the Protection Schedule page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#). If you already have repository information configured, and you selected the Advanced option in step 1, then the *Encryption* page appears. Proceed to step 22.

19. On the Repository page, do the following: If you already have a repository and want to store the data from this machine for protection in the existing repository, then do the following:
 - a. Select **Use an existing repository**.
 - b. Select an existing repository from the list.
 - c. Click **Next**.

The Encryption page appears. Skip to step 22 to optionally define encryption.

If you want to create a repository, then on the Repository page, enter the information described in the following table, and then click **Next** or **Finish**, as appropriate.

Table 46: Repository settings

Text Box	Description
Name	<p>Enter the display name of the repository.</p> <p>By default, this text box consists of the word Repository and a number, which corresponds to the number of repositories for this Core. For example, if this is the first repository, the default name is Repository 1. Change the name as needed.</p> <p>Repository names must contain between 1 and 40 alphanumeric characters, including spaces. Do not use prohibited characters or prohibited phrases.</p>
Location	<p>Enter the location for storing the protected data. This volume should be a primary storage location. The location could be local (a drive on the Core machine) or it could be a CIFS share network drive.</p> <p>For a CIFS share, the path must begin with \\ . When specifying the path, use only alphanumeric characters, the hyphen, and the period (only to separate host names and domains). The letters a to z are case-insensitive. Do not use spaces. No other symbols or punctuation characters are permitted.</p> <div> <p>CAUTION: Define a dedicated folder within the root for the storage location for your repository. Do not specify the root location. For example, use <code>E:\Repository\</code>, not <code>E:\</code>. If the repository that you are creating in this step is later removed, all files at the storage location of your repository are deleted. If you define your storage location at the root, all other files in the volume (e.g., <code>E:\</code>) are deleted, which could result in catastrophic data loss.</p> <p>For example, type <code>X:\Repository\Data</code>.</p> <p>When specifying the path, use only alphanumeric characters, the hyphen, and the period (only to separate host names and domains). You can use the backslash character only to define levels in the path. Do not use spaces. No other symbols or punctuation characters are permitted.</p> </div>
User name	Enter a user name for a user with administrative access. This information is only required when the repository location specified is a network path.
Password	Enter the password for the user with administrative access. This information is only required when the repository location specified is a network path.
Metadata path	<p>Enter the location for storing the protected metadata.</p> <p>For example, type <code>X:\Repository\Metadata</code>.</p> <p>When specifying the path, use only alphanumeric characters, the hyphen, and the period</p>

Text Box	Description
	(only to separate host names and domains). You can use the backslash character only to define levels in the path. Do not use spaces. No other symbols or punctuation characters are permitted.

20. On the Repository Configuration page, configure the amount of space on the disk to allocate to the repository. Quest recommends reserving 20 percent of the volume for metadata, which is the default. Optionally, using the slider control or the **Size** field, you can allocate more or less space on the volume for the repository.
21. Optionally, if you want to view and enter detailed information regarding storage location settings, select **Show advanced options**, and adjust the settings as described in the following table. To hide these options, clear the option **Show advanced options**. When you have completed your repository configuration, click **Next** or **Finish**, as appropriate.

Table 47: Storage configuration details

Text Box	Description
Bytes per sector	<p>Specify the number of bytes you want each sector to include. The default value is 512.</p> <p>i NOTE: Quest recommends setting the bytes per sector to match the physical sector size of the storage location in which the repository resides. For example, if the disk on the intended storage location has a 4096 byte sector size, change the bytes per sector setting to 4096.</p> <p>If using multiple storage locations with different sector sizes, Quest recommends retaining the default setting of 512 bytes per sector.</p>
Bytes per record	Specify the average number of bytes per record. The default value is 8192.
Write caching policy	<p>The write caching policy controls how the Windows Cache Manager is used in the repository and helps to tune the repository for optimal performance on different configurations. As of release 6.4, the default for this setting is Off.</p> <p>Set the value to one of the following:</p> <ul style="list-style-type: none"> • On. If set to On, Windows controls the caching. • Off. If set to Off, which is the default, Rapid Recovery controls the caching. • Sync. If set to Sync, Windows controls the caching as well as the synchronous input/output.

If you chose the **Advanced** option in Step 1, the Encryption page appears.

22. Optionally, on the Encryption page, if you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
- To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the **Select encryption key** drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 48: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

23. Optionally, on the Encryption page, to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**. This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.
24. Click **Finish** to save and apply your settings.

i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).

Protecting nodes in a cluster

This task requires that you first protect a cluster. For more information, see [Protecting a cluster](#).

This topic describes how to protect the data on a cluster node or machine that has a Rapid Recovery Agent installed. This procedure lets you add individual nodes to protection that you may have omitted when you protected a cluster.


1. In the Rapid Recovery Core Console, under Protected Machine, click the cluster with the nodes that you want to protect.
2. On the Summary page for the cluster, click the **Protected Nodes**.

3. On the Protected Nodes page, click **Protect Cluster Node**.
4. In the Protect Cluster Node dialog box, select or enter as appropriate the following information.

Table 49: Protect Cluster Node settings

Text Box	Description
Host	A drop-down list of nodes in the cluster available for protection.
Port	The port number on which the Rapid Recovery Core communicates with the Agent on the node.
User name	The user name of the domain administrator used to connect to this node; for example, example_domain\administrator or administrator@example_domain.com.
Password	The password used to connect to this machine.

5. To add the node, click **Connect**.
6. To start protecting this node with default protection settings, go to step 11.

 **NOTE:** The default settings ensure that all volumes on the machine are protected with a schedule of every 60 minutes.
7. In the Protect [Node Name] dialog box, if you want to use a repository other than the default setting, use the drop-down list to select a repository.
8. If you want to secure the recovery points for this cluster using Core-based encryption, use the drop-down list to select an encryption key.
9. If you do not want protection to begin immediately after completing this procedure, select **Initially pause protection**.
10. To enter custom settings (for example, to customize the protection schedule for the shared volumes), do the following:
 - a. To customize settings for an individual volume, next to the volume that you want to customize, click **Function** next to the relevant volume.
 - b. See [Creating custom protection schedules in Simple Mode](#).
11. Click **Protect**.

Creating custom protection schedules in Simple Mode

The procedure for creating a custom protection schedule from within a protection wizard is identical to the procedure for creating a protection schedule for an existing machine in Simple Mode. Protection schedules created in a wizard or in Simple Mode are not saved as templates. To create templates or multiple protection schedules, see [Creating multiple protection schedule periods in Advanced Mode](#). Complete the steps in this procedure to create custom schedules for using Rapid Recovery to back up data from protected machines.

To create custom protection schedules in Simple Mode

1. Complete one of the following options:
 - If using a protection wizard (Protect Machine, Protect Multiple Machines, Protecting a Cluster), on the Protection page of the wizard, select **Custom protection**, and then click **Next**.
 - If creating a protection schedule for a machine that is already protected, on the *Summary* page for the protected machine, expand the volumes of the protected machine, select the applicable volumes, and then click **Set a Schedule**.

The Protection Schedule page or dialog appears.

2. On the Protection Schedule page or dialog, complete one of the following options:
 - To set a protection period that runs on set days and at specified times, select **Periods**, and then continue to step 3.
 - To set a specific time to back up the machine every day, select **Daily protection time**, and then continue to step 7.
 - To take a backup after a specific number of days, select **Every number of days**, and then continue to step 7.
3. To change the interval schedule for any period, complete the following steps:
 - a. Create a span of time by selecting a **From** time and a **To** time.
 - b. For each period, click in the interval text box, and then enter an appropriate interval in minutes. For example, highlight the default interval of 60 and replace it with the value 20 to perform snapshots every 20 minutes during this period.
4. To customize snapshots for peak and off-peak business hours, complete the following steps: set an optimal interval for the peak range, select **Take snapshots for the remaining time**, and then set an off-peak interval by doing the following:
 - a. Select **Weekdays**.
 - b. Set the **From** and **To** times to create a span of time during your peak business hours.
 - c. In the **Every X minutes** box, enter an interval in minutes for how often Rapid Recovery should create recovery points during this span of peak business hours. For example, highlight the default interval of 60 and replace it with the value 20 to perform snapshots every 20 minutes during the time range you selected for this period.
 - d. To schedule snapshots during off-peak business hours, select **Take snapshots for the rest of the time**.
 - e. In the **Every X minutes** box, enter an interval in minutes for how often Rapid Recovery Core should create recovery points during this span of off-peak business hours. For example, because there is less business activity during these hours, and fewer changes to back up, you may decide to take fewer snapshots and keep the default interval of 60 minutes.
5. Continue to step 7.
6. To set a single time of day for a single backup to occur daily, select **Daily protection time** and then enter a time in format `HH:MM AM`. For example, to do a daily backup at 9:00 PM, enter 09:00 PM.
7. To define the schedule without beginning backups, select **Initially pause protection**. After you pause protection from the wizard, it remains paused until you explicitly resume it. Once you resume protection, backups occur based on the schedule you established. For more information on resuming protection, see [Pausing and resuming protection](#).

8. When you are satisfied with changes made to your protection schedule, click **Finish** or **Next**, as appropriate. If you are using a wizard, proceed to the next step in the wizard. Refer to the procedure for the appropriate wizard to complete any remaining requirements.

For more information, see the following related topics:

- [Understanding protection schedules](#)
- [Creating multiple protection schedule periods in Advanced Mode](#)

Creating multiple protection schedule periods in Advanced Mode


A protection schedule defines when backups are transferred from protected machines to the Rapid Recovery Core. Protection schedules are initially defined using the Protect Machine Wizard or the Protect Multiple Machines Wizard.

You can modify an existing protection schedule at any time from the Summary page for a specific protected machine.

i **NOTE:** For conceptual information about protection schedules, see [Understanding protection schedules](#). For information about protecting a single machine, see [Protecting a machine](#). For information about bulk protect (protecting multiple machines), see [About protecting multiple machines](#). For information on customizing protection periods when protecting an agent using either of these wizards, see [Creating custom protection schedules in Simple Mode](#). For information about modifying an existing protection schedule, see [Creating multiple protection schedule periods in Advanced Mode](#).

Complete the steps in this procedure to modify an existing protection schedule for volumes on a protected machine.

1. In the Rapid Recovery Core Console, from the list of protected machines, click the name of the machine with the protection schedule that you want to change.
2. On the *Summary* page for the machine you selected, in the Volumes pane, expand the volumes of the protected machine, select the applicable volumes, and then click **Set a Schedule**.
Initially, all volumes share a protection schedule.


To select all volumes at once, click in the checkbox in the header row. Click on  (Protection Group) to expand the volumes being protected, so you can view all volumes and select one or more.

i **NOTE:** It is best practice to protect, at minimum, the system reserved volume and the volume with the operating system (typically the C: / drive).

The Protection Schedule dialog box appears.

3. On the Protection Schedule dialog box, do one of the following:
 - If you previously created a protection schedule template and want to apply it to this protected machine, click **Advanced mode**, select the template from the **Templates** drop-down list, click **OK** to confirm, and then go to step 7.
 - If you want to remove an existing time period from the schedule, clear the check box next to each time period option, and then go to Options include the following:
 - **Weekdays (Mon - Fri)**: This range of time denotes a typical five-day work week.
 - **Weekends (Sat, Sun)**: This range of time denotes a typical weekend.

If you want to save a new protection schedule as a template, click **Advanced mode**, and then continue to step 4.

4. A period is a specified span of time during which you determine how many minutes should pass between each snapshot taken. When the weekday start and end times are from 12:00 AM to 11:59 PM, then a single period exists. To change the start or end time of a defined period, do the following:
 - a. Select the appropriate time period.
 - b. To change the start time for this period, use the clock icon under **Start Time**.
For example, use the arrows to show a time of 08:00 AM.
 - c. To change the end time for this period, use the clock icon under **End Time**.
For example, use the arrows to show a time of 06:00 PM.
 - d. Change the interval according to your requirements. For example, if defining a peak period, change the interval from 60 minutes to 20 minutes to take snapshots three times hourly.
A blue bar provides a visual representation of this interval.
 5. If you defined a period other than 12:00 AM to 11:59 PM in step 7, and you want backups to occur in the remaining time ranges, you must add more periods to define protection by doing the following:
 - a. Under the appropriate category, click **Add Period**.
 - b. Click the clock icon and select the desired start and end times, as appropriate.
For example, set a start time of 12:00 AM and an end time of 07:59 AM.
 - c. Change the interval according to your requirements. For example, if defining an off-peak period, change the interval from 60 minutes to 120 minutes to take snapshots every two hours.
 6. If needed, continue to create more periods, setting start and end times and intervals as appropriate.
-  **NOTE:** If you want to remove a period you added, click the trash icon to the far right of that period, and then click **Yes** to confirm.
7. To create a template from the schedule you set, click **Save as a Template**.
 8. In the *Save Template* dialog box, enter a name for the template, and then click **Save**.
 9. When your protection schedule meets your requirements, click **Apply**.
The protection Schedule dialog box closes.

For more information, see [Understanding protection schedules](#).

Pausing and resuming protection

When you pause protection, you temporarily stop all transfers of data from the selected machine to the Rapid Recovery Core. You can pause protection for any protected machine:

- When establishing protection using the Protect Machine Wizard or the Protect Multiple Machines Wizard.
- From the Protected Machines drop-down menu in the left navigation area of the Rapid Recovery Core (pausing protection for all protected machines).
- From the *Protected Machines* page (accessible when you click on the Protected Machines menu).
- From a specific protected machine in the Protected Machines drop-down menu.
- From the top of every page for a specific protected machine.

If you pause protection using the Protect Machine Wizard or the Protect Multiple Machines Wizard, protection is paused until explicitly resumed. If protection is paused using this procedure, you are prompted to specify whether to pause until resumed, or to pause for a designated amount of time (specified in any combination of days, hours and minutes).



- If you pause protection for a period of time, then when that time expires, the system automatically resumes protection based on the protection schedule.
- If you select **Pause until resumed**, then protection is paused until explicitly resumed using the Resume function described in this procedure.





If protection for a machine is paused, and you force a snapshot, a single backup snapshot is captured. If it is the first backup, a base image is captured; otherwise an incremental image is captured. Until you explicitly resume protection using this procedure, paused protection remains suspended.

You can resume protection for any paused protected machine:

- From the Protected Machines drop-down menu in the left navigation area of the Rapid Recovery Core (resuming protection for all protected machines).
- From a specific protected machine in the Protected Machines drop-down menu.
- From the *Protected Machines* page (accessible when you click on the Protected Machines menu).
- From the top of every page for a specific protected machine.

Use the procedure below to pause or to resume protection, as appropriate.


1. From the Rapid Recovery Core Console, to pause protection for all machines, click the **Protected Machines** drop-down menu in the left navigation area, and then do the following:
 - a. Select  **Pause Protection**.
The *Pause Protection* dialog box appears.
 - b. Select the appropriate setting using one of the options described below, and then click **OK**.
 - If you want to pause protection until you explicitly resume it, select **Pause until resumed**.
 - If you want to pause protection for a specified period, select **Pause for** and then, in the **Days**, **Hours**, and **Minutes** controls, type or select the appropriate pause period as appropriate.
2. To resume protection for all machines, do the following:
 - a. Select  **Resume Protection**.
The *Resume Protection* dialog box appears.
 - b. In the *Resume Protection* dialog box, select **Yes**.
The *Resume Protection* dialog box closes, and protection is resumed for all machines.

3. To pause protection for a single machine, then in the left navigation area, click the drop-down menu to the right of the machine you want to affect, and then do the following:
 - a. Select  **Pause Protection**.
The *Pause Protection* dialog box appears.
 - b. Select the appropriate setting using one of the options described below, and then click **OK**.
 - If you want to pause protection until you explicitly resume it, select **Pause until resumed**.
 - If you want to pause protection for a specified period, select **Pause for** and then, in the **Days**, **Hours**, and **Minutes** controls, type or select the appropriate pause period as appropriate.
4. To resume protection for a single machine, do the following:
 - a. Select  **Resume Protection**.
The *Resume Protection* dialog box appears.
 - b. In the *Resume Protection* dialog box, select **Yes**.
The *Resume Protection* dialog box closes, and protection is resumed for the selected machine.
5. To pause protection for a single machine from the machine pages, navigate to the machine that you want to affect.
The *Summary* page displays for the selected machine.
 - a. At the top of the page, click  **Pause**.
The *Pause Protection* dialog box appears.
 - b. Select the appropriate setting using one of the options described below, and then click **OK**.
 - If you want to pause protection until you explicitly resume it, select **Pause until resumed**.
 - If you want to pause protection for a specified period, select **Pause for** and then, in the **Days**, **Hours**, and **Minutes** controls, type or select the appropriate pause period as appropriate.
6. If you want to resume protection, do the following:
 - a. At the top of the page, click  **Resume**.
 - b. In the *Resume Protection* dialog box, click **Yes**.
The *Resume Protection* dialog box closes, and protection resumes for the selected machine.

About protecting multiple machines

You can add two or more Windows machines for protection on the Rapid Recovery Core simultaneously using the Protect Multiple Machines Wizard. To protect your data using Rapid Recovery, you need to add the workstations and servers for protection in the Rapid Recovery Core Console; for example, your Exchange server, SQL Server, Linux server, and so on.

As with protecting individual machines, protecting multiple machines simultaneously requires you to install the Rapid Recovery Agent software on each machine you want to protect.

 **NOTE:** As an exception to this rule, if protecting virtual machines on a VMware/ESXi or Hyper-V host, you can use agentless protection. For more information, including restrictions for agentless protection, see [Understanding Rapid Snap for Virtual](#).

Protected machines must be configured with a security policy that makes remote installation possible.

To connect to the machines, they must be powered on and accessible.

There is more than one method to deploy the Agent software to multiple machines simultaneously. For example:

- You can install the Rapid Recovery Agent software to multiple machines using the Deploy Agent Software Wizard. For more information, see [Using the Deploy Agent Software Wizard to deploy to one or more machines](#).
- You can deploy the Rapid Recovery Agent software as part of the Protect Multiple Machines Wizard workflow.

The process of protecting multiple machines includes optional steps that you can access if you select an advanced configuration. Advanced options include repository functions and encryption. For example, you can specify an existing Rapid Recovery repository to save snapshots, or you can create a new repository. You can also specify an existing encryption key (or add a new encryption key) to apply to the data saved to the Core for the machines you are protecting.

The workflow of the Protect Multiple Machines Wizard may differ slightly based on your environment. For example, if the Rapid Recovery Agent software is installed on the machines you want to protect, you are not prompted to install it from the wizard. Likewise, if a repository already exists on the Core, you are not prompted to create one.

When protecting multiple machines, follow the appropriate procedure, based on your configuration. See the following options for protecting multiple machines:



- [Protecting multiple machines on an Active Directory domain](#)
- [Protecting multiple machines on a VMware vCenter/ESXi virtual host](#)
- [Protecting multiple machines on a Hyper-V virtual host](#)
- [Protecting multiple machines manually](#)

Protecting multiple machines on an Active Directory domain

The protecting procedure requires the following tasks be completed before you begin:

- The Rapid Recovery Agent has been deployed and installed on the machine you want to protect.
- The machine you want to protect was restarted after the Agent installation.
- A repository has been created and is accessible from the Rapid Recovery Core.


Use this procedure to simultaneously protect one or more machines on an Active Directory domain.

1. From the Rapid Recovery Core Console button bar, click the **Protect**  drop-down menu, and then click  **Protect Multiple Machines**.
The Protect Multiple Machines Wizard opens.
2. On the *Welcome* page, select one of the follow installation options:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you want to define a repository or if you want to establish encryption, select **Advanced (show optional steps)**.
Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.
3. Click **Next**.


4. On the *Connection* page of the wizard, from the **Source** drop-down list, select **Active Directory**.
5. Enter the domain information and credentials as described in the following table.

Table 50: Domain information and credentials

Text Box	Description
Host	The host name or IP address of the Active Directory domain.
User name	The user name used to connect to the domain; for example, Administrator or, if the machine is in a domain, [domain name]\Administrator).
Password	The secure password used to connect to the domain.

6. Click **Next**.
7. On the *Select Machines* page of the wizard, select the machines you want to protect.
The system automatically verifies each machine you selected.
8. Click **Next**.
If the *Protection* page appears next in the Protect Multiple Machines Wizard, skip to step 12.
If the Agent software is not yet deployed to the machines you want to protect, or if any of the machines you specified cannot be protected for another reason, then the selected machines appear on the *Warnings* page.
9. Optionally, on the *Warnings* page of the wizard, you can verify any machine by selecting the machine and then clicking Verify in the toolbar.
10. Optionally, on the *Warnings* page, select After Agent installation, restart the machines automatically.
 **NOTE:** Quest recommends this option. You must restart agent machines before they can be protected.
11. If the status indicates that the machine is reachable, click **Next** to install the Rapid Recovery Agent software.
The *Protection* page appears.
12. Optionally, on the *Protection* page, if you want a name other than the IP address to display in the Rapid Recovery Core Console for this protected machine, then in the **Display name** field, type a name in the dialog box.
You can enter up to 64 characters. Do not use the special characters described in the topic [prohibited characters](#). Additionally, do not begin the display name with any of the character combinations described in the topic [prohibited phrases](#).
13. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
Schedule options are added to the wizard workflow.

14. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in step 2 and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** to see repository and encryption options.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to step 15 to choose which volumes to protect.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** and proceed to step 17 to see repository and encryption options.
15. On the *Protection Volumes* page, select which volumes you want to protect. If any volumes are listed that you do not want to include in protection, click in the **Check** column to clear the selection. Then click **Next**.

 **NOTE:** It is best practice to protect, at minimum, the system reserved volume and the volume with the operating system (typically the C : / drive).
16. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).
17. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.

18. On the *Encryption* page, do one of the following:

- If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the Select encryption key drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 51: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

- If you want to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**.

i **NOTE:** This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.

19. Click **Finish** to save and apply your settings.

i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).

The Rapid Recovery Agent software is deployed to the specified machines, if necessary, and the machines are added to protection on the Core.



Protecting multiple machines on a VMware vCenter/ESXi virtual host

While you can protect multiple machines simultaneously using this wizard, all the machines you protect in a single installation using the Protect Multiple Machines wizard must either use the Rapid Recovery Agent software, or use agentless protection. This process describes protection using the Agent software. For the procedure to protect multiple machines simultaneously using agentless protection, see [Protecting vCenter/ESXi virtual machines using agentless protection](#).

The protecting procedure requires the following tasks be completed before you begin:

- The Rapid Recovery Agent has been deployed and installed on the machine you want to protect.
- The machine you want to protect was restarted after the Agent installation.
- A repository has been created and is accessible from the Rapid Recovery Core.

Use this procedure to simultaneously protect one or more machines on a VMware vCenter/ESXi virtual host using Rapid Recovery Agent.

1. From the Rapid Recovery Core Console button bar, click the **Protect**  drop-down menu, and then select  **Protect Multiple Machines**.
The Protect Multiple Machines Wizard opens.
2. On the *Welcome* page, select one of the follow installation options:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you want to define a repository or if you want to establish encryption, select **Advanced (show optional steps)**.

Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.
3. Click **Next**.
4. On the *Connection* page of the wizard, from the **Source** drop-down list, select **vCenter / ESXi**.
5. Enter the host information and logon credentials as described in the following table.


Table 52: vCenter/ESXi connection settings

Text Box	Description
Host	The name or IP address of the VMware vCenter Server/ESXi virtual host.
Port	The port used to connect to the virtual host. The default setting is 443.
User name	The user name used to connect to this machine; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator). Enter the user name or, to use a set of credentials saved to Credentials Vault, use the drop-down list and select a user name. Optionally, to save your credentials to Credentials Vault, click the plus sign next to the text box. For more information, see Credentials Vault .
Password	The secure password used to connect to this virtual host.


6. To use agentless protection, select **Use Rapid Snap for Virtual host-based protection**, and then see [Protecting vCenter/ESXi virtual machines using agentless protection](#).
7. Click **Next**.
8. On the *Select Machines* page, do one of the following:
 - From the visible list of VMs, select the VMs you want to protect.
 - To navigate through the VMware tree structure to locate more VMs, click **View Tree** and then select either **Hosts and Clusters** or **VMs and Templates**. Select the VMs you want to protect.

A notification appears if Rapid Recovery detects that a machine is offline or does not have VMware Tools installed.

9. Click **Next**.
10. On the *Adjustments* page, enter the credentials for each machine in the following format:
hostname::username::password.


 **NOTE:** Enter one machine on each line.

11. Click **Next**.
If the *Protection* page appears next in the Protect Multiple Machines Wizard, skip to step 15.
If the Agent software is not yet deployed to the machines you want to protect, or if any of the machines you specified cannot be protected for another reason, then the selected machines appear on the *Warnings* page.
12. Optionally, on the *Warnings* page, you can verify any machine by selecting the machine and then clicking **Verify** in the toolbar.
13. Optionally, on the *Warnings* page, select **After Agent installation, restart the machines automatically**.

 **NOTE:** Quest recommends this option. You must restart agent machines before they can be protected.

14. If the status indicates that the machine is reachable, click **Next** to install the agent software.
The *Protection* page appears.
15. Optionally, on the *Protection* page, if you want a name other than the IP address to display in the Rapid Recovery Core Console for this protected machine, then in the **Display name** field, type a name in the dialog box.
You can enter up to 64 characters. Do not use the special characters described in the topic [prohibited characters](#). Additionally, do not begin the display name with any of the character combinations described in the topic [prohibited phrases](#).
16. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
Schedule options are added to the wizard workflow.

17. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in step 2 and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** to see repository and encryption options.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to the next step to choose which volumes to protect.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** and proceed to step 20 to see repository and encryption options.
18. On the *Protection Volumes* page, select which volumes you want to protect. If any volumes are listed that you do not want to include in protection, click in the **Check** column to clear the selection. Then click **Next**.

 **NOTE:** It is best practice to protect, at minimum, the system reserved volume and the volume with the operating system (typically the C : \ drive).
19. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).
20. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.

21. On the *Encryption* page, do one of the following:

- If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the Select encryption key drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 53: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

- If you want to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**.

i **NOTE:** This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.

22. Click **Finish** to save and apply your settings.

i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).


The Rapid Recovery Agent software is deployed to the specified machines, if necessary, and the machines are added to protection on the Core.

Protecting vCenter/ESXi virtual machines using agentless protection

This procedure requires that you first create a repository that is accessible from the Rapid Recovery Core. For more information, see [Understanding repositories](#).

Complete the following procedure to agentlessly protect one or more ESXi virtual machines (VMs).

NOTE: Quest recommends that VMware Tools be installed on virtual machines (VMs) you want to protect on vSphere or ESXi hosts. When VMware Tools are installed on a VM using a Windows operating system (OS), the backups that the Rapid Recovery Core captures use Microsoft Volume Shadow Services (VSS). For information on the behavior of agentless VMs with or without VMware Tools, see [Benefits of installing hypervisor tools for agentless protection](#).

1. From the Rapid Recovery Core Console button bar, click the **Protect** drop-down menu, and then select  **Protect Multiple Machines**.
The Protect Multiple Machines Wizard opens.
2. On the *Welcome* page, select one of the follow installation options:
 - If you do not have multiple repositories defined for this Core, or you do not need to establish encryption, select **Typical**.
 - If you have multiple repositories defined, or if you want to establish encryption, select **Advanced (show optional steps)**.

Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.
3. Click **Next**.
4. On the *Connection* page of the wizard, from the **Source** drop-down list, select **vCenter/ESXi**.
5. Enter the host information and logon credentials as described in the following table.

Table 54: vCenter/ESXi connection settings

Text Box	Description
Host	The name or IP address of the virtual host.
Port	The port used to connect to the virtual host. The default setting is 443.
User name	The user name used to connect to this machine; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator). Enter the user name or, to use a set of credentials saved to Credentials Vault, use the drop-down list and select a user name. Optionally, to save your credentials to Credentials Vault, click the plus sign next to the text box. For more information, see Credentials Vault .
Password	The secure password used to connect to this virtual host.

6. Ensure that **Use Rapid Snap for Virtual host-based protection** is selected. (This option is selected by default).
7. Click **Next**.

8. On the *Select Machines* page, select the VMs you want to protect. You can use the drop-down menu to display a tree of **Hosts and Clusters** or of **VMs and Templates** exactly as they appear in your vCenter/ESXi environment.

CAUTION: Quest recommends that you limit agentless protection to no more than 200 VMs at once. For example, do not select more than 200 VMs when using the Protect Multiple Machines Wizard. Protecting more than 200 VMs results in slow performance. There is no limit to how many VMs a Core can agentlessly protect over time. For example, you could protect 200 VMs today and another 200 VMs tomorrow.

NOTE: VMware Changed Block Tracking (CBT) must be enabled on each of the VMs you want to protect. If it is not enabled, Rapid Recovery automatically enables CBT to ensure protection.

9. If you want to automatically protect new VMs when they are added to the host, select **Auto protect new machines**, and then complete the following steps.
 - a. Click **Next**.
 - b. On the *Auto Protection* page, select any containers in which you expect to add new machines.

NOTE: You may need to switch between views of **Hosts and Clusters** and **VMs and Templates**.

10. Click **Next**.

11. On the *Protection Rules* page, select any of the following options:

Table 55: ESXi and vCenter agentless protection options

Option	Description
Protect machine if it is orphaned by this Core	Lets the Core protect a machine that was previously protected but was then removed from protection because the hypervisor became unreachable. This option is selected by default.
Protect machine if it already has recovery points	Shows existing recovery points alongside the new recovery points after protection. This option is selected by default.
Protect machine agentlessly if it is already protected with the Rapid Recovery Agent	If a Core detects that a machine is already protected by the Rapid Recovery Agent, this option permits duplicate protection (both agentlessly and with the Agent). The protected VM must be powered on and VMware Tools must be installed. This option is selected by default.
Protect machine if it is paired with a different Core	Protects the VM with this Core and discontinues protection from the other Core.
Delete old VMware snapshots in order to enable Changed Block Tracking	Lets the Core delete previous VMware snapshots, including snapshots created by a user or another program, if required to enable Changed Block Tracking (CBT).
Save rules	Saves the selected rules to use for future VM agentless protection on this hypervisor host. These rules apply to machines protected automatically or by using the Protect Multiple Machines wizard.

12. Click **Next**.

13. On the *Protection* page, select the appropriate protection schedule settings as described below:

- To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core takes snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
- To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
Schedule options are added to the wizard workflow.

14. Proceed with your configuration as follows:

- If you specified default protection, then click **Next** and continue to step 16 to the ABM Settings page.
- If you specified custom protection, then click **Next** and continue to the next step to configure a protection schedule.

15. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).
16. Optionally, on the ABM Settings page, select **Enable Active Block Mapping**, and then complete the following information:

Table 56: Active Block Mapping settings

Option	Description
Enable Active Block Mapping	Lets you enable or disable the ABM feature.
Enable swap file blocks exclusion	Excludes the content of system files, such as pagefile.sys, hyberfill.sys, and swapfile.sys, from the backup.
Exclude subdirectories	<p>Lets you exclude specific files by specifying '<file name>' or '<folder>\<subfolder>\<file name>'.</p> <p>Only the files will be excluded. The folders or subfolders that contained excluded files are included in the mount point, with no contents.</p> <p>i NOTE: This option may affect the performance of the "determining data" phase of transfers.</p>
+ Add	If you opted to exclude subdirectories, click Add and enter the location in the Path table for each item you want to exclude.

For more information, see [Understanding Active Block Mapping](#).

i **NOTE:** Active Block Mapping only supports NTFS file systems.

17. Click **Next**.
18. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in step 2 and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** to see repository and encryption options.
19. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.

20. On the *Encryption* page, do one of the following:

- If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the **Select encryption key** drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 57: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

21. Click **Finish** to save and apply your settings.

i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).

Protecting multiple machines on a Hyper-V virtual host

The protecting procedure requires the following tasks be completed before you begin:

- The Rapid Recovery Agent has been deployed and installed on the machine you want to protect.
- The machine you want to protect was restarted after the Agent installation.
- A repository has been created and is accessible from the Rapid Recovery Core.

Use this procedure to simultaneously protect one or more machines on a Hyper-V virtual host.

CAUTION: If you use agentless protection, Quest recommends that you limit protection to no more than 200 VMs at once. For example, do not select more than 200 VMs when using the **Protect Multiple Machines Wizard**. Protecting more than 200 VMs results in slow performance. There is no limit to how many VMs a Core can agentlessly protect over time. For example, you could protect 200 VMs today and another 200 VMs tomorrow.

1. From the Rapid Recovery Core Console button bar, click the **Protect** drop-down menu, and then select **Protect Multiple Machines**.
The Protect Multiple Machines Wizard opens.
2. On the *Welcome* page, select one of the follow installation options:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you want to define a repository or if you want to establish encryption, select **Advanced (show optional steps)**.
Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.
3. Click **Next**.
4. On the *Connection* page of the wizard, from the **Source** drop-down list, select one of the following options:
 - Hyper-V Server
 - Hyper-V Cluster
5. Enter the host information and logon credentials as described in the following table.


Table 58: Hyper-V connection settings

Text Box	Description
Host	The name or IP address of the virtual host.
Port	The port used to connect to the virtual host. The default setting is 8006.
User name	The user name used to connect to the virtual host; for example, Administrator or, if the machine is in a domain, [domain name]\Administrator.
Password	The secure password used to connect to this virtual host.

6. To use agentless protection, select **Use Rapid Snap for Virtual host-based protection**, and then see [Protecting Hyper-V virtual machines using host-based protection](#).
7. Click **Next**.
8. On the *Machines* page, select the VMs that you want to protect.
9. Optionally, if you want to automatically protect new VMs when they are added to the host, select **Auto protect new virtual machines**.
10. Click **Next**.
11. On the *Adjustments* page, enter the credentials for each machine in the following format:
hostname::username::password.

NOTE: Enter one machine on each line.


12. Click **Next**.
 - If the *Protection* page appears next in the Protect Multiple Machines Wizard, skip to [Step 15](#).
 - If the Agent software is present on the machines you want to protect, or if the machines you specified cannot be protected for another reason, then the selected machines appear on the Warnings page.
13. Optionally, on the *Warnings* page, you can verify any machine by selecting the machine and then clicking **Verify** in the toolbar.
14. Optionally, on the *Warnings* page, select **After Agent installation, restart the machines automatically**.

 **NOTE:** Quest recommends this option. You must restart agent machines before they can be protected.
15. If the status indicates that the machine is reachable, click **Next** to install the agent software. The *Protection* page appears.
16. Optionally, on the *Protection* page, if you want a name other than the IP address to display in the Rapid Recovery Core Console for this protected machine, then in the **Display name** field, type a name in the dialog box.

You can enter up to 64 characters. Do not use [prohibited characters](#). Additionally, do not begin the display name with any of these [prohibited phrases](#).
17. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.

With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.

Schedule options are added to the wizard workflow.
18. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in step 2 and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** to see repository and encryption options.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to step 19 to choose which volumes to protect.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, and default protection, then click **Next** and proceed to step 21 to see repository and encryption options.
19. On the *Protection Volumes* page, select which volumes you want to protect. If any volumes are listed that you do not want to include in protection, click in the **Check** column to clear the selection. Then click **Next**.


 **NOTE:** It is best practice to protect, at minimum, the system reserved volume and the volume with the operating system (typically the C: \ drive).
20. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).


21. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.
22. On the *Encryption* page, do one of the following:
 - If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the Select encryption key drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 59: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

- If you want to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**.

 **NOTE:** This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.
23. Click **Finish** to save and apply your settings.

 **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).



The Rapid Recovery Agent software is deployed to the specified machines, if necessary, and the machines are added to protection on the Core.

Protecting Hyper-V virtual machines using host-based protection

This procedure requires that you first create a repository that is accessible from the Rapid Recovery Core. For more information, see [Understanding repositories](#).

The Rapid Snap for Virtual feature lets you protect Hyper-V virtual machines or clusters agentlessly by installing the Rapid Recovery Agent on only the Hyper-V host instead of every virtual machine (VM).

CAUTION: Quest recommends that you limit agentless protection to no more than 200 VMs at once. For example, do not select more than 200 VMs when using the Protect Multiple Machines Wizard. Protecting more than 200 VMs results in slow performance. There is no limit to how many VMs a Core can agentlessly protect over time. For example, you could protect 200 VMs today and another 200 VMs tomorrow.

1. From the Rapid Recovery Core Console button bar, click the **Protect**  drop-down menu, and then select  **Protect Multiple Machines**.
The Protect Multiple Machines Wizard opens.
2. On the *Welcome* page, select one of the follow installation options:
 - If you do not have multiple repositories defined for this Core, or you do not need to establish encryption, select **Typical**.
 - If you have multiple repositories defined, or if you want to establish encryption, select **Advanced (show optional steps)**.


Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.
3. Click **Next**.
4. On the *Connection* page of the wizard, from the **Source** drop-down list, select one of the following options:
 - Hyper-V Server
 - Hyper-V Cluster

5. Enter the host information and logon credentials as described in the following table.

Table 60: Hyper-V connection settings

Text Box	Description
Host	The name or IP address of the virtual host.
Port	The port used to connect to the virtual host. The default setting is 8006.
User name	The user name used to connect to this machine; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator). Enter the user name or, to use a set of credentials saved to Credentials Vault, use the drop-down list and select a user name. Optionally, to save your credentials to Credentials Vault, click the plus sign next to the text box. For more information, see Credentials Vault .
Password	The secure password used to connect to this virtual host.

6. Depending on your choice from step 4, ensure that **Use Rapid Snap for Virtual host-based protection** is selected. (This option is selected by default).
7. Click **Next**.
8. On the *Select Machines* page, select the VMs that you want to protect.

 **CAUTION:** Quest recommends that you limit agentless protection to no more than 200 VMs at once. For example, do not select more than 200 VMs when using the Protect Multiple Machines Wizard. Protecting more than 200 VMs results in slow performance. There is no limit to how many VMs a Core can agentlessly protect over time. For example, you could protect 200 VMs today and another 200 VMs tomorrow.
9. Optionally, if you want to automatically protect new VMs when they are added to the host, select **Auto protect new virtual machines**.
10. Click **Next**.

11. On the *Protection Rules* page, select any of the following options:

Table 61: Hyper-V agentless protection options

Option	Description
Install the Agent on the host on which the virtual disks for VMs are located	Installs the Agent on the same machine as the virtual disks, which may be located on Hyper-V Server, a Hyper-V cluster node, SMB Server, or Scale-Out File Server (SOFS). If necessary, the host restarts automatically after the installation is complete. This option is selected by default.
Upgrade the Agent on the hosts on which virtual disks for VMs are located	If the Agent is located on the same host as the virtual disks, this option upgrades the Agent to the latest release. Virtual disks may be located on Hyper-V Server, a Hyper-V cluster node, SMB Server, or Scale-Out File Server. If necessary, the host restarts automatically after the upgrade is complete. This option is selected by default.
Protect SMB/SOFS server if a disk is located on an SMB or SOFS server	If any of the virtual disks for VMs are located on an SMB Server or Scale-Out File Server, then the server that hosts the share is also protected. This option is selected by default.
Protect machine if it is orphaned by this Core	Lets the Core protect a machine that was previously protected but was then removed from protection because the hypervisor became unreachable. This option is selected by default.
Protect machine if it already has recovery points	Shows existing recovery points alongside the new recovery points after protection. This option is selected by default.
Protect machine agentlessly if it is already protected with the Rapid Recovery Agent	If a Core detects that a machine is already protected by the Rapid Recovery Agent, this option permits duplicate protection (both agentlessly and with the Agent). The protected VM must be powered on and Integration Services must be installed.
Protect machine with both supported and unsupported virtual disk types	The Core does not support VHD or VHD Set virtual disk formats. If selected, only the support virtual disk types will be protected. If not selected, virtual machines with an unsupported disk format will not be protected.
Save Rules	Saves the selected rules to use for future VM agentless protection on this hypervisor host. These rules apply to machines protected automatically or by using the Protect Multiple Machines wizard.

12. When satisfied with the set of rules selected, click **Next**.

13. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
Schedule options are added to the wizard workflow.
14. Proceed with your configuration as follows:
 - If you specified default protection, then click **Next** and skip to step 17 to the *ABM Settings* page.
 - If you specified custom protection, then click **Next** and continue to the next step to configure a protection schedule.
15. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).
16. Optionally, on the *ABM Settings* page, select **Enable Active Block Mapping**, and then complete the following information:

Table 62: Active Block Mapping settings

Option	Description
Enable Active Block Mapping	Lets you enable or disable the ABM feature.
Enable swap file blocks exclusion	Excludes the content of system files, such as pagefile.sys, hyberfill.sys, and swapfile.sys, from the backup.
Exclude subdirectories	<p>Lets you exclude specific files by specifying '<file name>' or '<folder>\<subfolder>\<file name>'.</p> <p>Only the files will be excluded. The folders or subfolders that contained excluded files are included in the mount point, with no contents.</p> <p>i NOTE: This option may affect the performance of the "determining data" phase of transfers.</p>
+ Add	If you opted to exclude subdirectories, click Add and enter the location in the Path table for each item you want to exclude.

For more information, see [Understanding Active Block Mapping](#).

i **NOTE:** Active Block Mapping only supports NTFS file systems. Any type of dynamic disk is not supported.

17. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in step 2 and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in step 2, then click **Next** to see repository and encryption options.
18. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.
19. On the *Encryption* page, do one of the following:
 - If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the **Select encryption key** drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 63: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

- If you want to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**.

i **NOTE:** This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.

20. Click **Finish** to save and apply your settings.


i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).

Protecting multiple machines manually

The protecting procedure requires the following tasks be completed before you begin:

- The Rapid Recovery Agent has been deployed and installed on the machine you want to protect.
- The machine you want to protect was restarted after the Agent installation.
- A repository has been created and is accessible from the Rapid Recovery Core.

Use this procedure to manually specify details for multiple machines that you want to protect simultaneously using the Agent software. The details identify each machine on the network uniquely, and include connection information and credentials. This approach is often used when protecting Linux machines. However, using this process, you can protect only Windows machines, only Linux machines, or a combination of both.

1. From the Rapid Recovery Core Console button bar, click the **Protect** ▾ drop-down menu, and then select  **Protect Multiple Machines**.

The Protect Multiple Machines Wizard opens.

2. On the *Welcome* page, select one of the follow installation options:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you want to define a repository or if you want to establish encryption, select **Advanced (show optional steps)**.

Optionally, if you do not wish to see the *Welcome* page for the Protect Machine Wizard in the future, select the option **Skip this Welcome page the next time the wizard opens**.

3. Click **Next**.
4. On the Connection page of the wizard, from the **Source** drop-down list, select **Manually**.
5. Click **Next**.
6. On the Select Machines page, for each machine you want to protect, enter the machine details in the dialog box. Press **Enter** to separate information for each machine you want to add. Use the format `hostname::username::password::port`. The port setting is optional. The default port for installing Agent on Windows machines is 8006. For Linux machines, the default port is number 22 (SSH port). Examples include:

```
10.255.255.255::administrator::&11@yYz90z
```

```
Linux-host-00-2::administrator::p@$w0rD::22
```

7. Click **Next**.

If the Protection page appears next in the Protect Multiple Machines Wizard, skip to step 11.

If the Agent software is not yet deployed to the machines you want to protect, or if any of the machines you specified cannot be protected for another reason, then the selected machines appear on the Warnings page.

8. Optionally, on the Machines Warnings page, you can verify any machine by selecting the machine and then clicking **Verify** in the toolbar.
9. Optionally, on the Machines Warnings page, select **After Agent installation, restart the machines automatically**.

CAUTION: Quest recommends this option. You must restart agent machines before they can be protected. Restarting ensures that the Agent service is running, and that proper kernel module is used to protect the machine, if relevant.

10. If the status indicates that the machine is reachable, click **Next** to install the Agent software.
The Protection page appears.
11. Optionally, on the *Protection* page, if you want a name other than the IP address to display in the Rapid Recovery Core Console for this protected machine, then in the **Display name** field, type a name in the dialog box.
You can enter up to 64 characters. Do not use the special characters described in the topic [prohibited characters](#). Additionally, do not begin the display name with any of the character combinations described in the topic [prohibited phrases](#).
12. Select the appropriate protection schedule settings as described below:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
With a default protection schedule, the Core will take snapshots of all volumes on the protected machine once every hour. To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the *Summary* page for the specific protected machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
Schedule options are added to the wizard workflow.
13. Proceed with your configuration as follows:
 - If you selected a Typical configuration for the Protect Machine Wizard in [step 2](#) and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - If you selected Advanced configuration for the Protect Machine Wizard in [step 2](#), and default protection, then click **Next** to see repository and encryption options.
 - If you selected a Typical configuration for the Protect Machine Wizard and specified custom protection, then click **Next** and continue to [step 14](#) to choose which volumes to protect.
 - If you selected Advanced configuration for the Protect Machine Wizard in [step 2](#), and default protection, then click **Next** and proceed to [step 16](#) to see repository and encryption options.



14. On the *Protection Volumes* page, select which volumes you want to protect. If any volumes are listed that you do not want to include in protection, click in the **Check** column to clear the selection. Then click **Next**.
 **NOTE:** Typically, it is good practice to protect, at minimum, the System Reserved volume and the volume with the operating system (typically the C drive).
15. On the *Protection Schedule* page, define a custom protection schedule and then click **Next**. For details on defining a custom protection schedule, see [Creating custom protection schedules in Simple Mode](#).
16. On the *Repository* page, select the repository you want to use to store recovery points for this machine, and then click **Next**.
17. On the *Encryption* page, do one of the following:
 - If you want to use encryption keys for data stored in the repository, select **Encrypt the data at rest in a repository**, and then do one of the following:
 - To select an existing encryption key to apply to all new data stored in your repository, select **Encrypt data using Core-based encryption with an existing key**, and from the **Select encryption key** drop-down menu, select the encryption key.
 - To define a new encryption key at this time to apply to all future data stored in your repository, select **Encrypt data using Core-based encryption with a new key**, and then enter information about the key as described in the table below:

Table 64: Define new encryption key

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

- If you want to encrypt all data that transports over a network, select **Encrypt the data in transport over a network**.

 **NOTE:** This option is enabled by default, so if you do not want to encrypt data in this fashion, clear this option.


18. Click **Finish** to save and apply your settings.

i **NOTE:** The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) transfers to the repository indicated in your Rapid Recovery Core following the schedule you defined, unless you specified that the Core should initially pause protection. For information on pausing and resuming protection, see [Pausing and resuming protection](#).

The Rapid Recovery Agent software is deployed to the specified machines, if necessary, and the machines are added to protection on the Core.


Monitoring the protection of multiple machines

You can monitor the progress as Rapid Recovery applies the protection policies and schedules to the machines.

1. In the Rapid Recovery Core Console, navigate to the Rapid Recovery *Home* page and then click  (Events).

The *Events* page displays, broken down by Tasks, Alerts, and Events. As volumes are transferred, the status, start times, and end times display in the Tasks pane.

You can also filter tasks by status (active, waiting, completed, queued, and failed). For more information, see [Viewing tasks](#).

i **NOTE:** To only see tasks that are waiting to be performed, make sure that you select the  (Waiting Tasks) icon.

As each protected machine is added, an alert is logged, which lists whether the operation was successful or if errors were logged. For more information, see [Viewing alerts](#).

For information on viewing all events, see [Viewing a journal of all logged events](#).

Enabling application support

After a VM has been placed under agentless protection, you can support the Exchange or SQL application installed on that machine.

Before you begin, the following prerequisites must be in place.

- **Protect the VM with the Rapid Recovery Core.** The option to enable application support is not available during the protection process. The button to enable this capability is displayed on multiple pages in the UI after the SQL or Exchange machine is placed under protection. For more information, see [Protecting vCenter/ESXi virtual machines using agentless protection](#) or [Protecting Hyper-V virtual machines using host-based protection](#).
- **Enable remote WMI access.** To allow WMI access, you must install and configure Windows Remote Management on the target virtual machine (VM). For more information, see the Microsoft knowledge base article at [https://msdn.microsoft.com/en-us/library/aa384372\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/aa384372(v=vs.85).aspx).
- **Grant administrative rights remotely to local users.** Configure LocalAccountTokenFilterPolicy by running the following administrative command prompt:

```
winrm quickconfig
```

- **Obtain WMI root namespace access authorization.** For the Core to connect to applications without the Agent, the VM must allow incoming network traffic on TCP ports 135 and 445, and to dynamically assigned ports 1024 to 1034.
- **Allow remote access to SQL Server and Exchange.** This prerequisite varies depending on the application you want to support.
 - SQL: Configure the Windows Firewall by opening ports 1433 and 1434.
 - Exchange: Open the TCP and UDP 389 ports.
- **Integrate the guest with the host.** Integration is done by installing the optimization package specific to the hypervisor:
 - For ESXi VMs, use VMware Tools, which are required for agentless ESXi VM protection.
 - For Hyper-V, use the Integration Services bundle, which is preinstalled on most Windows Server operating systems.

Complete the following steps to enable application support for agentlessly protected VMs.

1. From the Core Console, go to the Machines page.
2. Complete one of the following actions:
 - To enable application support for single VM, click the Actions menu for that VM, and then click **Enable application support**.
 - To enable application support for multiple VMs simultaneously, select the VMs, click the **Application Support** drop-down, and then click **Enable application support**.
3. In the Edit Application Support dialog, enter the credentials for the VM.

A green icon is displayed next to the name of the protected machine for which application support is enabled. If you want to add application credentials, you can do so by clicking **SQL** or **Exchange** at the top of the Summary page for the specific machine.

Settings and functions for protected Exchange servers

If you are protecting a Microsoft Exchange Server in your Core, there are additional settings you can configure in the Rapid Recovery Core Console, and there are additional functions you can perform.

A single setting, Enable automatic mountability check, is available in the Core Console related to Exchange Server. If enabled, Exchange server mountability checks are conducted automatically. This setting is available when the status for the protected machine is green (active) or yellow (paused).

For more information, see [About Exchange database mountability checks](#).

You can also perform a mountability check on demand, from the Recovery Points pane on a protected Exchange server machine. For more information, see [Forcing a mountability check of an Exchange database](#).

Following are functions you can perform for an Exchange server protected by the Core.

- **Specify Exchange server credentials.** Rapid Recovery Core lets you set credentials so the Core can authenticate to the Exchange server to obtain information.
For more information about setting credentials for Exchange servers, see [Setting credentials for an Exchange server machine](#).

- **Truncate Exchange logs.** When you force log truncation of Exchange server logs, this process identifies the available space and reclaims space on the protected Exchange server.
For more information about truncating Exchange server logs on demand, see [Forcing log truncation for an Exchange machine](#). This process can also be performed as part of the nightly jobs.
- **Force a mountability check of an Exchange database.** This function checks that Exchange databases are mountable, to detect corruption and alert administrators so that all data on the Exchange server can be recovered successfully.
For more information about forcing a mountability check on demand, see [Forcing a mountability check of an Exchange database](#).
You can also force a mountability check to occur automatically after each snapshot. For more information about mountability checks, see [About Exchange database mountability checks](#).
- **Force a checksum check of Exchange Server recovery points.** This function checks the integrity of recovery points containing Exchange database files.
For more information about forcing a checksum check on demand, see [Forcing a checksum check of Exchange database files](#).


You can truncate Exchange logs and force a checksum check as part of nightly jobs. For more information about the tasks you can schedule as nightly jobs, see [Understanding nightly jobs](#). For information on configuring nightly jobs, see [Configuring nightly jobs for the Core](#).

Setting credentials for an Exchange server machine

In order to set login credentials, an Exchange server must be present on a protected volume. If Rapid Recovery does not detect the presence of an Exchange server, the Set Credentials function does not appear in the Core Console.

Once you protect data on a Microsoft Exchange server, you can set login credentials in the Rapid Recovery Core Console.

Complete the steps in this procedure to set credentials for each Exchange Server.

1. In the left navigation area of the Rapid Recovery Core Console, select the protected Exchange server machine for which you want to set credentials.
The Summary page appears for the protected Exchange server.
2. On the Summary page, from the links at the top of the page, click the downward-facing arrow  to the right of the Exchange menu, and then from the resulting drop-down menu, select Set Credentials.
The Edit Exchange Credentials dialog box for the protected Exchange server appears.
3. In the Edit Exchange Credentials dialog box, enter your credentials as follows:
 - a. In the User name text field, enter the user name for a user with permissions to the Exchange server; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator).
 - b. In the Password text field, enter the password associated with user name you specified to connect to the Exchange server.
 - c. Click OK to confirm the settings and close the dialog box.

Forcing log truncation for an Exchange machine

In order to force log truncation, an Exchange database must be present on a protected volume. If Rapid Recovery does not detect the presence of a database, the log truncation check does not appear in the Core Console.

When you force log truncation for a protected Exchange Server, the size of the logs are reduced. Complete the steps in this procedure to force log truncation on demand.

1. In the left navigation area of the Rapid Recovery Core Console, select the protected Exchange server for which you want to force log truncation
The Summary page for the protected machine appears.
2. At the top of the page, click the Exchange drop-down menu and select Force Log Truncation.
3. In the resulting dialog box, click to confirm that you want to force log truncation.
The dialog box closes. The system starts truncating the Exchange server logs. If Toast alerts are enabled for this type of event, you see a message that the log truncation process starts.

About Exchange database mountability checks

When using Rapid Recovery to back up Microsoft Exchange Servers, mountability checks can be performed on all Exchange databases after every snapshot. This corruption detection feature alerts administrators of potential failures and ensures that all data on the Exchange servers will be recovered successfully in the event of a failure.

To enable or disable this feature, go to the **Settings** menu for a protected machine, and set the **Enable automatic mountability check** option to **Yes** or **No**, respectively. For more information about modifying settings for a protected machine, see [Viewing and modifying protected machine settings](#).

Mountability checks are not part of nightly settings. However, if the automatic mountability check is enabled, and if the Truncate Exchange logs nightly job is enabled, then the mountability check is triggered after the completion of log truncation.

You can also perform a mountability check on demand, from the *Recovery Points* pane on a protected Exchange server machine. For more information, see [Forcing a mountability check of an Exchange database](#).


NOTE: The mountability checks only apply to Microsoft Exchange Server 2007, 2010, 2013 and 2016. Additionally, the Rapid Recovery Agent service account must be assigned the Organizational Administrator role in Exchange.

Forcing a mountability check of an Exchange database

In order to force a mountability check, an Exchange database must be present on a protected volume. If Rapid Recovery does not detect the presence of a database, the mountability check function does not appear in the Core Console.

Complete the steps in this procedure to force the system to perform a mountability check for a specific Exchange server recovery point on demand.

1. In the left navigation area of the Rapid Recovery Core Console, select the protected Exchange server machine for which you want to force the mountability check, and then click the **Recovery Points** menu.
2. Scroll down to the Recovery Points pane.


3. Navigate through the recovery points to find the desired recovery point. Optionally, click the ▶ arrow to the right of a recovery point in the list to expand the view.
In the expanded recovery point information, you can see volumes included in the recovery point.
4. In the Recovery Points pane, from the row representing the correct recovery point, click , and from the drop-down menu, select **Force Mountability Check**.
5. In the resulting dialog box, click to confirm that you want to force a mountability check.
The dialog box closes. The system performs the mountability check. If Toast alerts are enabled for this type of event, you see a message that the mountability check starts.

For instructions on how to view the status of the mountability check, see [Viewing events using tasks, alerts, and journal pages](#).

Forcing a checksum check of Exchange database files

In order to force a checksum check, an Exchange database must be present on a protected volume. If Rapid Recovery does not detect the presence of a database, the checksum check function does not appear in the Core Console.

Complete the steps in this procedure to force the system to perform a checksum check for a specific Exchange server recovery point.

1. In the left navigation area of the Rapid Recovery Core Console, select the protected Exchange server for which you want to force a checksum check, and then click the **Recovery Points** menu.
The Recovery Points page appears for the protected Exchange server.
2. Scroll down to the Recovery Points pane.
3. Navigate through the recovery points to find the desired recovery point. Optionally, click the ▶ arrow to the right of a recovery point in the list to expand the view.
In the expanded recovery point information, you can see volumes included in the recovery point.
4. In the Recovery Points pane, from the row representing the correct recovery point, click , and from the drop-down menu, select **Force Checksum Check**.
5. In the resulting dialog box, click to confirm that you want to force a checksum check.
The dialog box closes. The system performs the checksum check. If Toast alerts are enabled for this type of event, you see a message that the checksum check starts.

For instructions on how to view the status of the checksum check, see [Viewing events using tasks, alerts, and journal pages](#).

Settings and functions for protected SQL servers

If you are protecting a Microsoft SQL Server in your Core, there are additional settings you can configure in the Rapid Recovery Core Console, and there are additional functions you can perform.

A single setting, Attachability, is available in the Core Console related to SQL Server.

Rapid Recovery Core lets you perform a SQL attachability check to verify the integrity of recovery points containing SQL databases. This action checks the consistency of SQL databases and ensures that all supporting MDF (data) and LDF (log) files are available in the backup snapshot.

In previous releases, SQL attachability checks have historically required a licensed version of SQL Server on the Core machine. Rapid Recovery Core now provides the ability to perform SQL attachability checks from an instance of SQL Server on the Core, or from a licensed version of SQL Server on a protected SQL Server machine.

The attachability settings let you specify which licensed version of SQL Server is used to perform this check. For more information about configuring attachability settings, see [Managing Core SQL attachability settings](#).

For more information on SQL attachability, see [About SQL attachability](#).

Following are functions you can perform for a SQL server protected by the Core.

- **Specify SQL Server credentials.** Rapid Recovery Core lets you set credentials so the Core can authenticate to the SQL server to obtain information. You can set credentials for a single protected SQL Server machine, or set default credentials for all protected SQL Servers.
For more information about setting credentials for SQL servers, see [Setting credentials for a SQL Server machine](#).
- **Truncate SQL logs.** When you force log truncation of SQL Server logs, this process identifies the available space on the protected server. This process does not reclaim any space.
For more information about truncating SQL Server logs on demand, see [Forcing log truncation for a SQL machine](#).
- **Force an attachability check of a SQL Server.** This function checks the consistency of SQL databases and ensures that all supporting MDF (data) and LDF (log) files are available in the backup snapshot.
For more information about forcing an attachability check for SQL servers on demand, see [Forcing a SQL Server attachability check](#).

Other than specifying credentials, each of the functions described in the preceding list can be accomplished on demand, and can also be configured to occur as part of the nightly jobs performed for the Core. For more information about the tasks you can schedule as nightly jobs, see [Understanding nightly jobs](#). For information on configuring nightly jobs, see [Configuring nightly jobs for the Core](#).

Setting credentials for a SQL Server machine

You must add the SQL Server machine to protection on the Rapid Recovery Core before performing this procedure. For more information about protecting machines, see [Protecting a machine](#).

Once you protect data on a Microsoft SQL Server machine, you can set login credentials for a single instance, or for all SQL Servers, in the Rapid Recovery Core Console.

Complete the steps in this procedure to set credentials for each SQL Server.

1. In the left navigation area of the Rapid Recovery Core Console, select the protected SQL Server machine for which you want set credentials.
The Summary page displays for the protected SQL Server.

2. On the Summary page, from the links at the top of the page, click the downward-facing arrow ▼ to the right of the SQL menu, and then from the resulting drop-down menu, do one of the following:
If you want to set default credentials for all SQL Server database instances, click Set Default Credentials for All Instances, and in the Edit Default Credentials dialog box, do the following:
 - a. In the User name text field, enter the user name for a user with permissions to all associated SQL servers; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator).
 - b. In the Password text field, enter the password associated with the user name you specified to connect to the SQL server.
 - c. Click OK to confirm the settings and close the dialog box.

If you want to set credentials for a single SQL Server database instance, click the display name of the protected SQL Server machine, and then in the Edit Instance Credentials dialog box, do the following:

- a. Select the credential type (Default, Windows, or SQL)
- b. In the User name text field, enter the user name for a user with permissions to the SQL server; for example, Administrator (or, if the machine is in a domain, [domain name]\Administrator).
- c. In the Password text field, enter the password associated with the user name you specified to connect to the SQL server.
- d. Click OK to confirm the settings and close the dialog box.

Forcing log truncation for a SQL machine

Log truncation is available for machines that use SQL Server. Complete the steps in this procedure to force log truncation.

NOTE: When conducted for a SQL machine, truncation identifies the free space on a disk, but does not reduce the size of the logs.

1. In the left navigation area of the Rapid Recovery Core Console, select the machine for which you want to force log truncation.
The *Summary* page appears for the protected machine.
2. From the *Summary* page (or from any page for this protected machine), at the top of the page, click the **SQL** drop-down menu and select **Force Log Truncation**.
3. Click **Yes** to confirm that you want to force log truncation.

About SQL attachability

The SQL attachability feature lets the Rapid Recovery Core attach SQL master database files (.MDF files) and log database files (.LDF files) to a snapshot of a protected SQL Server. The snapshot is captured using a local instance of Microsoft SQL Server.

Issues relevant for Rapid Recovery users protecting SQL Server machines include which instance of SQL Server performs attachability, and the method of performing SQL attachability (on demand, or as part of nightly jobs).

The attachability check lets the Core verify the consistency of the SQL databases and ensures that all MDF and LDF files are available in the backup snapshot.

Attachability checks can be run on demand for specific recovery points, or as part of a nightly job.

To perform the SQL attachability check on demand, see [Forcing a SQL Server attachability check](#). To perform SQL attachability once daily, at the time specified for your nightly job operations, enable the option Check attachability for

SQL databases in nightly jobs. For more information about setting nightly jobs for the Core, see [Configuring nightly jobs for the Core](#). For more information about setting nightly jobs for a specific machine (in this case, a protected SQL Server), see [Customizing nightly jobs for a protected machine](#).

In previous versions, SQL attachability required a local instance of Microsoft SQL Server to be installed and configured on the Core machine. Rapid Recovery Core now lets you choose to perform the attachability check from a SQL Server instance on the Core, or from a SQL Server instance on a protected SQL Server machine. The instance you select must be a fully licensed version of SQL Server, procured from Microsoft or through a licensed reseller. Microsoft does not allow the use of passive SQL licenses.

Whichever SQL Server instance you specify is then used for all attachability checks. Attachability is synchronized between Core settings and nightly jobs. For example, if you specify using the Core instance of SQL Server for nightly jobs, on-demand attachability checks then also use the Core. Conversely, if you specify using a SQL Server instance on a specific protected machine, all on-demand and nightly attachability checks then use the local instance on the protected machine.

Select the SQL Server instance to use as part of global Core settings. For more information, see [Managing Core SQL attachability settings](#).

NOTE: Performing the attachability check from a protected SQL Server machine requires the Rapid Recovery Agent software to be installed on that server. Agentless protection is not supported for SQL attachability.

Attachability in Rapid Recovery Core supports SQL Server 2008, 2008 R2, 2012, 2014, and 2019. The account used to perform the test must be granted the sysadmin role on the SQL Server instance.

The SQL Server on-disk storage format is the same in both 64-bit and 32-bit environments and attachability works across both versions. A database that is detached from a server instance that is running in one environment can be attached on a server instance that runs in another environment.

NOTE: The version of SQL Server on the Core must be equal to or newer than the SQL Server version on all of the protected machines with SQL Server installed.

Forcing a SQL Server attachability check

In order to force an attachability check, a SQL database must be present on a protected volume. If Rapid Recovery does not detect the presence of a database, the attachability check function does not appear in the Core Console.

Complete the steps in this procedure to force the system to perform an attachability check for a specific SQL server recovery point.

1. In the left navigation area of the Rapid Recovery Core Console, select the protected SQL Server machine for which you want to force the attachability check, and then click the **Recovery Points** menu.
2. Scroll down to the Recovery Points pane.
3. Navigate through the recovery points to find the desired recovery point. Optionally, click the ▶ arrow to the right of a recovery point in the list to expand the view.
In the expanded recovery point information, you can see volumes included in the recovery point.
4. In the Recovery Points pane, from the row representing the correct recovery point, click ⚙️, and from the drop-down menu, select **Force Attachability Check**.
5. In the resulting dialog box, click to confirm that you want to force an attachability check.
The dialog box closes. The system performs the attachability check.

For instructions on how to view the status of the attachability check, see [Viewing events using tasks, alerts, and journal pages](#).

Managing protected machines

This section describes how to view, configure and manage the protected machines in your Rapid Recovery environment.

Topics include:

- [About managing protected machines](#)
- [Viewing protected machines](#)
- [Configuring machine settings](#)
- [Managing machines](#)
- [Understanding custom groups](#)

About managing protected machines

You can manage protected machines from the Rapid Recovery Core Console, including the following tasks:

- You can view protected machines in the Rapid Recovery Core Console using options described in the topic [Viewing protected machines](#).
- You can configure settings for a particular machine, which supersede Core default settings. Some configuration tasks include changing hypervisor host or VM settings, accessing system information, modifying transfer settings, customizing nightly jobs, or configuring notifications for events. For more information, see [Configuring machine settings](#).
- You can remove a machine or cluster from protection, view license information for a protected machine, or diagnose problems by viewing the log file for a protected machine. For more information on these and other tasks, see [Managing machines](#).
- You can view and manage data saved in the Core. For more information, see [Managing snapshots and recovery points](#).

Viewing protected machines

From the Home page on the Rapid Recovery Core Console, when viewing the Summary Tables view, you can see summary information for any machines protected by the Core in the Protected Machines pane.



NOTE: A software agent acts on behalf of the user to take specific actions. Protected machines are sometimes referred to as agents, since they run the Rapid Recovery Agent software to facilitate data backup and replication on the Rapid Recovery Core.

You can view the status, the display name for each machine, which repository it uses, the date and time of the last snapshot, how many recovery points exist in the repository for the machine, and the total amount of storage space the snapshots use in the repository.

To manage aspects of any protected machine, start by navigating to the machine you want to view, configure, or manage. From the Home page, there are three ways to navigate to a protected machine:

- You can click on the IP address or display name of any protected machine from the Protected Machines pane. This takes you to the *Summary* page for the selected protected machine.
- In the left navigation area, you can click on the title of the Protected Machines menu. The Protected Machines page appears. On this page, you can see summary information about each machine. For a detailed description of this page, see [Viewing summary information for a protected machine](#).
- In the left navigation area, under the Protected Machines menu, you can click any protected machine IP address or display name. This takes you to the Summary page for the selected protected machine. For a detailed description of this page, see [Viewing summary information for a protected machine](#).

Viewing cluster summary information

Complete the steps in this procedure to view summary information about a cluster including information about the associated quorum for the cluster.

1. In the Rapid Recovery Core Console, under Protected Machines, click the cluster you want to view. The Summary page for the machine appears.
2. On the Summary page, you can view such information as the cluster name, cluster type, quorum type (if applicable), and the quorum path (if applicable). This page also shows at-a-glance information about the volumes in this cluster, including size and protection schedule. If applicable, you can also view SQL Server or Exchange Server information for a different cluster.
3. To view the most current information, click Refresh.

For information about viewing summary and status information for an individual machine or node in the cluster, see [Viewing protected machines](#).

Configuring machine settings

Once you have protected a machine in your Rapid Recovery Core, you can easily view and modify the settings that govern the behavior of that protected machine. When you modify settings for a specific machine, those settings supersede the behavior set at the Core level.

You can view and configure the following machine settings in the Rapid Recovery Core Console:

- **General.** General machine configuration settings include display name, host name, port, encryption key, repository, and links to a hypervisor host. For information about configuring general settings for a machine, see [Viewing and modifying protected machine settings](#).
- **Credentials.** You can view the current logged-in user name for the protected hypervisor host, and you can update credentials (user name and password) to connect to the host. This setting appears only for vCenter/ESXi or Hyper-V hypervisor hosts.
- **Transfer Queue.** This setting establishes the maximum number of concurrent transfers from a single Hyper-V host. This option appears only for protected Hyper-V hosts.
- **Nightly Jobs.** The subset of Core nightly job settings that appear for a specific protected machine allow you to supersede nightly job settings set at the Core level. This includes rollout, which lets you manage the retention policy. Some settings may differ based on the type of machine that is protected.

- **Transfer.** Settings specific to managing data transfer processes for the selected protected machine. For information about the types of data transfer affected by these settings, see [About modifying transfer settings](#). This setting appears only for machines protected with Agent, or agentlessly protected machines (not hypervisor hosts).
- **ABM.** This setting, when enabled, lets Rapid Recovery Core protect only active blocks in backup snapshots. This option is not available for machines protected with Rapid Recovery Agent. Agentlessly protected machines can use ABM settings of the host or can have customized settings per machine. For more information, see [Understanding Active Block Mapping](#).
- **Excluded Writers.** These settings let you exclude writers. These are machine-specific. A writer is a specific API published by Microsoft to allow other software components to participate in using Microsoft Volume Shadow Services (VSS). Each of the writers in Rapid Recovery that participate in volume snapshots are listed in the Excluded Writers settings. In the event that a writer is interfering with or precluding successful backup transfers, these can be disabled one by one. Quest recommends leaving these settings alone, unless you are otherwise directed by a Quest Data Protection Support representative.
- **License details.** These are details about the license for the specific protected machine. These settings report information from the Core and the Rapid Recovery License Portal. These settings are read-only. To change these settings, update your license information between the Core and the license portal. See your license administrator for details. For more information, see the *Rapid Recovery License Portal User Guide*. These settings are not available for hypervisor hosts.
- **Hyper-V.** This setting, available only on agentlessly protected machines on a Hyper-V server, lets Rapid Recovery Core try to create a VSS snapshot during transfer. If this operation fails, it lets Core create a checkpoint.
- **vSphere.** These settings let Rapid Recovery Core manage some aspects of protected vSphere hosts. One setting lets Core delete user-created VMware snapshots required before capturing Rapid Recovery snapshots). One setting allows transfer of volumes with invalid used capacity. The third setting lets Core take quiesced snapshots.
- **Auto Protection.** This setting, when enabled, results in automatic agentless protection of any new hypervisor guest VMs added to the Hyper-V or vCenter/ESXi host.

The procedure for viewing or changing machine-level settings is identical for general, excluded writers, and license details. For more information, see [Viewing and modifying protected machine settings](#).

The procedure for modifying nightly jobs for a machine is different. For information about configuring nightly job settings for a machine, see [Customizing nightly jobs for a protected machine](#).

The procedure for modifying vSphere settings differs slightly. For more information, see [Configuring vSphere settings](#).

In some cases, you may want to adjust the data transfer rate for a protected machine. For more information, see [About modifying transfer settings](#).

Viewing and modifying protected machine settings

Machine settings help determine the behavior of a machine protected by the Core. When you modify settings for a specific machine, those settings supersede the behavior set at the Core level.

Likewise, a protected Hyper-V virtual host has different machine settings than the virtual machines it manages. For more information, see [Viewing summary information for a hypervisor or cluster host](#).

Complete the steps in this procedure to view and modify general settings, transfer settings, settings for excluded writers, and licensing settings for a protected machine.



NOTE: To view and modify nightly job settings, see [Customizing nightly jobs for a protected machine](#).

1. In the Rapid Recovery Core Console, under the Protected Machines menu, click the IP address or machine name for the machine you want to view or modify.



The *Summary* page for the selected machine displays.

2. Click the **Settings** menu.

The Settings page displays, showing settings for the selected machine. Optionally, to display setting categories from anywhere on the page, click the appropriate hyperlink on the left side of the page.

When you click on a setting you want to change, that setting becomes an editable control. Do one of the following:

- When the control is a drop-down menu, click the downward arrow to list the options, and select the desired option from the menu.
- When the control is a text field, enter a value.
- When the option displays **Yes** or **No**, click the value, which is replaced by a check box. For a setting of Yes, select the check box. For a value of No, clear the check box.
- When the option displays a time value (for example, showing hours, minutes, and seconds), you can click on each component and type a new value or use the up and down arrows to select new values.

For each setting, when satisfied with your changes, click  to confirm and save the change and exit edit mode, or click  to exit edit mode without saving.

3. To modify general settings for a protected machine, click the appropriate setting, and then enter the configuration information as described in the following table.

Table 65: General settings for a protected machine

Text Box	Description
Display name	<p>Enter a display name for the machine.</p> <p>This is the name that displays for a protected machine in the Rapid Recovery Core Console. You can enter up to 64 characters. By default, this is the host name of the machine. You can change this to something more user-friendly if needed. Do not use prohibited characters or prohibited phrases.</p>
Host name	<p>This is the name of the protected machine as it appears in the machine's metadata.</p> <p>i NOTE: Do not change this setting, as doing so could break the connection between the protected machine and the Core.</p>
Repository	<p>This setting only appears for protected Hyper-V clusters, which support shared virtual hard disks. Displays the repository configured on the Rapid Recovery Core in which recovery points of shared virtual hard disks for the protected Hyper-V cluster are stored.</p>
Port	<p>Enter a port number for the machine.</p> <p>The port is used by the Rapid Recovery Core service to communicate with this machine. The default port is 8006.</p>
Encryption key	<p>If you want an encryption key that is already defined for this Rapid Recovery Core to be applied to the data for every volume on this protected machine, you can specify the encryption key here. The key must be unlocked. If no encryption keys exist, you can add an encryption key. For more information on managing encryption keys, see Managing encryption keys.</p> <p>If the volumes on this protected machine are encrypted, you can change to a different encryption key. Alternatively, you can disassociate an encryption key by selecting (none) from the Encryption key drop-down menu.</p> <p>i NOTE: After you apply an encryption key, change an encryption key, or disassociate an encryption key for a protected machine, Rapid Recovery takes a new base image upon the next scheduled or forced snapshot.</p>
Repository	<p>Select a repository for the recovery points.</p> <p>Displays the repository configured on the Rapid Recovery Core in which to store the data from this machine.</p> <p>The repository volume can be local (on storage attached to the Core server), or on a volume on a CIFS shared location.</p> <p>i NOTE: The Repository setting on this page can only be changed if there are no recovery points or if the previous repository is missing.</p>

Text Box	Description
Hypervisor	<p>This setting only appears for virtual machines. The value of this setting indicates whether the selected virtual machine is associated as the child of a protected hypervisor host.</p> <ul style="list-style-type: none"> • If you see the IP address or display name of a protected host, the association exists. This result indicates that the protected machine is not unnecessarily consuming a license from your available license pool. • If you want to remove the association, click the IP address or display name of the hypervisor host, select Not linked from the drop-down menu, and click the check mark to confirm your change. Afterward, this protected VM will consume a license from your pool. • If you see "Not linked," the machine is not currently associated in your Core as a VM on a protected host, and consumes a license from your pool. If you want to save a license, and the host is protected by Agent, you can associate it by clicking Not linked, and selecting the host from the drop-down menu.
OS version	<p>This setting only appears for guest virtual machines associated with a protected hypervisor host. Generally, Rapid Recovery detects and displays the operating system running on the protected VM.</p>

4. To modify nightly job settings for a protected machine, see [Customizing nightly jobs for a protected machine](#).
5. To modify Exchange settings for a protected Exchange server, in the Exchange Server Settings section, click Enable automatic mountability check, and do the following:
 - To enable automatic mountability checks, select the check box, and then click ✓.
 - To disable automatic mountability checks, clear the check box, and then click ✕.

For more information about automatic mountability checks, see [About Exchange database mountability checks](#).

6. To modify transfer settings for a protected machine, click the appropriate setting, and then enter the configuration information as described in the following table.






 **NOTE:** For conceptual information about transfer settings, see [About modifying transfer settings](#).

Table 66: Transfer Settings for a protected machine

Text Box	Description
 Restore Default	This control restores all transfer settings to the system default settings.
Priority	Sets the transfer priority between protected machines. Enables you to assign priority by comparison with other protected machines. Select a number from 1 to 10, with 1 being the highest priority. The default setting establishes a priority of 5.  NOTE: Priority is applied to transfers that are in the queue.
Maximum Concurrent Streams	Sets the maximum number of TCP links that are sent to the Core to be processed in parallel per protected machine, for machines protected in a DVM repository.  NOTE: Quest recommends setting this value to 8. If you experience dropped packets, try increasing this setting.
Maximum Concurrent Writes	Sets the maximum number of simultaneous disk write actions per protected machine connection.  NOTE: Quest recommends setting this to the same value you select for Maximum Concurrent Streams. If you experience packet loss, set slightly lower—for example, if Maximum Current Streams is 8, set this to 7.
Use Core Default Maximum Retries	Select this option to use default retries number for each protected machine, if some of the operations fail to complete.
Maximum Segment Size	Specifies the largest amount of data, in bytes, that a computer can receive in a single TCP segment. The default setting is 4194304. Do not change this setting from the default unless directed to do so by a Quest Support representative.
Maximum Transfer Queue Depth	Specifies the amount of commands that can be sent concurrently. The default setting is 64. You can adjust this to a higher number if your system has a high number of concurrent input/output operations.
Outstanding Reads per Stream	Specifies how many queued read operations will be stored on the back end. This setting helps to control the queuing of protected machines. The default setting is 0.

Text Box	Description
Transfer Data Server Port	Sets the port for transfers. The default setting is 8009.
Transfer Timeout	Specifies in minutes and seconds the amount of time to allow a packet to be static without transfer.
Snapshot Timeout	Specifies in minutes and seconds the maximum time to wait to take a snapshot.
Snapshot Cleaning Timeout	Specifies in minutes and seconds the maximum time for process of deleting VSS snapshot on a protected machine.
Network Read Timeout	Specifies in minutes and seconds the maximum time to wait for a read connection. If the network read cannot be performed in that time, the operation is retried.
Network Write Timeout	Specifies the maximum time in seconds to wait for a write connection. If the network write cannot be performed in that time, the operation is retried.
Encrypt snapshot data	<p>Specifies whether data transported between the protected machine and the Core is encrypted. This option is enabled by default. This setting applies to data in transit over a network. When this option is enabled, all snapshot data transported to a DVM repository are encrypted.</p> <p>i NOTE: Quest recommends setting this to Yes when data between your Core and protected machines must flow over the public or untrusted networks such as the internet.</p>

7. To modify settings for excluded writers, click the appropriate setting, and then select a writer if you want to exclude it.

i **NOTE:** Because the writers that appear in the list are specific to the machine you are configuring, you will not see all writers in your list.

8. To modify the synthetic incremental setting, click Enable synthetic incremental, and do the following:
 - To enable synthetic incremental backups, select the check box, and then click ✓.
 - To disable synthetic incremental backups, clear the check box, and then click ✕.

For more information, see <https://support.quest.com/rapid-recovery/kb/331728/rapid-recovery-synthetic-incremental-feature-overview>.

9. License details for a protected machine are read-only. License detail information is described in the following table.

Table 67: License details for a protected machine


Text Box	Description
Expiration Date	Indicates the expiration date of the license for the selected protected machine.
License Status	Indicates the current status of the license for the selected protected machine.
License Type	Indicates the type of the license for the selected protected machine.
Agent type	Indicates if the current protected machine is a physical or virtual agent.

Related topics:

- [Changing the settings for a Hyper-V host or node](#)
- [Changing the settings for a Hyper-V protected virtual machine](#)
- [Changing the vSphere settings for a VMware protected virtual machine](#)

Changing the settings for a Hyper-V host or node

This procedure applies to Hyper-V hosts or nodes that use Rapid Recovery Rapid Snap for Virtual (agentless protection) to protect virtual machines (VMs).


A Hyper-V host that is using Rapid Snap for Virtual (agentless protection) to protect VMs is indicated in the left navigation area by the host icon .

The settings for a Hyper-V host with VMs that are protected agentlessly are not the same as a typical protected machine. All changes made to the settings for a host apply to the VMs on that host.


1. On the Core Console, under Protected Machines in the left navigation area, click the Hyper-V host whose settings you want to change.
The Summary page for the host opens.
2. On the menu bar for the host, click **Settings**.
The Settings page opens.
3. Under General, click the setting you want to change.
The setting you selected becomes editable, as a text field or a drop-down menu.

4. Enter the configuration information as described in the following table.

Table 68: General settings information

Text Box	Description
Display Name	The name that displays for a protected machine in the Rapid Recovery Core Console. You can enter up to 64 characters. By default, it is the host name of the machine. You can change the display name to something more user-friendly if needed. Do not use prohibited characters or prohibited phrases .
Host Name	The name of the protected machine as it appears in the machine's metadata.  NOTE: Do not change this setting, as doing so could break the connection between the protected machine and the Core.

5. Under Transfer Queue, to change the number of transfer jobs that can occur on the host at one time, click the setting for **Maximum concurrent transfers**.

 **NOTE:** For best performance, it is recommended that the maximum concurrent transfers for the Hyper-V host or node be set to 1, which is the default setting.

6. Under Nightly Jobs, to change the settings for the available nightly jobs, click **Change**. The Nightly Jobs windows appears.
7. Enter the configuration information as described in the following table.


Table 69: Nightly Jobs settings information

Text Box	Description
Clear orphaned registry keys on protected Hyper-V host	Removes the unnecessary files from the registry that result from attaching and detaching virtual disks during data transfers.
Check integrity of recovery points	Conducts an integrity check of each recovery point created for the virtual machines on the Hyper-V host.

8. Click **OK**.
9. Under Auto Protection, to determine whether to automatically protect new virtual machines when they are added to the Hyper-V host, click the setting for **Auto protect new virtual machines**.

Changing the settings for a Hyper-V protected virtual machine

This procedure applies to Hyper-V virtual machines (VMs) that are protected using Rapid Recovery Rapid Snap for Virtual (agentless protection).

A Hyper-V VM that is being protected by Rapid Snap for Virtual (agentless protection) is indicated in the left navigation area by the host icon .

The settings for a Hyper-V agentless VM the same as a typical protected machine with the exception of the Hyper-V section at the bottom of the Settings page. The following task provides instructions for only the Hyper-V section settings. For all other protected machine settings, see [Viewing and modifying protected machine settings](#).

1. On the Rapid Recovery Core Console, in the left navigation area under Protected Machines, click the Hyper-V VM whose settings you want to change.
The Summary page for the VM opens.
2. On the menu bar for the host, click **Settings**.
The Settings page opens.
3. In the list on the left side, click **Hyper-V**.
The setting you selected becomes editable, as a text field or a drop-down menu.
4. Under Hyper-V, click **Snapshot configuration**.
The setting you selected becomes editable a drop-down menu.
5. From the drop-down menu, select one of the options described in the following table.

Table 70: Hyper-V settings information

Text Box	Description
Try to create VSS snapshot during transfer first, if it fails, create a checkpoint	If the VSS snapshot succeeds, the recovery point will be in an application-consistent state. If the VSS snapshot fails and a checkpoint is created, the recovery point will be in a crash-consistent state.
Do not create VSS snapshot during transfer	Generates a recovery point in a crash-consistent state.
Use only VSS snapshots during transfers. If VSS snapshot creation fails, the entire transfer will fail	Generates only application-consistent recovery points. If the VSS snapshot fails, no recovery point is generated.

Changing the vSphere settings for a VMware protected virtual machine

This procedure applies to VMware ESXi or Workstation virtual machines (VMs) that are protected using Rapid Recovery Rapid Snap for Virtual (agentless protection).

The settings for a VMware VM that is protected agentlessly include the same settings that are used for a typical protected machine, with one exception. The vSphere section of the Settings page includes settings that apply only to agentlessly protected VMware VMs. The following task provides instructions for only the vSphere section of the Settings page. For all other protected machine settings, see [Viewing and modifying protected machine settings](#).

1. On the Rapid Recovery Core Console, under Protected Machines in the left navigation area, click the display name of the Hyper-V host you want to change.
The Summary page for the host opens.
2. On the menu bar for the host, click **Settings**.
The Settings page opens.
3. In the list on the left side, click **vSphere**.
The setting you selected becomes editable, as a text field or a drop-down menu.
4. Under vSphere, click the setting that you want to change.
The setting you selected becomes editable, as a text field or a drop-down menu.

5. Enter the configuration information as described in the following table.

Table 71: vSphere settings information

Text Box	Description
Allow Rapid Recovery to delete user created in VMware	The default setting is No .
Allow transfer for volumes with invalid used capacity	The default setting is Yes .
Allow quiesced snapshots	The default setting is Yes .

Understanding Active Block Mapping

Active Block Mapping (ABM) is a patent-pending technology that filters out inactive blocks of data from managed images, thereby letting Rapid Recovery protect only the active blocks, which optimizes function and performance. This feature is only available for agentless (Rapid Snap for Virtual) protection of ESXi or vCenter virtual machines (VMs) and Hyper-V servers and clusters.

ABM delivers a query to the file system header of a volume. The query returns a list of active blocks within the image. For this reason, ABM only works with NTFS file systems. When protecting ESXi and vCenter VMs, ABM can be combined with Changed Block Tracking (CBT), to read only active and changed blocks when taking incremental or differential snapshots.

When configuring agentless protection of a VM with the Protect Multiple Machines wizard, you have the option to enable ABM. If you opt to automatically protect new VMs added to the specified host, then the ABM rule also applies to any new VMs subsequently added to protection on the host.

You can change your ABM choice at any time in the **Settings** page for the host or VM. For more information, see [Changing ABM settings](#).

Changing ABM settings

To change the Active Block Mapping (ABM) settings for a supported hypervisor host or virtual machine (VM), complete the following steps.

i **NOTE:** For more information, see [Understanding Active Block Mapping](#).

1. From the Core Console, in the left navigation area, click the host or VM for which you want to change the ABM settings.
2. From the *Summary* page of the machine, click **Settings**.
3. From the list of machine settings on the left side of the *Settings* page, click **ABM Settings**.
4. Do one of the following:
 - a. If changing ABM settings for an agentlessly protected ESXi or Hyper-V virtual machine, proceed to step 5.
 - b. If changing ABM settings for a supported hypervisor host, skip to step 6.




5. If you are on the *Settings* page of a VM, select whether to use the settings from the host (the default option) or to use the settings for this machine.
 - If you want this VM to use the same ABM settings as the host, select **Use the settings of the Hyper-V host**. If you made any changes, click **Apply**. This procedure is complete.
 - If you want this VM to use different ABM settings than the host, select **Use the settings of this protected machine**, click **Apply**, and then continue to the next step.
6. Review the ABM settings and change them if desired. For each of the following options, click the current setting, select (or clear) the selection, and then click  to accept the change or click  to cancel.

Table 72: Active Block Mapping settings


Option	Description
Enable Active Block Mapping	Lets you enable or disable the ABM feature.
Enable swap file blocks exclusion	Excludes the content of system files, such as pagefile.sys, hyberfill.sys, and swapfile.sys, from the backup.
Exclude subdirectories	<p>Lets you exclude specific files by specifying '<file name>' or '<folder>\<subfolder>\<file name>'.</p> <p>Only the files will be excluded. The folders or subfolders that contained excluded files are included in the mount point, with no contents.</p> <div>  NOTE: This option may affect the performance of the "determining data" phase of transfers. </div>
+ Add	If you opted to exclude subdirectories, click Add and enter the location in the Path table for each item you want to exclude.

About modifying transfer settings

In Rapid Recovery, you can modify the settings to manage the data transfer processes for a protected machine. The transfer settings described in this section are set at the protected machine level. To affect transfer at the Core level, see [Modifying transfer queue settings](#).

There are three types of transfers in Rapid Recovery:

- **Snapshot.** Backs up the data on your protected machine. Two types of snapshots are possible: a base image of all protected data, and an incremental snapshot for data updated since the last snapshot. This type of transfer creates recovery points, which are stored on the repository associated with the Core. For more information, see [Managing snapshots and recovery points](#).
- **Virtual Machine Export.** Creates a virtual machine (VM) from a recovery point, containing all of the data from the backup of the protected machine, as well the operating system and drivers and associated data to ensure the VM is bootable. For more information, see [VM export](#).
- **Restore.** Restores backup information to a protected machine. For more information, see [About restoring volumes from a recovery point](#).


NOTE: The entire volume is always rewritten during restore of Windows systems using EFI system partitions.

Data transfer in Rapid Recovery involves the transmission of a volume of data along a network from protected machines to the Core. In the case of replication, transfer also occurs from the originating or source Core to the target Core.

Data transfer can be optimized for your system through certain performance option settings. These settings control data bandwidth usage during the process of backing up protected machines, performing VM export, or performing a restore. These are some factors that affect data transfer performance:

- Number of concurrent agent data transfers
- Number of concurrent data streams
- Amount of data change on disk
- Available network bandwidth
- Repository disk subsystem performance
- Amount of memory available for data buffering

You can adjust the performance options to best support your business needs and fine-tune the performance based on your environment. For more information, see [Throttling transfer speed](#).

Throttling transfer speed

When transferring backup data or replicated recovery points between protected machines and Cores over the network, you can intentionally reduce the speed of the transfer. This process is known as throttling.

When you throttle the transfer speed, you limit the amount of your network bandwidth dedicated to file transfers from Rapid Recovery. When setting up replication, for example, throttling can reduce the likelihood that the transfer of prior recovery points to the replicated Core consumes all of your network bandwidth.

! CAUTION: Throttling transfer speed is not always required or recommended. This information is provided to provide insight into a potential solution for performance issues in your Rapid Recovery environment. For example, sometimes, throttling may solve issues related to repeated transfer failures or network slowdowns caused by transferring a substantial amount of data for your protected or replicated Cores.

There are several factors involved in determining the best approach to throttling. The type of machine being protected is a key factor. For example, a busy Microsoft Exchange server has a much higher change rate than a seldom-used legacy web server.

The input and output capabilities of the storage volumes on your protected machines can also contribute to more or less efficiency.

The speed of your network is another critical factor, with many variables. The network backbone in place (for example, 1GbE versus 10GbE), architecture, configuration, intentional use of NIC teaming, and even the type of cables used can all affect network transfer speed. If your environment has a slower wide area network, and if transfer jobs fail for backup or replication, consider throttling the transfer speed using some of these settings.

Ultimately, the process of network throttling involves trial and error. Quest recommends that you adjust and test your transfer settings, and revisit these settings periodically to ensure that your settings continue to meet your needs.

Adjusting transfer speed should be accomplished on an individual machine basis. In the Core Console, navigate to a specific machine, select Settings, and adjust the Transfer speed. For specific information about viewing and changing these settings, see [Viewing and modifying protected machine settings](#).

That topic also includes descriptions of each of the settings used for throttling transfer. Those descriptions may be useful in determining which settings you should experiment with first.

The four main settings involved in throttling transfer speed are described in the following table:

Table 73: Transfer speed settings

Machine-Level Setting	Default Setting	Suggested Throttle Setting
Maximum Concurrent Streams	8	4
Maximum Concurrent Writes	8	4
Maximum Segment Size	4194304	2097152
Outstanding Reads per Stream	0	Start at 24

Quest recommends adjusting and testing the other settings prior to changing the default setting for outstanding reads per stream, unless directed otherwise by a Quest Support representative. When tuning and testing this setting, start with a value of 24.


When you specify limitations to protected machine transfer parameters, these limitations apply per job. If two transfer jobs occur simultaneously or overlap, twice the bandwidth is used. If four transfer jobs across the network overlap, four times the bandwidth is used; and so on.


Customizing nightly jobs for a protected machine


Nightly jobs can be configured at the Core level or at the machine level. When nightly jobs are set at the Core level, the changes are applied to all relevant machines protected by that Core. Changes made to the nightly jobs at the machine level supersede the changes made at the Core level, and apply only to the machines specified.

For a list of all nightly jobs, including descriptions and the scope available for each, see the topic [Understanding nightly jobs](#).

Complete the steps in the following procedure to make changes to the nightly jobs for a single protected machine.

1. In the Rapid Recovery Core Console, under the Machines menu, click the IP address or machine name for the machine for which you want to customize nightly jobs.
The Summary page for the selected machine appears.
2. Click the Settings menu.
The Settings page appears, showing configuration settings for the selected machine.
3. Optionally, click the **Nightly Jobs** link to scroll down in the Settings page to view nightly jobs settings.
4. Under the Nightly Jobs heading, click  **Change**.
The Nightly Jobs dialog box appears.
5. In the Nightly Jobs dialog box, select the jobs you want to include to run nightly, or clear the options you want to omit for this machine.

 **NOTE:** Options may vary by machine. For example, a protected machine using Exchange Server may include Check checksum of Exchange databases and Truncate Exchange logs.

 **NOTE:** For information about the Rollup setting, including setting a custom retention policy, see [Customizing retention policy settings for a protected machine](#).

6. Click **OK**.

i **NOTE:** The results of this procedure apply only to the selected protected machine. To apply elsewhere, repeat the procedure for each machine you want to customize. To change the nightly job settings for all machines protected by a Core, see [Configuring nightly jobs for the Core](#).

Understanding system information for a protected machine

You can view system information about each machine protected in your Rapid Recovery Core. To see how to access this information, see [Viewing system information for a protected machine](#)

i **NOTE:** This topic describes system information for a specific protected machine. For information about accessing system information for the Core, see [Viewing system information for the Core](#).

You can see the following information on the *System Information* page for each protected machine. Each of the following elements appears in its own pane, if applicable:

Table 74: Protected machine system information

Topic	Description
System Information	For the protected machine you are viewing, this pane includes the following information: Host name, OS version and architecture, physical memory, display name, domain name, and VM type. If available, click Summary on the bottom of this pane to return to the <i>Summary</i> page for this protected machine.
Volumes	This pane includes the volume name, device ID, file system, formatted capacity, and capacity used.
Processors	This pane includes the architecture, number of microprocessor cores and threads, clock speed, and description.
Network Adapters	This pane includes the physical or virtual network adapter type and speed.
IP Addresses	This pane includes the IP address and family.

Viewing system information for a protected machine

The Rapid Recovery Core Console provides you with easy access to system information about the machines protected on your Core. System information viewed at the machine level includes relevant information about the protected machine. For more information, see [Understanding system information for a protected machine](#).

Complete the steps in this procedure to view detailed system information for a protected machine.

1. Navigate to the Rapid Recovery Core Console, and from the protected machines menu in the left navigation area, click a protected machine name.
The Summary page appears for the selected protected machine.

2. On the Summary page, at the bottom of the Summary pane, click **System Information**.
The *System Information* page appears.

Managing machines

This section describes a variety of tasks you can perform to manage your protected machines managing your machines.

Topics include:

- [Removing a machine](#)
- [Removing a cluster from protection](#)
- [Downloading and viewing the log file for a protected machine](#)
- [Converting a protected cluster node to a protected machine](#)

Removing a machine


When you remove a machine from protection on the Rapid Recovery Core, you are presented with two options: you can keep the recovery points saved thus far to the Core, or you can remove the recovery points.

If you keep the recovery points, you have what is known as a “recovery points only” machine. Restore and mount operations continue to be available from the recovery points captured from the machine when it was under protection. However, backups for the protected machine no longer continue.

If you remove the recovery points, this action deletes any snapshot data for that formerly protected machine from the Rapid Recovery Core.

Complete the steps in the following procedure to remove a machine from protection in your Rapid Recovery environment.

1. From the Rapid Recovery Core Console, in the left navigation pane under Protected Machines, click the machine you want to remove.
2. On the Summary page for the relevant machine, click **Remove Machine**.
3. In the dialog box, if you want to also delete all recovery points for this machine from the repository, select **Remove with recovery points**.
4. To confirm your choice to remove the machine, click **Yes**.

 **CAUTION:** If you delete recovery points, you will no longer be able to restore data for that machine.

Rapid Recovery removes the machine from protection and cancels all active tasks for that machine.

Removing a cluster from protection

Complete the steps in the following procedure to remove a cluster from protection.

1. In the Rapid Recovery Core Console, under Protected Machines, click the cluster you wish to remove.
2. On the Summary page for the cluster, click **Remove Cluster**.

3. Optionally, in the dialog box, to remove all currently stored recovery points for this cluster from the repository, select **Remove with recovery points**.

CAUTION: If you delete recovery points, you will no longer be able to restore data for that cluster.

4. In the dialog box, click **Yes** to confirm.

Removing cluster nodes from protection

Complete the steps in the following procedures to remove cluster nodes from protection.

If you just want to remove a node from the cluster, see [Converting a protected cluster node to a protected machine](#).

1. In the Rapid Recovery Core Console, under Protected Machines, click the cluster node that you want to remove.
2. On the Summary page for the node, click **Remove Machine**.
The *Remove Node* dialog box appears.
3. Optionally, in the dialog box, to remove all currently stored recovery points for this cluster from the repository, select **Remove with recovery points**.

CAUTION: If you delete recovery points, you will no longer be able to restore data for cluster nodes.

4. In the dialog box, click **Yes** to confirm.

Removing all nodes in a cluster from protection

Complete the steps in this procedure to remove all nodes in a cluster from protection.

CAUTION: If you remove all cluster nodes, the cluster is also removed.

1. In the Rapid Recovery Core Console, under Protected Machines, click the cluster whose nodes you want to remove.
2. On the Summary page for the cluster, click **Protected Nodes**.
3. On the Protected Nodes page, select all of the nodes.
4. Click the **Remove Machines** drop-down menu, and then select one of the options described in the following table.

Table 75: Remove Nodes options

Option	Description
Remove and Keep Recovery Points	To keep all currently stored recovery points for this cluster.
Remove Recovery Points	To remove all currently stored recovery points for this cluster from the repository. CAUTION: If you delete recovery points, you will no longer be able to restore data for that cluster.

5. In the *Delete Nodes* dialog box, click **Yes** to confirm.




Viewing license information on a machine

You can view current license status information for a machine protected on the Rapid Recovery Core.

1. From the Rapid Recovery Core Console, under Protected Machines, click the machine that you want to modify.
The *Summary* page for the selected machine appears.
2. Click the **Settings** menu.
The *Settings* page appears, showing configuration settings for the selected machine.
3. Click the **Licensing** link to scroll down to the Licensing Details section of the *Settings* page.
The page refreshes, showing machine-specific licensing settings.

Downloading and viewing the log file for a protected machine

If you encounter any errors or issues with a protected machine, you can download the machine logs to view them or to share them with your Quest Support representative.

1. In the left navigation area of the Core Console, under the Protected Machines menu, click the  arrow if necessary to expand the list of protected machines.
2. For the relevant machine, then click the machine name.
The Summary page for the selected machine appears.
3. From the horizontal context-sensitive menu for the relevant protected machine, click **More** and then select  **Agent Log**.
The Download Agent Log page appears.
4. On the Download Agent Log page, click  **Click here to begin the download**.
5. In the Opening AgentAppRecovery.log dialog box, do one of the following:
To open the log file, select **Open with**, then select an application (such as Notepad) for viewing the text-based log file, and finally click **OK**.
The *AgentAppRecovery.log* file opens in the selected application.
To save the file locally, select **Save File** and then click **OK**.
The *AgentAppRecovery.log* file saves to your Downloads folder. It can be opened using any text editor.

Converting a protected cluster node to a protected machine

In Rapid Recovery, you can convert a protected cluster node to a protected machine so that it is still managed by the Core, but it is no longer part of the cluster. This is helpful, for example, if you need to remove the cluster node from the cluster but still keep it protected.

1. In the Rapid Recovery Core Console, navigate to the cluster that contains the machine you wish to convert, and then click Protected Nodes.
2. On the Protected Nodes page, from the specific node you want to convert, click the Actions drop-down menu and select Convert to Agent.
3. To add the machine back to the cluster, select the machine, and then on the Summary page, from the Actions menu, select Convert to Cluster Node, and then click Yes to confirm the action.

Understanding custom groups

The Rapid Recovery Core Console shows a Protected Machines menu in the left navigation area. This includes all machines or server clusters added to protection on your Rapid Recovery Core. Beneath this, other menus may appear, based on whether you include those objects in your Core. In the same manner, you can create a custom group, which displays as the last menu type in the left navigation area.

The main benefit of a custom group is the ability to group Core objects together in a logical container. This can help you organize and manage Core objects for a specific purpose (for example, by organization, cost center, department, geographical region, and so on).

The act of creating a group always adds one group member (for example, a protected machine or server cluster, a replicated machine, or a recovery points-only machine) to the new custom group. The object added is determined by your origin point when you create the group. Ideally, you would then add additional members to the group. Thereafter, you can perform group actions that apply to all like members of that custom group, as described in [Performing group actions](#).

Custom groups can include protected machines, server clusters, replicated machines, and recovery point-only machines. Server clusters behave the same as protected machines, with the exception that a server cluster and its nodes behave as a single entity. If you attempt to add a node from a server cluster to a group, the entire cluster is added.

A custom group may contain similar or dissimilar members. For groups of similar members, all group actions apply to all members of the group. For example, if you force a snapshot for a custom group of protected machines, each machine will be backed up. For groups with dissimilar members (for example, protected machines and replicated machines), if you apply a group action such as forcing replication, this will only apply to the replicated machines.

You can create one or more groups. A single protected machine or replicated machine can be included in one or more groups. This way, you can group machines on your Core in any way you choose, and can perform actions on that specific group.

Each custom group appears in the left navigation area, with a label you designate. Groups with standard protected machines appear first in the custom group; replicated machines appear below protected machines, as applicable. If there are any recovery point-only machines, these are listed below replicated machines.

In the left navigation area, the objects that are protected on the Core appear each in their own menu. Of these menus, custom groups appear last.

Including a machine in a group does not remove it from its original location. For example, if you have three protected machines called Agent1, Agent2, and Agent3, and you add Agent1 to CustomGroup1, then Agent1 appears in both locations.

Related topics:




- [Creating custom groups](#)
- [Modifying custom group names](#)
- [Removing custom groups](#)

- [Performing group actions](#)
- [Viewing all machines in a custom group on one page](#)



Creating custom groups

When you scroll your cursor over the name of any machine in the Protected Machines or replicated machines menu, you will see an arrow that opens a drop-down menu. From this menu, you can create a custom label.

Use the procedure below to create a custom group.

1. Navigate to the Rapid Recovery Core Console.
2. From the Protected Machines, replicated machines, or recovery points-only menu, do the following:
 - a. Place your cursor over a machine in the menu.
 - b. Click on the  (More) drop-down menu for that machine.
 - c. Scroll down to  **Label as**, scroll right, then click  **New Label**.

The *Create Label* dialog box appears.



3. In the **Name** text box, enter an appropriate label for your custom group.
Use a descriptive name that communicates the purpose of the group. For example, to group protected machines, replicated machines, and recovery point-only machines by department, type *Accounting Department*. You can rename a group later.
-  **NOTE:** Labels must be 50 or fewer characters. You can include a single space between words. You must provide a label for your custom group.
4. When you are satisfied with the label name, click **OK**.
The dialog box closes, and the custom group appears as the last element in the left navigation area.
 5. Optionally, you can add other protected machines, replicated machines, or recovery point-only machines to this group. Navigate to the machine name in the appropriate menu, click its drop-down menu, scroll down and select  **Label as**, and then click the name of the custom group.

You can now perform group actions on this group. For more information, see [Performing group actions](#).


Modifying custom group names


When you modify the name of a custom group, only the label changes. The machine names remain the same.

Use the procedure below to modify a custom group name.

1. Navigate to the Rapid Recovery Core Console.
2. In the Protected Machines menu, scroll your cursor over the custom group you want to modify.
3. Click on the  (More) drop-down menu for that group, and then click  **Edit**.
The Edit Label dialog box appears, within which the name of the custom group becomes editable.

4. In the **Name** field, update the text, or delete the existing label text and type a new label or your custom group.
Use a descriptive name that communicates the purpose of the group. For example, to group protected machines, replicated machines, and recovery point-only machines by geographic region, type **Tokyo**. You can rename a group later.





NOTE: Labels must be 50 or fewer characters. You can include a single space between words. You must provide a label for your custom group.
5. When you are satisfied with the label name, click **OK**.
The dialog box closes, and the modified custom group appears as the last element in the left navigation area.
6. Optionally, you can add other protected machines, replicated machines, or recovery point-only machines to this group. Navigate to the machine name in the appropriate menu, click its drop-down menu, scroll down and select  **Label as**, and then click the name of the custom group.

Removing custom groups

When you remove a custom group, you delete that group from the Protected Machines menu. The machines that were in the group are not removed, and can still be found in the appropriate standard menu.

Use the procedure below to remove a custom group.


1. Navigate to the Rapid Recovery Core Console.
2. In the Protected Machines menu, scroll your cursor over the custom group you want to remove.
3. Click on the  (More) drop-down menu for that group, and then click  **Remove label**.
You see a message asking to confirm the removal of the group.
4. Confirm the removal of the custom group.
The confirmation dialog box closes, and the custom group is removed from the navigation area.

Performing group actions

You can perform group actions on any group appearing in the left navigation area of the Rapid Recovery Core Console. If the group contains dissimilar members (for example, replicated machines and recovery points-only machines), then the actions you request will only be performed on the relevant group members.

Use the procedure below to perform group actions on a custom group.

1. Navigate to the Rapid Recovery Core Console.
2. In the Protected Machines menu, scroll your cursor over the custom group for which you want to perform a group action.

3. Click on the  (More) drop-down menu for that group, and then select an action as follows:
- To force an incremental snapshot or base image for all of the protected machines in the group, click **Force Snapshot** or **Force Base Image**, as appropriate.
 - To pause protection for all of the protected machines in the group, click **Pause Protection** and then specify resumption parameters.
 - To resume protection for all of the protected machines in the group for which protection has been paused, click **Resume Protection** and then confirm that you want to resume.
 - To refresh the information for all of the objects in the group, click **Refresh Metadata**.
 - To pause replication for all replicated machines in this group, under Replication, click **Pause**.
 - To resume replication for all replicated machines in this group for which replication has been paused, under Replication, click **Resume**.
 - To force replication for all replicated machines in this group, under Replication, click **Force**.
 - To remove replication for all replicated machines in this group, under Replication, click **Remove**.
 - To remove recovery points-only machines from this Core and discard the recovery points, under Recovery Points Only, click **Remove Recovery Points**.
 - For custom groups only, to modify the label for the custom group, select **Edit**.
 - For custom groups only, to remove the custom group from the navigation menu, select **Remove label**.

For more information, see the following related topics:

- [Forcing a snapshot](#)
- [Forcing replication](#)
- [Pausing and resuming replication](#)
- [Removing incoming replication from the target Core](#)
- [Modifying custom group names](#)

Viewing all machines in a custom group on one page

Clicking the name of a custom group takes you to a Machines page that lists all the machines in that custom group. You can then perform some functions on all machines from the Actions menu, or you can perform functions individually by selecting commands from each individual machine.

Snapshots and recovery points

This section describes how to use and manage the snapshots and recovery points generated by Rapid Recovery. It includes information about mounting, viewing, and forcing, as well as migrating and deleting recovery points.

Topics include:

- [Managing snapshots and recovery points](#)
- [Viewing the recovery points page of a protected machine](#)
- [Mounting a recovery point](#)
- [Dismounting recovery points](#)
- [Working with Linux recovery points](#)
- [Forcing a snapshot](#)
- [Removing recovery points](#)
- [Deleting an orphaned recovery point chain](#)
- [Migrating recovery points manually to a different repository](#)

Managing snapshots and recovery points

A recovery point is a collection of snapshots taken of individual disk volumes and stored in the repository. Snapshots capture and store the state of a disk volume at a given point in time while the applications that generate the data are still in use. In Rapid Recovery, you can force a snapshot, temporarily pause snapshots, and view lists of current recovery points in the repository as well as delete them if needed. Recovery points are used to restore protected machines or to mount to a local file system.

The snapshots that are captured by Rapid Recovery are done so at the block level and are application aware. This means that all open transactions and rolling transaction logs are completed and caches are flushed to disk before creating the snapshot.

Rapid Recovery uses a low-level volume filter driver, which attaches to the mounted volumes and then tracks all block-level changes for the next impending snapshot. Microsoft Volume Shadow Services (VSS) is used to facilitate application crash consistent snapshots.

Viewing the recovery points page of a protected machine

Complete the steps in the following procedure to view the full list of recovery points for a protected machine. For more information on viewing specific recovery points, see [Viewing recovery points for a machine](#).

i | **NOTE:** If you are protecting data from a DAG or CCR server cluster, the associated recovery points do not appear at the cluster level. They are only visible at the node or machine level.

1. In the Rapid Recovery Core Console, navigate to the protected machine for which you want to view recovery points.

2. From the menu at the top of the page, click **Recovery Points**.

The *Recovery Points* page appears, showing a Recovery Points Summary pane and a Recovery Points pane.







You can view summary information about the recovery points for the machine as described in the following table.


Table 76: Recovery point summary information

Info	Description
Total recovery points	Lists the total number of recovery points saved to the repository for this machine.
Total protected data	Indicates the amount of storage space used in the repository for these recovery points.
Repository	Lists the name of the repository in which these recovery points are stored.
Repository status	Graphically displays the amount of space consumed by the recovery points. Shows percentage of the repository used, the amount of space, and the total space of the repository. Click on the graph to see the amount of space remaining.

You can view information about the recovery points for the machine as described in the following table.

Table 77: Recovery point information

Info	Description
Icon	Graphic depiction of either a recovery point  or, if expanded, a volume  within the recovery point. Recovery points show an arrow  indicating that detail can be expanded (or, if expanded, an arrow  showing that the menu can be contracted).
Encrypted	Indicates if the recovery point is encrypted.
Status	Indicates current status of the recovery point.
Contents and More Information	<p>Lists the volumes included in the recovery point. For Exchange servers, click  to display information about the server.</p> <p>Hover over the  More Information icon to see the space usage and the file system in the recovery point or volume displayed.</p>
Type	Defines a recovery point as either a base image or an incremental (differential) snapshot.
Creation Date	Displays the date when the recovery point was created.
Size	Displays the amount of space that the recovery point consumes in the repository.

Info	Description
	The [More] drop-down menu lets you perform certain functions for the selected recovery point.

3. Optionally, expand a recovery point to view the protected volumes.

Understanding recovery point status indicators

Once a recovery point is captured for a protected SQL or Exchange server, the application displays a corresponding color status indicator in the Recovery Points grid. This grid appears in the Recovery Points pane when viewing recovery points for a specific machine. The color that displays is based on the check settings for the protected machine and the success or failure of those checks, as described in the following tables.



NOTE: For more information on viewing recovery points, see [Viewing the recovery points page of a protected machine](#).

Recovery status point colors for Exchange databases

The following table lists the status indicators that display for Exchange databases.

Table 78: Exchange database status indicators

Status Color	Description
White	Indicates that an Exchange database is not detected within the recovery point, volume, or volume group.
Yellow	Indicates that the Exchange database mountability checks have not yet been run.
Red	Indicates that either the mountability or checksum checks failed on at least one database.
Green	Indicates that the recovery point contains one or more database, and that mountability checks are enabled, and that mountability check passed or that the checksum check passed.

Recovery status point colors for SQL databases

The following table lists the status indicators that display for SQL databases.

Table 79: SQL database status indicators

Status Color	Description
White	Indicates that a SQL database is not detected within the recovery point, volume, or volume group.
Yellow	SQL database was offline, indicating that attachability checks were not possible and have not been performed.

Status Color	Description
Red	Indicates that the attachability check failed, or SQL database is offline.
Green	Indicates that the attachability check passed.



i **NOTE:** Recovery points that do not have an Exchange or SQL database associated with it appear with a white status indicator. In situations where both an Exchange and SQL database exists for the recovery point, the most severe status indicator displays for the recovery point.

Mounting a recovery point

In Rapid Recovery, you can mount a recovery point from the Core Console to access stored data through a local file system.

i **NOTE:** To mount a Linux recovery point with the `local_mount` utility, see [Mounting a recovery point volume on a Linux machine](#).

i **NOTE:** When mounting recovery points from data restored from a machine that has data deduplication enabled, you must also enable deduplication on the Core server.

1. From the Rapid Recovery Core Console, navigate to the machine that you want to mount to a local file system.
The *Summary* page appears for the selected protected machine.
2. From the navigation links at the top of the page, click the **Recovery Points** menu.
The *Recovery Points* page appears for the selected machine.
3. Optionally, in the Recovery Points pane, from the list of recovery points, click the right arrow  symbol to expand the recovery point detail, showing volumes included in the recovery point.
4. In the row for the recovery point that you want to mount, click the  (More) drop-down menu and select **Mount**.
The Mount Wizard appears, displaying the *Volumes* page.
5. On the *Volumes* page, select each volume of the recovery point that you want to mount, and then click **Next**.
The *Mount Options* page of the wizard appears.


6. In the *Mount Options* page, edit the settings for mounting a recovery point as described in the following table.

Table 80: Mount Options settings

Option	Description
Destination and other options Local folder	<p>Choose from one of the following options:</p> <p>Mount to the next available drive letter. This option will provide an alphabetic designation (for example, <code>F:\</code>) for the volume you want to mount, using the next available letter. This option is only accessible if you selected a single volume in step 5.</p> <p>Mount to a drive letter. This option assigns the alphabetic designation you select (for example, <code>Z:\</code>) for the volume you want to mount. The letter must not already be in use. This option is only accessible if you selected a single volume in step 5.</p> <p>Mount to a folder. Specify the path used to access the mounted recovery point. For example, select <code>C:\ProgramData\AppRecovery\MountPoints\MountPoint1</code>.</p>
Mount type	<p>Specify the way to access data for the mounted recovery point:</p> <ul style="list-style-type: none">• Read-only• Read-only with previous writes• Writable
Create a Windows share for this mount	<p>Optionally, select this check box to specify if the mounted recovery point can be shared, and then set access rights to it, including the Share name and Allowed groups.</p>

7. Click **Finish** to mount the recovery point.



i **NOTE:** If you want to copy directories or files from a mounted recovery point to another Windows machine, you can use Windows Explorer to copy them with default permissions or original file access permissions. For details, see [Restoring a directory or file using Windows Explorer](#) or [Restoring a directory or file and preserving permissions using Windows Explorer](#).


8. Optionally, while the task is in process, you can view its progress from the  **Running Tasks** drop-down menu on the Core Console, or you can view detailed information on the Events page. For more information about monitoring Rapid Recovery events, see [Viewing events using tasks, alerts, and journal pages](#).

Dismounting recovery points


Complete the steps in this procedure to dismount recovery points that are mounted on the Core.

i **NOTE:** When dismounting a recovery point mounted remotely, the action is referred to as *disconnecting*.

1. In the Rapid Recovery Core Console, from the icon bar, click  (More) and then select  **Mounts**. The *Mounts* page appears. There is a pane for Local Mounts (recovery points mounted from the Core) and another for Remote Mounts (recovery points mounted using the Local Mount Utility). In each pane, the respective mounted recovery points appears in a list.

2. To dismount local mounts, in the Local Mounts pane, do the following:
 - a. Select the local mount point or points you want to dismount.
 - To dismount all recovery points, click the check box in the title bar of the Local Mounts table. All mounts are selected.
 - To dismount one or more recovery points, click the check box in the first column of each row representing the mount point you want to disconnect.
 - b. Click  **Dismount**.
A confirmation dialog box appears.
 - c. Confirm that you want to dismount the selected recover points.
The local recovery points dismount.

i

NOTE: If toast alerts are enabled, you may see a message that the appropriate mount points are being dismounted.
3. To disconnect recovery points mounted remotely, in the Remote Mounts pane, do the following:
 - a. Select the remote mount point or points you want to disconnect.
 - To disconnect all recovery points, click the check box in the title bar of the Remote Mounts table to select all mount points.
 - To disconnect one or more recovery points, click the check box in the first column of each row representing the mount point you want to disconnect.
 - b. Click  **Disconnect**.
A confirmation dialog box appears.
 - c. Confirm that you want to disconnect the selected recover points.
The local recovery points disconnected.

i

NOTE: If toast alerts are enabled, you may see a message that the appropriate mount points are being dismounted.
4. Confirm that the previously mounted recovery points no longer appear in the Local Mounts or Remote Mounts list, as appropriate.

Working with Linux recovery points

The recommended and supported method to mount and unmount recovery points from a protected Linux machine is to use the `local_mount` utility.

- [Mounting a recovery point volume on a Linux machine](#)
- [Unmounting a recovery point on a Linux machine](#)

The procedures referenced above specifically address using `local_mount` to mount and unmount Linux recovery points.

i

NOTE: For managing Linux recovery points from the Rapid Recovery Core Console, see [Managing snapshots and recovery points](#).

Mounting a recovery point volume on a Linux machine

Using the `local_mount` utility in Rapid Recovery, you can remotely mount a volume from a recovery point as a local volume on a Linux machine.

i **NOTE:** When performing this procedure, do not attempt to mount recovery points to the `/tmp` folder, which contains the `rapidrecovery-vdisk` files.

1. Create a new directory for mounting the recovery point (for example, you can use the `mkdir` command).
2. Verify the directory exists (for example, by using the `ls` command).
3. Run the Rapid Recovery `local_mount` utility as root, or as the super user, for example:

```
sudo local_mount
```

4. At the Rapid Recovery mount prompt, enter the following command to list the protected machines.

```
lm
```

5. When prompted, enter the IP address or hostname of your Rapid Recovery Core server.
6. Enter the logon credentials for the Core server, that is, the user name and password.
A list of the machines that are protected by the Rapid Recovery server displays. Each machine is identified by the following: line item number, host/IP address, and an ID number for the machine.

For example: 7d658e5f-fa08-4600-95f0-5f486bc1b6a4#de0896fd-571a-4cc5-aeed-264d2c3c72f4#f377e145-dd4d-3ac3-5b15-37ce8f4913ba

7. Enter the following command to list the recovery points that are available for a specified machine:

```
lr <line_number_of_machine>
```

i **NOTE:** You can also enter the machine ID number in this command instead of the line item number.

A list of the base and incremental recovery points for the machine appears. The list includes the line item number, date and timestamp, location of volume, size of recovery point, and an ID number for the volume, which includes a sequence number at the end to identify the recovery point.

For example, 7d658e5f-fa08-4600-95f0-5f486bc1b6a4#de0896fd-571a-4cc5-aeed-264d2c3c72f4#f377e145-dd4d-3ac3-5b15-37ce8f4913ba:2

8. Enter the following command to select and mount the specified recovery point at the specified mount point/path.

```
m <volume_recovery_point_ID_number> <volume-letter> [flag] <path>
```

The flag in the command determines how to mount the recovery point. You can use one of the following options:

[r] - mount read-only (default). This flag lets you mount a recovery point but does not let you make changes to it.

[w] - mount writable. This flag lets you mount the recovery point and lets you make changes.

[v] - mount with previous writes. Mounting with the "v" flag lets you mount the recovery point and include any changes that were made during the previous writable mount but are not present in the recovery point.

[n] - do not mount nbd to <path>. A nbd (network block device) makes a socket connection between the Core and the protected machine when you perform a local mount. This flag lets you mount the recovery point without mounting the nbd, which is useful if you want to manually check the file system of the recovery point.

i **NOTE:** You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. In that case, you would use the machine line number (from the `lm` output), followed by the recovery point line number and volume letter, followed by the path, such as, `m <machine_line_number> <recovery_point_line_number> <volume_letter> <path>`. For example, if the `lm` output lists three protected machines, and you enter the `lr` command for number 2 and you mount the twenty-third recovery point volume b to `/tmp/mount_dir`, then the command would be: `m 2 23 b /tmp/mount_dir`.

i **NOTE:** If you are mounting a BTRFS volume from a compatible operating system (see the "Rapid Recovery release operating system installation and compatibility matrix" topic in the *Rapid Recovery System Requirements Guide*), then you must include the following parameter:

```
mount -o nodatasum,device=/dev/xxx /dev/xxx /mnt/yyy
```

9. To verify that the mount was successful, enter the following command, which should list the attached remote volume:

```
l
```

Unmounting a recovery point on a Linux machine

Complete the steps in this procedure to unmount a recovery point on a Linux machine.

1. Run the Rapid Recovery `local_mount` utility as root, or as the super user, for example:

```
sudo local_mount
```

2. At the Rapid Recovery mount prompt, enter the following command to list the protected machines.

```
lm
```

3. When prompted, enter the IP address or hostname of your Rapid Recovery Core server.

4. Enter the logon credentials (user name and password) for the Core server.

A list of the machines that are protected by the Rapid Recovery Core server displays.

5. Enter the following command to list the recovery points that are available for a specified machine:

```
lr <line_number_of_machine>
```

Note: Note that you can also enter the machine ID number in this command instead of the line item number. A list of the base and incremental recovery points for the machine will display and includes. The list includes the line item number, date and timestamp, location of volume, size of recovery point, and an ID number for the volume that includes a sequence number at the end, which identifies the recovery point.

For example: 7d658e5f-fa08-4600-95f0-5f486bc1b6a4#de0896fd-571a-4cc5-aeed-264d2c3c72f4#f377e145-dd4d-3ac3-5b15-37ce8f4913ba:2

6. Run the `l` or `list` command to obtain a list of mounted Network Block Device (NBD)-devices. If you mount any recovery point, you will get a path to NBD-device after executing the `l` or `list` command.
7. Enter the following command to unmount a recovery point.

```
umount <path_of_nbd-device>
```

8. Run the `l` or `list` command to verify that the unmount of the recovery point was successful.

Forcing a snapshot

Forcing a snapshot lets you force a data transfer for the current protected machine. When you force a snapshot, the transfer starts immediately or is added to the queue if other jobs are running.

You can choose from two types of snapshots.

If you select an incremental snapshot and there is no previous recovery point, a base image is captured. Forcing a snapshot does not change the timing for any schedules snapshots.

- A base image is a snapshot of all data on the selected volumes of the machine.
 - An incremental snapshot captures all data that has been changed since the last snapshot.
1. In the Rapid Recovery Core Console, navigate to the machine or cluster with the recovery point for which you want to force a snapshot.
 2. On the Summary page, in the Summary pane, click **Force Snapshot**.
The Force Snapshot dialog appears.
 3. In the Force Snapshot dialog box, in the check box, click one or more volumes or protection groups.
 4. Click **Force Snapshot** or **Force Base Image**, respectively.
 5. If you selected a base image, click to confirm that you want to take a base image.
A base image could take a substantial amount of time, based on the amount of data in the volumes you want to back up.
The snapshot you selected is queued and begins as soon as other jobs have completed.

Removing recovery points

You can easily remove recovery points for a particular machine from the repository. When you delete recovery points in Rapid Recovery, you can specify one of the following options.

- **Delete All Recovery Points.** Removes all recovery points for the selected protected machine from the Repository.
- **Delete a Range of Recovery Points.** Removes all recovery points in a specified range before the current, up to and including the base image, which is all data on the machine as well as all recovery points after the current until the next base image.

i | **NOTE:** You cannot recover the recovery points you have deleted. If you need the data stored in the recovery points, considering archiving the data first.

1. In the Rapid Recovery Core Console, under the **Protected Machines** menu, click the name or IP address of the machine for which you want to view and remove recovery points.
The Summary view for the selected protected machine appears.
2. Next to the machine name or IP address, click the Recovery Points menu.
The *Recovery Points* page for the selected machine appears.
3. Scroll down to the Recovery Points pane.
Options appear under the pane title, including Refresh, Delete Range, and Delete All.
4. To delete all currently stored recovery points, under the Recovery Points pane title, click **Delete All**, and in the confirmation dialog box, click to confirm deletion.
5. To delete a set of recovery points in a specific data range, do the following:
 - a. Under the Recovery Points pane title, click **Delete Range**.
The *Delete Recovery Points Within Range* dialog box appears.
 - b. In the *Delete Recovery Points Within Range* dialog box, in the **From** field, select the date and time from which you want to start deleting recovery points.
 - c. In the **To** field, select the date and time defining the last recovery point you want to delete.
 - d. Click **Delete**.
 - e. In the confirmation dialog box, click to confirm deletion.

Deleting an orphaned recovery point chain

An orphaned recovery point is an incremental snapshot that is not associated with a base image. Subsequent snapshots continue to build onto this recovery point; however, without the base image, the resulting recovery points are incomplete and are unlikely to contain the data necessary to complete a recovery. These recovery points are considered to be part of the orphaned recovery point chain. If this situation occurs, the best solution is to delete the chain and create a new base image.

For more information about forcing a base image, see [Forcing a snapshot](#).

1. In the Rapid Recovery Core Console, navigate to the protected machine for which you want to delete the orphaned recovery point chain.
2. From the menu at the top of the page, click **Recovery Points**.
3. In the Recovery Points pane, expand the orphaned recovery point.
This recovery point is labeled in the **Type** column as "Incremental, Orphaned."

4. Next to Actions, click **Delete**.
The *Delete Recovery Points* window appears.
5. In the *Delete Recovery Points* window, click **Yes**.

CAUTION: Deleting this recovery point deletes the entire chain of recovery points, including any incremental recovery points that occur before or after it, until the next base image. This operation cannot be undone.

The orphaned recovery point chain is deleted.

Migrating recovery points manually to a different repository

If you want to remove the recovery points of a protected machine from a repository without deleting them, you can migrate them to a different repository manually by using this procedure. This process involves archiving recovery points from the source repository, and then importing the archive into the target repository.

For example, you can perform this procedure if your existing repository is full, or if your needs change and you want to protect a machine using a different Core and repository.

CAUTION: If your repository was upgraded previously from AppAssure 5.3 or 5.4 and used replication, Quest recommends performing the Check Repository Job on each repository in that target Core before migration. Performing this job will preclude copying any data irregularities to the new destination repository. The Check Repository Job is only available in the UI if it is applicable to your Core, and could take a substantial amount of time to run.


1. In the Rapid Recovery Core Console, pause protection for any protected machines that have recovery points you want to migrate. For more information, see [Pausing and resuming protection](#).
2. Cancel all current operations for any protected machines that have recovery points you want to migrate, or wait for them all to complete.
3. Archive the recovery points for the machine or machines you paused. For more information, see [Creating an archive](#).
4. After archiving and verifying the archive, remove the existing recovery points for the protected machine you want to migrate. For more information, see [Removing recovery points](#).

NOTE: Without removing existing recovery points, you cannot change repositories for a protected machine.

5. Create a new repository for the migrated recovery points, or ensure a new destination repository exists. For more information, see [Creating a DVM repository](#).
If you want to use an existing repository, continue to [step 6](#).

6. Change the repository for each machine that you paused by completing the following steps:
 - a. On the Core Console, click the protected machine in the navigation tree.
 - b. On the *Summary* page of the protected machine, click **Settings**.
 - c. On the *Settings* page, in the General pane, click the **Repository** drop-down list, and then select the name of the repository you created in [step 4](#).

If you want to use an existing repository, select the name of an existing repository.

i **NOTE:** When migrating recovery points to an existing repository, ensure that the existing repository has enough free space to contain the migrated recovery points.
 - d. Click  to save the change to settings.
7. Resume protection for the machine or machines that you paused. For more information, see [Pausing and resuming protection](#).
8. Take a new base image for each of the protected machines you moved. For more information, see [Forcing a snapshot](#) and use the **Force Base Image** option.
9. Import the archived data for the machines you want to migrate. For more information, see [Importing an archive](#).

Managing privacy

This section describes the personal information that Rapid Recovery can collect, what that information is used for, and how you can control the privacy of that data.

Topics include:

- [General Data Protection Regulation compliance](#)
- [How Rapid Recovery uses personal information](#)
- [Non-phone-home license restrictions](#)
- [Obtaining and using non-phone-home licenses](#)

General Data Protection Regulation compliance

The General Data Protection Regulation (GDPR) is legislation crafted to strengthen and unify data protection for all individuals within the European Union (EU). It also addresses the export of personal data outside the EU, which makes it relevant to software manufacture in the US and other countries. It updates rules governing the handling of individuals' personal data. GDPR is being widely adopted throughout the software industry.

To comply with the GDPR, the collection of any personally identifiable information (PII) by Rapid Recovery has been carefully considered. Data collection has been streamlined, and the information collected and how it is used is clearly documented.

When installing the Rapid Recovery Core or running the Rapid Recovery Info Gathering Tool, you are provided a description of the information Rapid Recovery collects and our purposes for collecting the information.

If you accept the stated use of personal data, you can then associate a license (running in standard "phone-home" mode) with your Core. If you choose to decline the use of personal data described in the privacy policy, you must request a special "non-phone-home" license. After you receive that license and associate it with your Core, your PII will not be used, and certain functions (auto update, and enabling integration between the Core and the QorePortal) are disabled.

Regardless of the privacy option you selected during installation, from the Core General setting *Agree to use of personal data*, you can change this setting. To switch between phone-home and non-phone-home modes in either direction, you must have access to the appropriate license.

For more information about the GDPR, see the EU General Data Protection Regulation website at <https://eugdpr.org/the-regulation/>.

For more information about managing your privacy, see the following topics in the *Rapid Recovery User Guide*:

- Certain business rules apply when changing between phone-home and non-phone-home mode using the *Agree to use of personal data* general setting. For more information, see the topic [Configuring Core general settings](#).

- To see what information Rapid Recovery collects, in which circumstances, and why the information is collected, see [How Rapid Recovery uses personal information](#).
- To see what functions you cannot perform when using a non-phone-home license, see the topic [Non-phone-home license restrictions](#).
- To download a phone-home license, log into the Rapid Recovery License Portal. From the navigation menu, click **Licensing**, and from the drop-down menu on the top right, select **License Key**.
- To learn how to obtain a license in non-phone-home mode, see the topic [Obtaining and using non-phone-home licenses](#).

How Rapid Recovery uses personal information

As detailed in the *Rapid Recovery Installation and Upgrade Guide* topic "Understanding Rapid Recovery licenses," Rapid Recovery uses three types of software licenses: subscription, perpetual, and trial licenses.

There are two modes in which licenses can be used:

- **Phone-home mode.** All licenses are issued in phone-home mode unless otherwise requested. If you register a phone-home license with the Rapid Recovery License Portal, Rapid Recovery collects some personally identifiable information (PII). The information it collects, and how the information is used, are described below.
- **Non-phone-home mode.** Even when connected to the internet, if you obtain and register a non-phone-home license, Rapid Recovery will not share your PII. Using this mode precludes you from performing certain functions, as described in the topic [Obtaining and using non-phone-home licenses](#).

Subscription licenses can only be used in phone-home mode. Perpetual and trial licenses can be used in either phone-home or non-phone-home mode.

If you register a phone-home license, you give Rapid Recovery permission to collect the following PII:

- The IP addresses and hostnames of hosts that run on or interact with Rapid Recovery Core and Agent.
- The email addresses associated with Rapid Recovery licenses; and
- The consumption of licenses against the amount of licenses in the license pool.

This information is sent to Quest Software Inc. for the following purposes:

- To properly apply the appropriate license terms for the product;
- To provide customer support; for example, when you run the Info-Gathering Tool, logs and diagnostic data you specify is gathered to a local folder to send to Quest, or may be uploaded to Amazon, accessible only to Quest Data Protection Support.
- To notify users of available updates (if the Update settings on your Core specify an option other than **Never check for updates**); and
- To allow communication between your Core and the QorePortal. This communication can be enabled or disabled using the QorePortal setting on the Core. This portal lets licensed users with a current Support contract monitor the health of Cores and protected machines, manage multiple Cores, and generate reports on demand for the relevant Cores and protected machines.

You have the right to choose whether or not to share this information with Quest. First, when installing or updating the Rapid Recovery Core, you can choose whether to share this information on the Privacy Policy page of the

installer. Also, if you decide not to share information with Quest, you can change the Core general setting **Agree to use of personal data**. Changing this setting prompts you to enter a license. If you enter a non-phone-home license, auto-update is disabled, as is your connection to the QorePortal.

For more resources on this topic, see the following related links.

- For more information about the functions you cannot perform when using non-phone-home mode, see the topic [Non-phone-home license restrictions](#).
- For more information about obtaining a license that uses non-phone-home mode, contact the Quest licensing team by web form, as described in the topic [Obtaining and using non-phone-home licenses](#).
- For more information about changing your general settings, including the sharing of PII, see the topic [Configuring Core general settings](#).
- For more information about viewing licensing information for a single protected machine, see [Viewing license information on a machine](#).
- For more information about the types of licenses available, see the *Rapid Recovery Installation and Upgrade Guide* topic "Managing licenses" and its subtopics. For example, for more information about updating license key or file information, see "Updating or changing a license." To see how to add a license to an appliance, see "Adding a license." For more information about contacting the license portal server, see "Contacting the Rapid Recovery License Portal Server." In this release, these topics all appear only in the *Rapid Recovery Installation and Upgrade Guide*.
- For additional topics regarding the management of licenses, see the *Rapid Recovery License Portal User Guide*.

Non-phone-home license restrictions

Registering a Rapid Recovery license in non-phone-home mode precludes the Core from sharing your personal information. This includes email address, IP addresses, and license consumption information.

After you register a non-phone-home license, you will not be able to do the following:

- View license server information from the *Settings* page on the Core Console (since your Core is not permitted to communicate with the Rapid Recovery License Portal).
- Manage the consumption of licenses from the Core Console.
- Submit information to Quest Data Protection Support from the Rapid Recovery Info Gathering Tool.
- Monitor the health of Cores and protected machines, manage multiple Cores, and generate reports on demand for the relevant machines on multiple Cores from the QorePortal.
- Use the auto update feature to update a new version of Rapid Recovery Core (your Core is not notified that new versions are available).
- Use the auto update feature to directly update protected Linux machines using package managers such as yum, zypper, or apt (you can still download an installation package from an internet-accessible Linux machine and move installation files to the secured computer manually).

Obtaining and using non-phone-home licenses

If you obtain the non-phone-home license before you upgrade or install Rapid Recovery Core, transfer the license to the Core server. When you run the installer, on the Privacy Policy page, select the option to decline to share data, and when prompted, register the non-phone-home license.

If your Rapid Recovery Core is already registered with a phone-home key, access the **General** settings on the Core, change the setting **Agree to use of personal data** to **No**, and when prompted, register the non-phone-home key.


For more information or step-by-step instructions for changing General settings for your Core, see the topic [Configuring Core general settings](#).

For more information about managing licenses from the Rapid Recovery Core, see the topic "Understanding Rapid Recovery licenses" in the *Rapid Recovery Installation and Upgrade Guide*.

Complete the steps in this procedure to contact the Quest licensing team to obtain a non-phone-home license.

1. In a web browser, navigate to the Quest Licensing Assistance website at <https://support.quest.com/contact-us/licensing>.
2. From the **How can we help you** drop-down menu, select **Obtain a license for my product**.
3. From the **Select Product** drop-down menu, select **Rapid Recovery**.
4. From the **Product Version** drop-down menu, select the appropriate option.
For example, select 6.7.

5. In the Contact Information section of the form, add information as described in the following table.

Field Name	Description	Required Field
Business Email	Enter the email address to which you want the Quest licensing team to respond. If you have access to the email account associated with your Rapid Recovery license, use that address for fastest response.	Yes
Contact First Name	Enter your first name.	Yes
Contact Last Name	Enter your last name.	Yes
Company Name	Enter the name of the company associated with your Rapid Recovery license.	
US Federal	Select if your license is related to a US federal organization.	No
Country	Select your country.	Yes
Phone Number	Enter your phone number, including area code. If outside the US, include country code.	Yes
License Number (if available)		No
License Key (if available)	License keys were used in AppAssure 5.4.1 and earlier. This is typically a string of 30 characters (6 groups of 5 numbers and upper-case alphabetic letters, separated by hyphens).	No
Machine ID	The name of the registered Core machine.	No
Service Tag (if available)	Enter the service tag if available.	No
License Request Details	<p>Indicate in this field that, per GDPR, you want a non-phone-home license to replace your phone-home license to protect your PII.</p> <p> NOTE: By making this request, you agree that you will delete the phone-home license key when you receive and register the non-phone-home key. You also agree that you will not share this key.</p>	Yes
License File	If you have a phone-home license, you can attach the license file.	No

6. To validate your request and submit the form, select **I'm not a robot**, and then click **Next**.
The form is submitted, and you receive an email message with a Service Request (SR) number.

When you receive the non-phone-home license file, upload it to the Core server to register the license. For more information, see "Updating or changing a license" in the *Rapid Recovery Installation and Upgrade Guide*.

Encryption

This section describes the process of securing data in your environment using encryption keys and machine-level snapshot encryption settings.

Topics include:

- [Understanding encryption keys](#)
- [Encrypting data in transport over a network](#)
- [Applying or removing encryption keys](#)
- [Managing encryption keys](#)

Understanding encryption keys

The Rapid Recovery Core can encrypt snapshot data for all volumes within any repository using encryption keys that you define and manage from the Core Console.

Instead of encrypting the entire repository, Rapid Recovery lets you specify an encryption key for one or more machines protected on a single Rapid Recovery Core. Each active encryption key creates an encryption domain. There is no limit to the number of encryption keys you can create on the Core.

In a multi-tenant environment (when a single Core hosts multiple encryption domains), data is partitioned and deduplicated within each encryption domain. As a result, Quest recommends using a single encryption key for multiple protected machines if you want to maximize the benefits of deduplication among a set of protected machines.

You can also share encryption keys between Cores using one of three methods. One method is to export an encryption key as a file from one Rapid Recovery Core and import it to another Core. A second method is to archive data secured with an encryption key, and then import that archived data into another Rapid Recovery Core. The third method is to replicate recovery points from a protected machine using an encryption key. After you replicate protected machines, the encryption keys used in the source Core appear as replicated encryption keys in the target Core.

In all cases, once imported, any encryption key appears in the Core with a state of Locked. To access data from a locked encryption key, you must unlock it. For information about importing, exporting, locking or unlocking encryption keys, see the topic [Managing encryption keys](#).

Key security concepts and considerations include:

- Encryption is performed using 256-bit Advanced Encryption Standard (AES) in Cipher Block Chaining (CBC) mode that is compliant with SHA-3.
- Deduplication operates within an encryption domain to ensure privacy.
- Encryption is performed without impact on performance.
- You can apply a single encryption key to any number of protected machines, but any protected machine can only have one encryption key applied at a time.

- You can add, remove, import, export, modify, and delete encryption keys that are configured on the Rapid Recovery Core.

CAUTION: Rapid Recovery takes a new snapshot whenever you apply an encryption key to a protected machine. A new snapshot is also triggered after you disassociate an encryption key for a protected machine.

Encryption keys generated from the Rapid Recovery Core are text files that contain four parameters, as described in the following table:

Table 81: Components of an encryption key

Component	Description
Name	This value is equivalent to the key name given when adding a key in the Rapid Recovery Core Console.
Key	This parameter consists of 107 randomly generated English alphabetic, numeric, and mathematical operator characters.
ID	The key ID consists of 26 randomly generated upper-case and lower-case English characters.
Comment	The comment contains the text of the key description entered when the key was created.

Encrypting data in transport over a network

Rapid Recovery Core includes an encryption feature. You can encrypt all data in transport over a network. Quest recommends enabling this encryption setting when data between your Core and protected machines (or between two Cores such as for replication) must flow over the public or untrusted networks such as the internet.

While there is only a small performance cost involved in enabling this encryption, if your Cores and protected machines are within the confines of a private local area network, you can disable this option with confidence.

Please read the following information and adjust your environment accordingly.

By default, when protecting a machine using the Protect Machine wizard or the Protect Multiple Machines wizard, encryption for the data in transport over a network is enabled. If you select advanced options for the wizard, you can view the Encryption options. On the *Encryption* page of the wizard, if preferred, you can clear the option **Encrypt the data in transport over a network**.

NOTE: If you do not select Advanced options in the wizard, encryption for data in transport is enabled nevertheless.

After completing the relevant protection wizard, you can always enable or disable encryption for snapshot data by changing transfer settings at the machine level. Select the protected machine, click **Settings**, and under Transfer settings, for the setting **Encrypt snapshot data**, select **Yes** to enable encryption or select **No** to disable encryption during transport. For specific details, see [Viewing and modifying protected machine settings](#).

Applying or removing encryption keys

You can secure the data protected on your Core at any time by defining an encryption key and applying it to one or more protected machines in your repository. You can apply a single encryption key to any number of protected machines, but any protected machine can only use one encryption key at a time.

The scope of deduplication in Rapid Recovery is limited to protected machines using the same repository and encryption key. Therefore, to maximize the value of deduplication, Quest recommends applying a single encryption key to as many protected machines as is practical. However, there is no limit to the number of encryption keys you can create on the Core. Thus, if legal compliance, security rules, privacy policies, or other circumstances require it, you can add and manage any number of encryption keys. You could then apply each key to only one protected machine, or any set of machines in your repository.

Any time you apply an encryption key to a protected machine, or dissociate an encryption key from a protected machine, Rapid Recovery takes a new base image for that machine upon the next scheduled or forced snapshot. The data stored in that base image (and all subsequent incremental snapshots taken while an encryption key is applied) is protected by a 256-bit advanced encryption standard. There are no known methods for compromising this method of encryption.

If you change the name or passphrase for an existing encryption key currently used for a protected machine, then upon the next scheduled or forced snapshot, Rapid Recovery Core captures and reflects the updated properties of the key. The data stored in that image (and all subsequent incremental snapshots taken while an encryption key is applied) is protected by a 256-bit advanced encryption standard.


Once an encryption key is created and applied to a protected machine, there are two concepts involved in removing that encryption. The first is to disassociate the key from the protected machine. Optionally, once the encryption key is disassociated from all protected machines, it can be deleted from the Rapid Recovery Core.

This section includes the following topics:

- [Associating an encryption key with a protected machine](#)
- [Applying an encryption key from the Protected Machines page](#)
- [Disassociating an encryption key from a protected machine](#)

Associating an encryption key with a protected machine

You can apply an encryption key to a protected machine using either of two methods:

- **As part of protecting a machine.** When using this method, you can apply encryption to one or multiple machines simultaneously. This method lets you add a new encryption key, or apply an existing key to the selected machine or machines.
To use encryption when first defining protection for a machine, you must select the advanced options in the relevant Protect Machines Wizard. This selection adds an Encryption page to the wizard workflow. From this page, select **Enable encryption**, and then select an existing encryption key or specify parameters for a new key. For more information, see [Protecting a machine](#) or [About protecting multiple machines](#), respectively.
- **By modifying the configuration settings for a machine.** This method applies an encryption key to one protected machine at a time. There are two approaches for modifying configuration settings for a machine in the Rapid Recovery Core:
 - **Modify the configuration settings for a specific protected machine.** The encryption key you want to use for this approach must already exist on the Rapid Recovery Core, be a universal key type, and must be in an unlocked state. Encryption is part of the General settings. For more information, see [Viewing and modifying protected machine settings](#).
 - **Click the  Not Encrypted icon on the Protected Machines page.** Using this approach you can create and apply a new encryption key, or assign an existing unlocked universal key to the specified protected machine. For more information, see [Applying an encryption key from the Protected Machines page](#).

Applying an encryption key from the Protected Machines page

Once an encryption key has been added to aRapid Recovery Core, it can be used for any number of protected machines.



If you select an encryption key during the initial protection of one or more machines, that key is automatically applied to any machines you protect using that wizard. In such cases, this procedure is not required.

Perform this procedure:

- If you want to apply an existing, universal, unlocked encryption key to any protected machine in your Core.
- If you just added a new encryption key using the process described in the topic [Adding an encryption key](#) and want to apply that key to a protected machine.
- If encryption is already applied to a protected machine in your Core, but you want to change the key to a different universal, unlocked key available in your Core.

CAUTION: After you apply an encryption key to a protected machine, Rapid Recovery takes a new base image for that machine upon the next scheduled or forced snapshot.

1. Navigate to the Rapid Recovery Core Console and click **Protected Machines**.

The Protected Machines page appears, listing all the machines protected by this Core. An open lock  appears for any machine that does not have an encryption key applied. A closed lock  indicates that a protected machine has encryption applied.

2. In the Protected Machines pane, click the lock icon for the protected machine you want to configure. The Encryption Configuration dialog box appears.

3. Do one of the following:

- If you want to apply an existing encryption key to this machine, select **Encrypt data using Core-based encryption with an existing key**, and from the drop-down menu, select the appropriate key. Click **OK** to confirm.
- If you want to change an existing encryption key to a different universal, unlocked key, select **Encrypt data using Core-based encryption with a new key**, and from the drop-down menu, select the appropriate key. Click **OK** to confirm.
- If you want to create a new encryption key and apply it to this protected machine, select **Encrypt data using Core-based encryption with a new key**. Then enter the details for the key as described in the following table.

Table 82: New encryption key details

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters or prohibited phrases .
Description	Enter a descriptive comment for the encryption key. This information appears in the Description field when viewing a list of encryption keys in the Rapid Recovery Core Console. Descriptions may contain up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm Passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

4. Click **OK**.

The dialog box closes. The encryption key you specified has been applied to future backups for this protected machine, and the lock now appears as closed.

Optionally, if you want the encryption key applied immediately, force a snapshot. For more information, see [Forcing a snapshot](#).

! CAUTION: Rapid Recovery uses AES 256-bit encryption in the Cipher Block Chaining (CBC) mode with 256-bit keys. While using encryption is optional, Quest recommends that you establish an encryption key, and that you protect the passphrase you define. Store the passphrase in a secure location as it is critical for data recovery. Without a passphrase, data recovery is not possible.




Disassociating an encryption key from a protected machine

Once an encryption key is applied to a protected machine, all subsequent snapshot data stored in the Rapid Recovery Core is encrypted.

You can disassociate an encryption key from a protected machine. This action does not decrypt the existing backup data, but does result in a new base image for that machine at the time of the next scheduled or forced snapshot.

NOTE: If you want to remove an encryption key from the Core, as described in the topic [Removing an encryption key](#), you must first disassociate that encryption key from all protected machines.

Perform this procedure to disassociate an encryption key from a specific protected machine.

1. Navigate to the Rapid Recovery Core Console and click **Protected Machines**.
The *Protected Machines* page appears, listing all the machines protected by this Core. An open lock  appears for any machine that does not have an encryption key applied. A closed lock  indicates that a protected machine has encryption applied.
2. In the Protected Machines pane, click the  Encrypted icon for the protected machine you want to configure.
The Encryption Configuration dialog box appears.
3. Select **Encrypt data using Core-based encryption with an existing key**, and from the drop-down menu, select **(None)** and then click **OK**.
4. If you want to remove this encryption key from the Rapid Recovery Core, first repeat this procedure for all protected machines using this key. Then perform the procedure described in the topic [Removing an encryption key](#).

Managing encryption keys


To manage encryption keys for the Rapid Recovery Core, from the icon bar, click  (More) and then select **Encryption Keys**. The *Encryption Keys* page appears. For each encryption key added to your Rapid Recovery Core (if any have been defined yet), you see the information described in the following table.

Table 83: Information about each encryption key

UI Element	Description
Select Item	For each encryption key, you can select the checkbox to perform actions from the list of menu options above the table.
Name	The name associated with the encryption key.
Thumbprint	This parameter is a 26-character alphabetic string of randomly generated English upper and lower case letters that helps uniquely identify each encryption key.
Type	Type describes the origin point of an encryption key and its ability to be applied. An encryption key can contain one of two possible types: Universal. Universal type is the default condition when you create an encryption key. A key with a type of Universal, combined with a state of Unlocked, indicates that the key can be applied to a protected machine. You cannot manually lock a universal key type; instead, you must first change its type as described in the procedure Changing encryption key types . Replication. When a protected machine in a source Core has encryption enabled, and recovery points for that machine are replicated in a target Core, any encryption keys used in the source appear automatically in the target Core with a type of Replication. The default state after receiving a replicated key is locked. You can unlock an encryption key with a type of Replication by providing the passphrase. If a key has a type of Unlocked, you can manually lock it. For more

UI Element	Description
	information, see the topic Unlocking an encryption key .
State	<p>The state indicates whether an encryption key can be used. Two possible states include:</p> <ul style="list-style-type: none"> • Unlocked. An Unlocked state indicates that the key can be used immediately. For example, you can encrypt snapshots for a protected machine, or perform data recovery from a replicated recovery point on the target Core. • Locked. A Locked state indicates that the key cannot be used until it is unlocked by providing the passphrase. Locked is the default state for a newly imported or replicated encryption key. <p>If the state of an encryption key is locked, it must be unlocked before it can be used.</p> <p>If you previously unlocked a locked encryption key, and the duration to remain unlocked has expired, the state changes from unlocked to locked. After the key locks automatically, you must unlock the key again in order to use it. For more information, see the topic Unlocking an encryption key.</p>
Description	The description is an optional field that is recommended to provide useful information about the encryption key such as its intended use or a passphrase hint.

At the top level of the Encryption Keys pane, you can add an encryption key or import a key using a file exported from another Rapid Recovery Core. You can also delete keys selected in the summary table.

Once an encryption key exists for a Core, you can manage the existing keys by editing the name or description properties; changing the passphrase; unlocking a locked encryption key; or removing the key from the Rapid Recovery Core. You can also export a key to a file, which can be imported into another Rapid Recovery Core.

When you add an encryption key from the *Encryption Keys* page, the key appears in the list of encryption keys, but is not applied to a specific protected machine. For information on how to apply an encryption key you create from the Encryption Keys pane, or to delete a key entirely from the Rapid Recovery Core, see [Applying or removing encryption keys](#).

From the *Encryption Keys* pane, you can manage security for the backup data saved to the Core for any protected machine in your repository by doing the following:

- [Adding an encryption key](#)
- [Importing an encryption key](#)
- [Unlocking an encryption key](#)
- [Editing an encryption key](#)
- [Changing an encryption key passphrase](#)
- [Exporting an encryption key](#)
- [Removing an encryption key](#)
- [Changing encryption key types](#)
- [Applying an encryption key from the Protected Machines page](#)

Adding an encryption key

Rapid Recovery uses AES 256-bit encryption in the Cipher Block Chaining (CBC) mode with 256-bit keys. While using encryption is optional, Quest recommends that you establish an encryption key, and that you protect the

passphrase you define.

CAUTION: Store the passphrase in a secure location. Without a passphrase, you cannot recover data from encrypted recovery points.

After an encryption key is defined, you can use it to safeguard your data. Encryption keys can be used by any number of protected machines.

This step describes how to add an encryption key from the Rapid Recovery Core Console. This process does not apply the key to any machines currently being protected on the Core. You can also add an encryption key during the process of protecting a machine. For more information on adding encryption as part of protecting one machine, see [Protecting a machine](#). For more information on adding encryption to two or more machines while initially protecting them, see [About protecting multiple machines](#).

Complete the steps in this procedure to add an encryption key.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **⋮** (More) and then select **Encryption Keys**.
The Encryption Keys page appears.
3. Click **Add Encryption Key**.
The Create Encryption Key dialog box appears.
4. In the Create Encryption Key dialog box, enter the details for the key as described in the following table.

Table 84: Create encryption key details.

Text Box	Description
Name	Enter a name for the encryption key. Encryption key names must contain between 1 and 64 alphanumeric characters. Do not use prohibited characters and prohibited phrases .
Description	Enter a comment for the encryption key. This information appears in the Description field when viewing encryption keys from the Core Console. You can enter up to 254 characters. Best practice is to avoid using prohibited characters and prohibited phrases .
Passphrase	Enter a passphrase used to control access. Best practice is to avoid using prohibited characters . CAUTION: Record the passphrase in a secure location. Quest Data Protection Support cannot recover a passphrase. Once you create an encryption key and apply it to one or more protected machines, you cannot recover data if you lose the passphrase.
Confirm passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.




5. Click **OK**.
The dialog box closes and the encryption key you created is visible on the *Encryption Keys* page.
6. If you want to apply the encryption key to a protected machine, see [Applying an encryption key from the Protected Machines page](#).

Importing an encryption key

You can import an encryption key from another Rapid Recovery Core and use that key to encrypt data for a protected machine in your Core. To import the key, you must be able to access it from the Core machine, either locally or through your network. You must also know the passphrase for the encryption key.

Complete the steps in this procedure to import an encryption key.

NOTE: This procedure does not apply the key to any protected machines. For more information on applying the key, see [Applying an encryption key from the Protected Machines page](#).

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More) and then select  **Encryption Keys**.
The Encryption Keys page appears.
3. Click  **Import**.
The File Upload dialog box appears.
4. In the File Upload dialog box, navigate to the network or local directory containing the encryption key you want to import.
For example, navigate to the *Downloads* folder for the logged-in user.
The key filename starts with "EncryptionKey-", followed by the key ID, and ending in the file extension *.key*.
For example, a sample encryption key name is EncryptionKey-RandomAlphabeticCharacters.key.
5. Select the key you want to import, and then click **Open**.
6. In the Import Key dialog box, click **OK**.
The dialog box closes and the encryption key you imported is visible on the Encryption Keys page. If the encryption key was used to protect a volume before it was exported, the state of the key is Locked.

Unlocking an encryption key

Encryption keys may contain a state of unlocked or locked. An unlocked encryption key can be applied to a protected machine to secure the backup data saved for that machine in the repository. From a Rapid Recovery Core using an unlocked encryption key, you can also recover data from a recovery point.

When you import an encryption key into a Rapid Recovery Core, its default state is Locked. This is true regardless of whether you explicitly imported the key, or whether the encryption key was added to the Rapid Recovery Core either by replicating encrypted protected machines or by importing an archive of encrypted recovery points.

For encryption keys added to the Rapid Recovery Core by replication only, when you unlock a key, you can specify a duration of time (in hours, days, or months) for the encryption key to remain unlocked. Each day is based on a 24-hour period, starting from the time the unlock request is saved to the Rapid Recovery Core. For example, if the key is unlocked at 11:24 AM on Tuesday and the duration selected is 2 days, the key automatically re-locks at 11:24 AM that Thursday.

NOTE: You cannot use a locked encryption key to recover data or to apply to a protected machine. You must first provide the passphrase, thus unlocking the key.

You can also lock an unlocked encryption key, ensuring that it cannot be applied to any protected machine until it is unlocked. To lock an encryption key with a state of Universal, you must first change its type to Replicated.

If an unlocked encryption key is currently being used to protect a machine in the Core, you must first disassociate that encryption key from the protected machine before you can lock it.

Complete the steps in this procedure to unlock a locked encryption key.


1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **☰** (More) and then select **Encryption Keys**.
The Encryption Keys page appears. The State column indicates which encryption keys are locked.
3. Locate the encryption key you want to unlock, click its drop-down menu **☰**, and select **Unlock**.
The Unlock Encryption Key dialog box appears.
4. In the dialog box, in the **Passphrase** text box, enter the passphrase to unlock this key.
5. To specify the length of time that the key remains unlocked, in the Duration option, do one of the following:
 - To specify that the key remains unlocked until you explicitly lock it, select **Until locked manually**.
 - To specify that the key remains locked for a duration which you configure in hours, days, or months, do the following:
 - Select the number field and enter a value between 1 and 999.
 - Specify the duration number as hours, days, or months, respectively.
 - Then click **OK**.
This option is available for encryption keys added by replication.
The dialog box closes and the changes for the selected encryption key are visible on the *Encryption Keys* page.
6. To specify that the key remains locked until a date and time that you specify, do the following:
 - Select the **Until** option.
 - In the text field or using the calendar and clock widgets, explicitly specify the data and time you want the encryption key to lock.
 - Then click **OK**.
This option is available for encryption keys added by replication.
The dialog box closes and the changes for the selected encryption key are visible on the *Encryption Keys* page.


Locking an encryption key

When an encryption key state is locked, it cannot be applied to any protected machine until it is unlocked. To lock an encryption key with a type of Universal, you must first change its type to Replicated.

Complete the steps in this procedure to lock an encryption key.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **☰** (More) and then select **Encryption Keys**.
The Encryption Keys page appears. The State column indicates which encryption keys are unlocked, and shows the type for each key.
3. Locate the encryption key you want to lock. If its type is Universal, then click its drop-down menu **☰**, and select **Change the type to Replicated**.
The Change Encryption Key Type dialog box appears.
4. In the dialog box, confirm that you want to change the key type to **Replicated**.


5. If you successfully changed the encryption key status to Replicated, then click its drop-down menu , and select **Lock**.
The *Lock Encryption Key* dialog box appears.
6. In the dialog box, confirm that you want to lock the key.
The dialog box closes, and the state of the selected encryption key is now locked.



 **NOTE:** This option is available for encryption keys added by replication.

Editing an encryption key

After an encryption key is defined, you can edit the name of the encryption key or the description of the key. These properties are visible when you view the list of encryption keys in the Encryption Keys pane.


Complete the steps in this procedure to edit the name or description of an existing unlocked encryption key.


 **CAUTION:** After you edit the name or description an encryption key that is used to protect one or more machines, Rapid Recovery takes a new base image. That base image snapshot occurs for that machine upon the next scheduled or forced snapshot.


1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More) and then select **Encryption Keys**.
The Encryption Keys page appears.
3. Locate the encryption key you want to edit, and do the following:
 - If the key is locked, you must first unlock it. See [Unlocking an encryption key](#).
 - If the key is unlocked, proceed as described below.
4. Click the drop-down menu  for the specified encryption key, and select **Edit**.
The Edit Encryption Key dialog box appears.
5. In the dialog box, edit the name or the description for the encryption key, and then click **OK**.
The dialog box closes, and the changes for the selected encryption key are visible on the Encryption Keys page.


Changing an encryption key passphrase

To maintain maximum security, you can change the passphrase for any existing encryption key. Complete the steps in this procedure to change the passphrase for an encryption key.

 **CAUTION:** After you edit the name or description an encryption key that is used to protect one or more machines, Rapid Recovery takes a new base image. That base image snapshot occurs for that machine upon the next scheduled or forced snapshot.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More) and then select **Encryption Keys**.
The Encryption Keys page appears.



3. Locate the encryption key you want to update, click its drop-down menu , and select **Change passphrase**.
The Change Passphrase dialog box appears.
4. In the dialog box, in the **Passphrase** text box, enter the new passphrase for the encryption.
5. In the **Confirm passphrase** text box, re-enter the identical passphrase.
6. Click **OK**.
The dialog box closes and the passphrase is updated.
7. Optionally, if you use a hint in the Description field, edit the encryption key to update the hint. For more information, see [Editing an encryption key](#).

 **CAUTION:** Rapid Recovery uses AES 256-bit encryption in the Cipher Block Chaining (CBC) mode with 256-bit keys. Quest recommends that you record the passphrase in a secure location and keep this information updated. Quest Data Protection Support cannot recover a passphrase. Without the passphrase, you cannot recover information from encrypted recovery points.

Exporting an encryption key


You can export an encryption key from any Rapid Recovery Core with the express purpose of using it in another Core. When you perform this procedure, the key is saved to the *Downloads* folder for the active Windows user account.

Complete the steps in this procedure to export an encryption key.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (More) and then select **Encryption Keys**.
The Encryption Keys page appears.
3. Locate the encryption key you want to export, click its drop-down menu , and select **Export**.
The Opening EncryptionKey-[name.key] dialog box appears.
4. In the dialog box, select **Save File** to save and store the encryption keys in a secure location, and then click **OK**.
The encryption key downloads as a text file to the default location, such as the *Downloads* folder of the active Windows user account.
5. Optionally, if you want to import this key into a different Core, copy the file to a location accessible from that Core.

Removing an encryption key

When you remove an encryption key from the Encryption Keys page, the key is deleted from the Rapid Recovery Core.

 **NOTE:** Removing an encryption key does not decrypt the recovery points already saved using the key. You must still retain and provide the passphrase for the key to recover data for existing encrypted recovery points.

You cannot remove an encryption key that is already associated with any protected machine. You must first view the encryption settings for each protected machine using the key, and disassociate the encryption key you want to remove. For more information, see the topic [Disassociating an encryption key from a protected machine](#).

Complete the steps in this procedure to remove an encryption key.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **⋮** (More) and then select **Encryption Keys**.
The Encryption Keys page appears.
3. Locate the encryption key you want to remove. Click its drop-down menu **⋮**, and select **Remove**.
The Remove Encryption Key dialog box appears. You see a message confirming the action to remove the encryption key.
4. In the dialog box, confirm that you want to remove the encryption key.

i **NOTE:** Removing an encryption key does not decrypt the recovery points already saved using the key. You must still retain and provide the key to recover data for existing encrypted recovery points. The dialog box closes, and the encryption key you removed no longer appears on the Encryption Keys page.

Changing encryption key types

Encryption keys list one of two possible types on the Encryption Keys pane: Universal or Replication. The type indicates the likely origin of the encryption key, and determines whether you can change its details or passphrase. You can modify these attributes only if the type is Universal. If you need to modify these attributes for a key with Replicated type, you must change its type to Universal using this procedure. When you change the type of an encryption key to Universal, it is unlocked manually and can be used to encrypt other protected machines.

i **NOTE:** You must know the passphrase to change the type from Replicated to Universal.


Encryption keys also have two possible states: Locked or Unlocked. The state controls your ability to apply an encryption key to a protected machine, or to restore data from a recovery point with encryption. You can change the type of an encryption key manually only if the state is Unlocked.

When you first create an encryption key, its type is Universal, and its state is Unlocked. You can use such a key immediately (for example, to encrypt backups for a protected machine). However, a Universal key type cannot be locked manually. If you want to manually lock an encryption key with a type of Universal, you must change the type to Replicated using this procedure.

You cannot change an encryption key type if it is already in use encrypting recovery points for one or more protected machine.

Follow this procedure to change an encryption key type.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **⋮** (More) and then select **Encryption Keys**.
The Encryption Keys page appears. Any encryption keys accessible to the Core appear in a summary table. Each lists a type of Universal or Replicated.
3. Locate the encryption key you want to update.
4. If you want to change a Universal encryption key to Replication, do the following:
 - a. Click its drop-down menu **⋮**, and select **Change the type to Replicated**.
The Change Encryption Key Type dialog box appears. You see a message confirming that you want to change the type to Replicated.
 - b. In the dialog box, confirm that you want to change the type to Replication.
The dialog box closes, and the encryption key type updates to Replication.

5. If you want to change a Replication encryption key to Universal, do the following:
 - a. Click its drop-down menu , and select **Change the type to Universal**
The Change Encryption Key Type dialog box appears. You see a message confirming that you want to change the type to Universal.
 - b. In the dialog box, in the **Passphrase** text box, enter the passphrase and then click **OK** to confirm that you want to change the type to Universal.
The dialog box closes, and the encryption key type updates to Universal.

Authentication

This section describes the two options available for logging in to the Rapid Recovery Core Console.

Topics include:

- [Understanding SAML single sign-on](#)
- [Credentials Vault](#)

Understanding SAML single sign-on

Rapid Recovery offers the ability to link a single sign-on identity provider (IdP), to the Rapid Recovery Core using the SAML 2.0 protocol. Use of this feature is optional. Rapid Recovery supports only service provider-initiated login.

To configure Rapid Recovery for SAML, see [Configuring SAML settings](#).

Prerequisite

Before you configure the SAML feature, establish an account with one of the following compatible IdPs:

- Microsoft Azure Active Directory (Azure AD)
- Okta
- OneLogin

For more information, see the white paper *Configuring SAML single sign-on authentication for Quest Rapid Recovery* and the documentation provided by your IdP.

Credentials Vault

This section describes the Credentials Vault feature of Rapid Recovery.

Topics include:

- [Understanding the Credentials Vault](#)
- [Adding accounts to the Credentials Vault](#)
- [Viewing or changing accounts saved in the vault](#)
- [Using credentials from the vault](#)

Understanding the Credentials Vault

Credentials Vault is a usability feature of Rapid Recovery release 6.8 and later that manages account login credentials used within the Rapid Recovery Core Console. Use of this feature is optional.

When performing operations such as adding a machine or cluster to protection, setting up virtual export or replication, connecting to a repository, archiving or restoring archived recovery points, and so on, you are prompted to enter account credentials. For each user account, credentials include the user name, password, and a description field to identify the account. After you enter your credentials, if you choose to, you can add them to the Credentials Vault.

Thereafter, the next time you want to perform an operation in the Core Console that uses the same account, instead of manually entering your user name and password, you can select the account from a drop-down menu.

The Credentials Vault simplifies management of your passwords. For example, if your organization has a security policy mandating password changes at frequent intervals, one visit to the Credentials Vault page can let you easily update your password for each user account accessed from the Rapid Recovery Core Console.

The Credentials Vault is unobtrusive. Sections of the Core Console UI that are enabled for the Credentials Vault include a + sign next to the User name field when prompted for credentials.

As its name implies, the Credentials Vault includes security features. For example:

- Credential information in the vault is encrypted.
- Once entered and saved, passwords are not displayed. This reduces the chance of exposure of individuals' passwords when multiple users access the Core.
- By design, Cloud credentials are managed separately in the Core Console.

At any time, you can open the Credentials Vault page in the Core Console to view and manage accounts saved in the Credentials Vault. If no accounts have been entered yet, optionally, you can add them directly from this page.

- For information on adding accounts to the vault, see [Adding accounts to the Credentials Vault](#).
- For more information on viewing and modifying account credentials held in the vault, see [Viewing or changing accounts](#).
- For information on using account credentials saved in the vault, see [Using credentials from the vault](#).

Command Line and PowerShell scripts exist to support this feature. For more information about the Credentials Vault, see the most recent edition of the *Rapid Recovery Commands and Scripting Reference Guide*.

Adding accounts to the Credentials Vault

You can add accounts to the Credentials Vault from the Credentials Vault page, or from practically any Rapid Recovery Core Console window or wizard in which account credentials are requested.

Follow this procedure to add accounts to the Credentials Vault.

1. Do one of the following:
 - a. If you are on the Credentials Vault page of the Core Console, click **+ Add New Account**.
 - b. If you are viewing a Credentials Vault-enabled wizard page, window, or dialog box in the Rapid Recovery Core Console, next to the User name field, click **+**.

The *Add New Account* dialog box opens.

If you already entered your account user name and password, those fields are populated. By default, the Description field populates with the current system date and time.

2. In the **User name** field, if required, enter the user name for this account.
3. In the **Password** field, if required, enter the password for this account.
4. In the **Description** field, enter a meaningful text description of this account. Do not skip this step.




CAUTION: Replace the default text with a unique text string that clearly describes the account it represents.


Quest strongly recommends adding well-planned descriptions for accounts held in the vault. Consider the following points:




- Some users will have two or more accounts saved to the Credentials Vault with the same user name. Particularly in these cases, it is the description field that lets you identify the correct account in the vault.
 - For security purposes, passwords saved to the vault are never displayed.
 - Since passwords are not displayed, you cannot rely on the combination of user name and password to later identify the purpose of the account.
 - If you add the same credential to the vault, you can later merge them.
5. When satisfied with your selections, click **OK**.
The *Add New Account* dialog box closes, and your account credentials information is saved securely to the vault.

Viewing or changing accounts saved in the vault

Follow this procedure to view accounts in the Credentials Vault, or to edit, merge, or remove accounts.

1. From the Rapid Recovery Core Console, click the  (More) menu, and then select  **Credentials Vault**.
The *Credentials Vault* page appears. For each account, the user name, description, and utilization appears.
2. In the **User name** field, if you want to see when the account information was last modified, click .
3. In the **Description** field, read the descriptive text to help identify which account this record represents.
4. In the **Utilization** field, see how many systems reference this account from the vault.

5. To edit any record, do the following:
 - a. From the row for the appropriate account, click  (More) and then select **Edit**.
The *Edit Account* dialog box appears.
 - b. If you want to change the user name associated with this account, in the **User name** field, update the information.
 - c. If you want to change the password currently saved for this account, in the **Password** field, enter the appropriate password.


NOTE: Since passwords are not viewable, if you have any concerns about the password associated with this account, simply re-enter the appropriate password in the **Password** field.
 - d. If you want to update or change the description, enter the updated information in the **Description** field.
 - e. When satisfied, click **OK**.
6. To merge account records (for example, if you created more than one entry for the same account), do the following:
 - a. From the row for the appropriate account, select  **Merge to account**.
The *Merge to Account* dialog box appears.
 - b. From the **Target account** field, select the other account with which you want to merge this account record.
 - c. Click **Merge**.
The *Merge to Account* dialog box closes, the screen refreshes, and the account records are merged.
7. To remove an account from the Credentials Vault, from the row for the appropriate account, click  (More) and then select  **Remove**.

Using credentials from the vault

Many actions in the Rapid Recovery Core Console require you to enter account credentials.

After accounts have been added to the Credentials Vault, when prompted to authenticate, you can view the list of accounts and select an account with one click, rather than manually entering your account user name and password.

Follow this procedure to use an account from the Credentials Vault.

1. From a location on the Rapid Recovery Core Console in which you are asked for credentials, click the downward-facing arrow  in the **User name** field to expand the view.

The Credentials Vault drop-down grid appears. Each row shows the user name and description associated with an account held in the vault.
2. If necessary, scroll through the list to identify the account for which you want to enter credentials. Then click on the row for the appropriate account.
The grid closes, and the account information is passed to the window or dialog box. Since passwords are hidden, the password field is not shown.
3. Complete the function requiring credentials.

Replication

This section describes how to configure and manage the replication of protected data from a Rapid Recovery source Core to a Rapid Recovery target Core for disaster recovery.

Topics include:

- Replication with Rapid Recovery
- Recovery point chains and orphans
- When replication begins
- Determining your seeding needs and strategy
- Performance considerations for replicated data transfer
- Viewing incoming and outgoing replication
- Configuring replication
- Replicating to a self-managed target Core
- Replicating to a third-party target Core
- Adding a machine to existing replication
- Consuming the seed drive on a target Core
- Managing replication settings
- Removing outgoing replication from the source Core
- Removing incoming replication from the target Core
- Recovering replicated data

Replication with Rapid Recovery

This section provides conceptual and procedural information to help you understand and configure replication in Rapid Recovery.

Replication is the process of copying recovery points from one Rapid Recovery Core and transmitting them to another Rapid Recovery Core for disaster recovery purposes. The process requires a paired source-target relationship between two or more Cores.

The source Core copies the recovery points of selected protected machines, and then asynchronously and continually transmits that snapshot data to the target Core.

Unless you change the default behavior by setting a replication schedule, the Core starts a replication job immediately after completion of every backup snapshot, checksum check, mountability check, and attachability check. Log truncation of any type also triggers a replication job, as does checking the integrity of recovery points or of an Oracle database. If any of these actions are included in nightly jobs, then completion of nightly jobs also triggers a replication job. For more information, see [Scheduling replication](#).

i **NOTE:** When you replicate data for a cluster, you must replicate the entire cluster. For example, if you select a node to replicate, the cluster is automatically selected. Likewise, if you select the cluster, all nodes in that cluster are also selected.

For optimum data security, administrators usually use a target Core at a remote disaster recovery site. You can configure outbound replication to a company-owned data center or remote disaster recovery site (that is, a “self-managed” target Core). Or, you can configure outbound replication to a third-party managed service provider (MSP) or cloud provider that hosts off-site backup and disaster recovery services. When replicating to a third-party target Core, you can use built-in work flows that let you request connections and receive automatic feedback notifications. Replication is managed on a per-protected-machine basis. Any machine (or all machines) protected or replicated on a source Core can be configured to replicate to a target Core.

Possible scenarios for replication include:

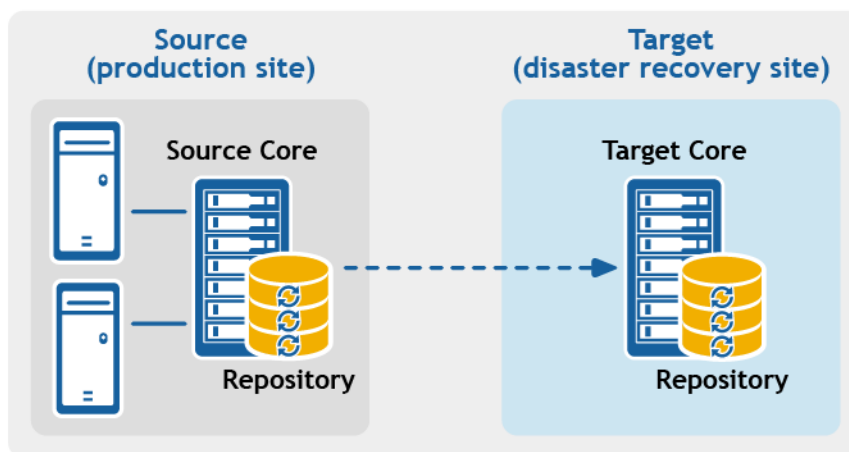
- **Replication to a local location.** The target Core is located in a local data center or on-site location, and replication is maintained at all times. In this configuration, the loss of the Core would not prevent a recovery.
- **Replication to an off-site location.** The target Core is located at an off-site disaster recovery facility for recovery in the event of a loss.
- **Replication to Microsoft Azure.** The target Core is located in an Azure-hosted virtual machine and the repository is located in an Azure Blob Container. For more information, see [Azure repositories](#).
- **Mutual replication.** Two data centers in two different locations each contain a Core and are protecting machines and serving as the off-site disaster recovery backup for each other. In this scenario, each Core replicates the protected machines to the Core that is located in the other data center. This scenario is also called cross replication.
- **Hosted and cloud replication.** Rapid Recovery MSP partners maintain multiple target Cores in a data center or a public cloud. On each of these Cores, the MSP partner lets one or more of their customers replicate recovery points from a source Core on the customer’s site to the MSP’s target Core for a fee.

NOTE: In this scenario, customers only have access to their own data.

Possible replication configurations include:

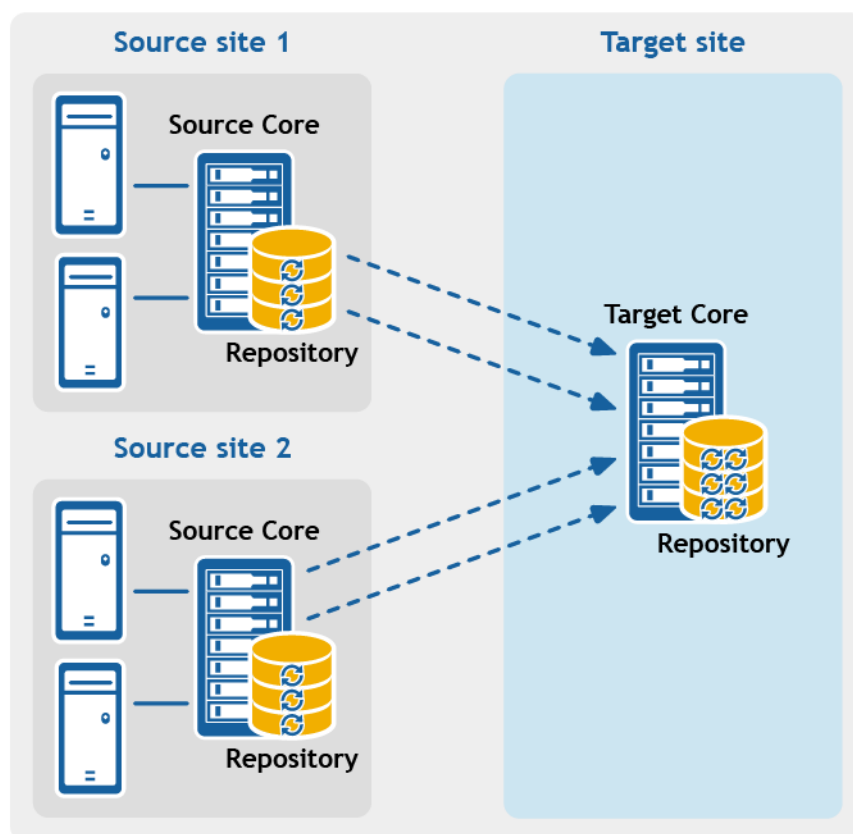
- **Point-to-point replication.** Replicates one or more protected machines from a single source Core to a single target Core.

Figure 1: Point-to-point replication configuration



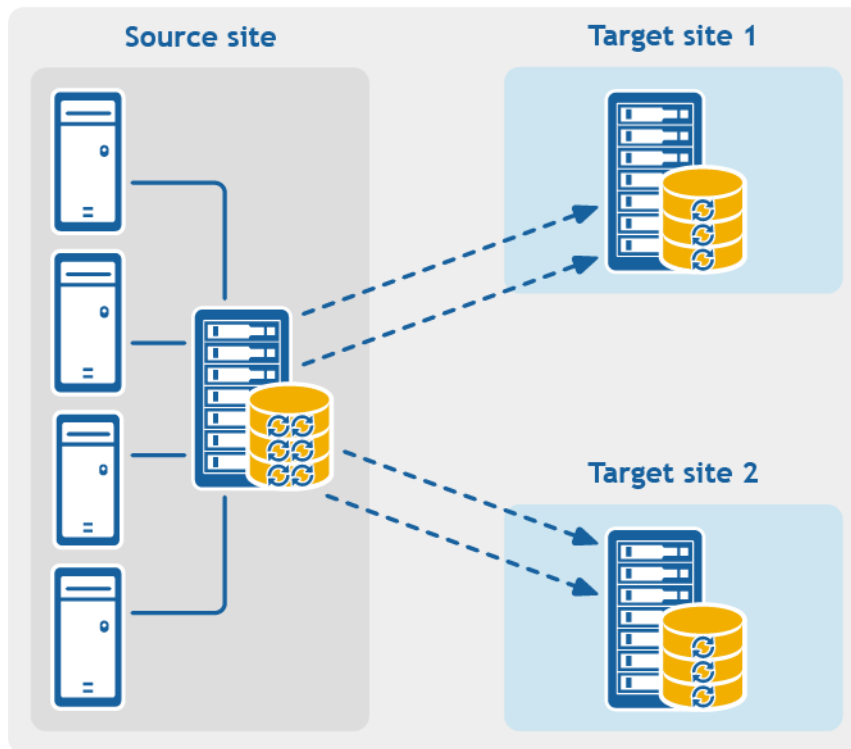
- **Multipoint-to-point replication.** Replicates protected machines from multiple source Cores to a single target Core.

Figure 2: Multipoint-to-point replication configuration



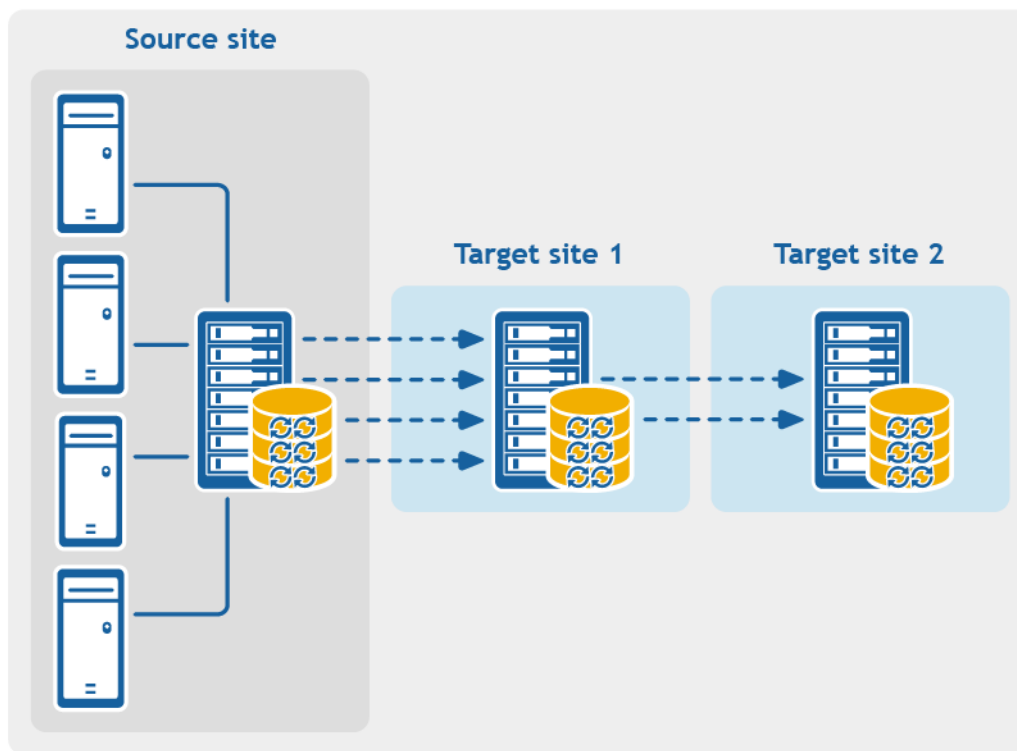
- **Point-to-multipoint replication.** Replicates one or more protected machines from a single source Core to more than one target Core.

Figure 3: Point-to-multipoint replication configuration



- **Multi-hop replication.** Replicates one or more protected machines from one target Core to another target Core, producing additional failover or recovery options on the replicated Core.

Figure 4: Multi-hop replication configuration



Recovery point chains and orphans

Rapid Recovery captures snapshots of a protected machine, and saves the data to a repository as a *recovery point*. The first recovery point saved to the Core is called a *base image*. The base image includes the operating system, applications, and settings for each volume you choose to protect, as well as all data on those volumes. Successive backups are *incremental snapshots*, which consist only of data changed on the protected volumes since the last backup. The base image plus all incremental snapshots together form a complete *recovery point chain*.

From a complete recovery point chain, you can restore data with ease and confidence, using the full range of recovery options available to Rapid Recovery. These options include file-level restore, volume-level restore, and bare metal restore.

Since logically you cannot restore from data that does not exist, in the case of an incomplete recovery point chain, you cannot restore data at the volume level or perform a bare metal restore. In such cases, you can still restore any data that does exist in a recovery point at the file level.

If the information you want to restore from a recovery point is in a previous backup that is not available to the Core (an earlier incremental snapshot or the base image), the recovery point is said to be *orphaned*. Orphaned recovery points are typical in some replication scenarios.

For example, when you first establish replication, your options for restoring data from the replicated recovery points are limited. Until all backup data from the source Core is transmitted to the target Core, creating full recovery point chains from the orphans, you can only perform file-level restore.

When replication begins

By default, replication transfer jobs are automatically queued by the Core immediately after each regularly scheduled backup transfer completes. Thus, unless the replication schedule for a protected machine is customized, its replication schedule is based on its standard backup snapshot schedule.

When you first set up replication, if one or more recovery points exist on the source Core, the replication process begins immediately, unless:

- You select the option to initially pause replication, or
- You select the option to use a seed drive to perform the initial transfer.

If you pause replication initially, replication begins when you explicitly resume replication.

If you set up replication and specify the use of a seed drive, replication to the target Core begins with the next regularly scheduled backup snapshot.

i **NOTE:** You can also force a backup of the protected machine after establishing replication. This causes replication to begin immediately after the protected machine snapshot completes.

If you specify a seed drive when you set up replication, only future backup transfers are replicated. If you want existing recovery points from the original protected machine to exist on the target Core, you must seed data from the protected machine. To seed data, create a seed drive from the source Core, and then consume the seed drive on the target Core.

You can also customize the replication schedule for a protected machine. For example, if you use the default protection schedule of one backup per hour, you can specify that the source Core replicate to the target Core at a different schedule (for example, once daily at 2AM).

Determining your seeding needs and strategy

The following topics discuss restoring from replicated data and whether you need to seed recovery point data from the source Core.

When seeding data is required

When you first establish replication, unless you specify the use a seed drive, the source Core begins transmitting all of the recovery points for the selected machines to the target Core. Transmitting your data over the network can take a good deal of time. Factors involved include the speed of your network, the robustness of your network architecture, and the amount of data to be transmitted to the target Core. For example, if the backup data on the source Core measures 10GB and the WAN link transfers 24Mbps, the transfer could take approximately one hour to complete.

Based on the amount of information you want to copy to the target Core, the seed drive can add up to hundreds or thousands of gigabytes of data. Many organizations choose not to consume the network bandwidth required, and instead opt to define and consume a seed drive. For more information, see [Performance considerations for replicated data transfer](#).

If you specify the use of a seed drive when defining replication, then only recovery points saved to the source Core *after* you establish replication are replicated to the target Core. Backups saved on the source Core *before*

replication was established will not be present on the target Core until you explicitly *seed* the data, using the following process.

To avoid slowing down your network with an intensive transfer of historical data, seed your prior backup data to the target Core using a **seed drive**. A seed drive is an archive file that **copies** a set of deduplicated base images and incremental snapshots from the source Core. The seed drive file contains the full set of previous recovery points for the protected machines you want to replicate from the source Core to the target Core.

Move the seed drive file to a storage volume which you then make available to the target Core. Then you **consume** the information from the seed drive. This involves attaching the volume with the seed drive image to the target Core and importing the data to the repository from the Core Console. This process repairs orphans, uniting incremental snapshots replicated to the target Core with their base images, to form one or more complete recovery point chains. This process is sometimes called copy-consume.

Seeding data from your source Core is not always required. For example:

- If you are setting up replication for a new Rapid Recovery Core, seeding is not required.
- If the data from previous snapshots are not critical for your replicated data, and you only need to recover data saved after replication is set up, seeding is not required.

i | **NOTE:** In this case, Quest recommends capturing a new base image immediately before or immediately after setting up replication. This step ensures a full recovery point chain exists on the target Core from which to restore data in the future.

- If you captured a base image immediately before setting up replication, and only have a need to restore from data captured after that date, seeding is not required.
- If you set up replication without specifying a seed drive, then the snapshot data transmits over the network from the source Core to the target Core.

If one of these situations applies to you, you do not need to seed data. In such cases, replication can be completed entirely from the source Core.

If you set up replication for a Core with existing recovery points and may need to restore at the volume level, want to perform a BMR, or want to restore data from an earlier base image or incremental snapshot, seeding is required. In these situations, consider your seeding needs and strategy. Review the information in this topic and decide whether you will seed to your target Core, and which approach you will use.

Approaches to seeding data

If you want your replicated machines on a target Core to have access to data saved previously on the original source Core, seed your target Core using one of the following approaches:

1. **Seed to the target Core over a network connection.** Specify the use of a seed drive when you define replication. You can then share the folder containing the seed drive with the target Core, and consume the seed drive file over the network. For large data or slow connections, seeding by this method can take a substantial amount of time and consume substantial network bandwidth.

i | **NOTE:** Quest does not recommend seeding large amounts of data over a network connection. Initial seeding potentially involves very large amounts of data, which could overwhelm a typical WAN connection.

2. **Transfer backup data from the source Core using physical storage media.** Transfer the seed drive file to a portable external removable storage device. This approach is typically useful for large sets of data or sites with slow network connections. Seeding using this method requires you to perform the following steps:
 - a. Create a seed archive from the source Core, saving it to removable media.
 - b. Transport the seed drive to the physical location of the target Core.
 - c. Attach the drive to the target Core.
 - d. Consume the data from the seed drive to the repository of the target Core.

If replicating to a third-party Core, once your media is received by the MSP, a data center representative typically attaches the media and notifies you when it is ready for you to consume (or import) the seed data into the Core.

i NOTE: Because large amounts of data need to be copied to the storage device, an eSATA, USB 3.0, or other high-speed connection is recommended. If the total size of the seed data archive is larger than the space available on the removable media, the archive can span across multiple devices.

3. **For source and target Cores stored on virtual hosts, transfer backup data using a virtual hard disk.** If your source Core and target Core are both on a virtual host, you can define and consume a seed drive on virtual storage media. Seeding using this method requires you to perform the following steps:
 - a. Create a seed drive file from the source Core, saving it to a virtual storage volume.
 - b. Detach the volume from the source Core and attach it to the target Core.
 - c. Consume the data from the seed drive to the repository of the target Core.

i NOTE: While replication of incremental snapshots can occur between the source and target Cores before seeding is complete, the replicated snapshots transmitted from the source to the target remain orphaned until the initial data is consumed, and they are combined with the replicated base images.

Related links

- For details on the process of consuming the seed drive, see the topic [Consuming the seed drive on a target Core](#).
- For more information about orphaned recovery points, see [Deleting an orphaned recovery point chain](#).
- For information on preparing a seed drive, see [Consuming the seed drive on a target Core](#).
- For more information on seeding a physical drive to the Azure data center, see the Azure document center. See article <https://azure.microsoft.com/en-us/documentation/articles/storage-import-export-service/>.

Performance considerations for replicated data transfer

If the bandwidth between the source and target Cores cannot accommodate the transfer of stored recovery points, set up replication and specify the use of a seed drive. This process seeds the target Core with base images and recovery points from the selected servers protected on the source Core. The seeding process can be performed at any time. Seeding can be performed as part of the initial transfer of data, which serves as the foundation for regularly scheduled replication. You can also seed data for a previously replicated machine if replication has been paused or deleted. In this case, the "Build recovery point chains" option lets you copy not-yet replicated recovery

points to a seed drive.

When preparing for replication, consider the following factors:

- **Change rate.** The change rate is the rate at which the amount of protected data is accumulated. The rate depends on the amount of data that changes on protected volumes and the protection interval of the volumes. Some machine types typically have a higher change rate, such as an Exchange email server. One way to reduce the change rate is to reduce the protection interval.
- **Bandwidth.** The bandwidth is the available transfer speed between the source Core and the target Core. It is crucial that the bandwidth be greater than the change rate for replication to keep up with the recovery points snapshots create. For very large data transfers from Core to Core, multiple parallel streams may be required to perform at wire speeds up to the speed of a 1GB Ethernet connection.

i **NOTE:** Bandwidth that ISPs specify is typically the total available bandwidth. All devices on the network share the outgoing bandwidth. Make sure that there is enough free bandwidth for replication to accommodate the change rate.

- **Number of protected machines.** It is important to consider the number of machines protected per source Core and how many you plan to replicate to the target. You are not required to replicate every machine protected on the source Core; Rapid Recovery lets you replicate on a per-protected machine basis, so you can choose to replicate only certain machines, if you want. If all protected machines on a source Core must be replicated, the change rate is typically higher. This factor is relevant if the bandwidth between the source and target Cores is insufficient for the amount and size of the recovery points being replicated.

The maximum change rate for WAN connection types is shown in the following table, with examples of the necessary bandwidth per gigabyte for a reasonable change rate.

Table 85: Examples of bandwidth per gigabyte

Broadband	Bandwidth	Max Change Rate
DSL	768 Kbps and up	330MB per hour
Cable	1 Mbps and up	429MB per hour
T1	1.5 Mbps and up	644MB per hour
Fiber	20 Mbps and up	8.38GB per hour

i **NOTE:** For optimum results, adhere to the recommendations listed in the preceding table.

If a link fails during data transfer, replication resumes from the previous failure point of the transfer (once link functionality is restored).

Depending on your network configuration, replication can be a time-consuming process. Ensure that you account for enough bandwidth to accommodate replication, other Rapid Recovery transfers such as backups, and any other critical applications you must run.

If you experience issues successfully transferring data over the network, especially for certain protected or replicated machines, considering adjusting the rate of data transfer for those machines. For more information, see [About modifying transfer settings](#) and [Throttling transfer speed](#).

About replication and encrypted recovery points


While the seed drive does not contain backups of the source Core registry and certificates, the seed drive does contain encryption keys from the source Core if the recovery points being replicated from source to target are encrypted. The replicated recovery points remain encrypted after they are transmitted to the target Core. The owners or administrators of the target Core need the passphrase to recover the encrypted data.

About retention policies for replication

Retention policies on the source and target Cores are not synchronized. Rollup and on-demand deletion perform independently on each Core on initial action, as well as when running nightly jobs.

For more information on retention policies, see [Managing retention policies](#).

Viewing incoming and outgoing replication

If you click the  Replicate icon from the icon bar, the Replication page appears. This page gives you an understanding of replication from the scope of this Core. It includes two panes:


- The Outgoing Replication pane lists any machines protected in this Core that are replicated on another Core.
- The Incoming Replication pane lists the machines replicated on this Core, and the source Core from which these machines are replicated.

This section describes the information shown in these panes.

Information about outgoing replication from this Rapid Recovery Core is described in the following table.


Table 86: Information about outgoing replication

UI Element	Description
Select item	For each row in the summary table, you can select the checkbox to perform actions from the list of menu options above the table.
Type	Shows the machine type. You can expand a target Core to show individual replicated machines.
Status indicator	Status of replication. Colored circles in the Status column show whether a replicated machine is online or unreachable. If you hover the cursor over the colored circle, the status condition is displayed. Status conditions include green (replication established and online), yellow (replication paused), red (authentication error), and gray (offline or unreachable).
Replication Name	The display name of the Core machine to which machines from this source Core are replicated.
Machines	Lists the number of machines replicated to the selected target Core.
Sync	The date and time of the last replication transfer to the target Core.

UI Element	Description
	When you click the More drop-down menu in this column, you see a list of actions to perform to affect the specific replication relationship.

You can perform actions on two or more of the target Cores listed in the Outgoing Replication grid. To perform actions on multiple target Cores, select the checkbox for each Core in the grid, and then, from the menu above the grid, select the action you want to perform. You can perform the actions described in the following table.


Table 87: Global actions available in the Outgoing Replication pane

UI Element	Description
Add Target Core	Lets you define another target Core to replicate machines protected on this source Core.
Refresh	Refreshes the information shown in the table.
Force	Forces replication.
Pause	Pauses established replication.
Resume	Resumes paused replication.
Copy	Opens the replication wizard, letting you copy existing recovery points for selected protected machines to a seed drive.
Delete	Deletes outgoing replication.
Seed Drives	This menu option appears if data was copied to a seed drive when replication was set up. Displays information about the seed drive file, including the data and time the seed drive was saved. Collapsible menus indicate the target Core and the protected machines from which the seed drive files were generated.
	When you click the More drop-down menu in this column, you see a list of actions to perform to affect the specific replication relationship.

Information about incoming replication from another Core is described in the following table.

Table 88: Information about incoming replication

UI Element	Description
Select item	For each row in the summary table, you can select the checkbox to perform actions from the list of menu options above the table.
Type	Shows the machine type. You can expand a source Core to show individual replicated machines.
Status indicator	Status of replication. Colored circles in the Status column show whether a replicated machine is online or unreachable. If you hover the cursor over the colored circle, the status condition is displayed. Status conditions include green (replication established and online), yellow (replication paused), red (authentication error), and gray (offline or unreachable).
Replication Name	The display name of the source Core machine containing protected machines that are replicated on this target Core. This name can be optionally specified when establishing replication on the source Core using the Replication Wizard.

UI Element	Description
Machines	Lists the number of machines protected on the source Core that are replicated to this target Core.
Sync	The date and time of the last replication transfer from the source Core.
	When you click the More drop-down menu in this column, you see a list of actions to perform to affect the specific replication relationship.

You can perform actions on two or more of the source Cores listed in the Incoming Replication grid. To perform actions on multiple source Cores, select the checkbox for each Core in the grid, and then, from the menu above the grid, select the action you want to perform. You can perform the actions described in the following table.

Table 89: Global actions available in the Incoming Replication pane

UI Element	Description
Refresh	Refreshes the information shown in the table.
Force	Forces replication.
Pause	Pauses established replication.
Resume	Resumes paused replication.
Delete	Deletes incoming replication.

Configuring replication

To replicate data using Rapid Recovery, you must configure the source and target Cores for replication. After you configure replication, you can then replicate protected machine data, monitor and manage replication, and perform recovery.

The version of Rapid Recovery Core on the server used as the target can be equal to or later than the version of Rapid Recovery Core installed on the source Core. The source Core must never run a version of Rapid Recovery more recent than the target Core runs.

When upgrading Cores that use replication, do the following:

- Always upgrade the target Core first, then upgrade the source Core, and lastly upgrade the Agent software on your protected machines.
- If using the automatic update feature, to ensure the proper upgrade order is followed, set up automatic update on target Cores only. After each automatic update, manually upgrade source Cores before updating Rapid Recovery Agent on your protected machines.

For more information on updating cores configured for replication, see the *Rapid Recovery Installation and Upgrade Guide*, including topics "Upgrading factors to consider" and "Automatically installing updates."



NOTE: When you replicate data for a cluster, you must replicate the entire cluster. For example, if you select a node to replicate, the cluster is automatically selected. Likewise, if you select the cluster, all nodes in that cluster are also selected.

Performing replication in Rapid Recovery involves performing the following operations:

- Set up a repository on the target Core. For more information on adding a repository to the target Core, see [Creating a DVM repository](#).
- Configure self-managed replication. For more information on replicating to a self-managed target Core, see [Replicating to a self-managed target Core](#).
- Configure third-party replication. For more information on replicating to a third-party target Core, see [Replicating to a third-party target Core](#).
- Replicate an existing protected machine. For more information on replicating a machine that is already protected by the source Core, see [Adding a machine to existing replication](#).
- Consume the seed drive. For more information on consuming seed drive data on the target Core, see [Consuming the seed drive on a target Core](#).
- Set the replication priority for a protected machine. For more information on prioritizing the replication of protected machines, see [Setting replication priority for a protected machine](#).
- Set a replication schedule for a protected machine. For more information on setting a replication schedule, see [Scheduling replication](#).
- Monitor replication as needed. For more information on monitoring replication, see [Viewing incoming and outgoing replication](#).
- Manage replication settings as needed. For more information on managing replication settings, see [Managing replication settings](#).
- Recover replicated data in the event of disaster or data loss. For more information on recovering replicated data, see [Recovering replicated data](#).

Replicating to a self-managed target Core

This configuration applies to replication to an off-site location and to mutual replication. The following steps are prerequisite:

- The Rapid Recovery Core must be installed on all source and target machines.
- If you are configuring Rapid Recovery for multi-point to point replication, you must perform this task on all source Cores and the one target Core. For descriptions of these replication configurations, see [Replication](#).
- If you need to create a seed drive and transfer it to a physical removable storage volume to perform the initial transfer of existing recovery points, you must have a suitable portable storage device prepared. You must also have physical access to the source Core machine, to attach the drive to copy the seed drive archive.
- If using a seed drive in a self-managed target Core, you or a trusted administrator must have physical access to the target Core.
- If using an Azure repository on a target Core, you have access and permissions for the Azure container.

A self-managed target Core is a Core to which you have access. For example, a self-managed Core is often managed by your company at an off-site location, or is hosted at a different geographic location than the source Core. Replication can be set up entirely on the source Core, unless you choose to seed your data using a seed drive. In such cases, you must create a seed drive using this procedure, and later attach the seed drive to the target Core to consume the archived recovery point data. For more information, see [Determining your seeding needs and strategy](#).

Complete the steps in the following procedure to configure your source Core to replicate to a self-managed target Core.



1. Navigate to the Rapid Recovery Core Console of the source Core.
2. On the button bar, click  **Replicate**.
The Replication Wizard appears.
3. On the Target Core page of the Replication Wizard, if you are establishing replication with a target Core that has been paired with this source Core previously, select **Use an existing target Core**, and then select the appropriate target Core from the drop-down list. Skip to [step 5](#).
4. On the Target Core page of the Replication Wizard, if you are establishing replication with a target Core from this source Core for the first time, select **I have my own Target Core**, and then enter the information as described in the following table.

Table 90: Target Core information

Text Box	Description
Host Name	Enter the host name or IP address of the Core machine to which you are replicating.
Port	Enter the port number on which the Rapid Recovery Core will communicate with the machine. The default port number is 8006.
User Name	Enter the user name for accessing the machine.
Password	Enter the password for accessing the machine.

5. Click **Next**.

 **NOTE:** If no repository exists on the target Core, a warning appears notifying you that you can pair the source Core with the target Core, but that you are unable to replicate agents (protected machines) to this location until a repository is established. For information about how to set up a primary repository to a Core, see [Creating a DVM repository](#).
6. On the Request page, enter a name for this replication configuration; for example, SourceCore1. This is the display name used for the Incoming Replication pane on the target Core's *Replication* page.
7. Click **Next**.
8. On the Protected Machines page, select the protected machines you want to replicate, and then use the drop-down lists in the Repository column to select a repository for each protected machine.

9. If you want to perform the seeding process for the transfer of the base data, complete the following sub-steps. If you do not want to seed your data, proceed to [step 10](#).

i **NOTE:** Because large amounts of data need to be copied to the portable storage device, Quest recommends using an eSATA, USB 3.0, or other high-speed connection to the portable storage device.

- a. On the Protected Machines page of the Replication Wizard, select **Use a seed drive to perform the initial transfer**.

- If you currently have one or more protected machines replicating to a target Core, you can include these protected machines on the seed drive by selecting Include already replicated recovery points in the seed drive.
- If you do not want replication to begin immediately after completing this procedure, select **Initially pause replication**.

i **NOTE:** If you select this option, replication does not begin until you explicitly resume it. For more information, see [Pausing and resuming replication](#).


- b. Click **Next**.

- c. On the Seed Drive Location page of the Replication Wizard, use the **Location type** drop-down list to select from the following destination types:

- Local
- Network
- Cloud

- d. In the Location field, enter the details for the seed drive file as described in the following table, based on the location type you selected in [step c](#).

Table 91: Archive details

Option	Text Box	Description
Local	Output location	Enter the location for the output. It is used to define the location path where you want the seed drive archive to reside; for example, D:\work\archive.
Network	Output location	Enter the location for the output. It is used to define the location path where you want the seed drive archive to reside; for example, \\servername\sharename.
	User Name	Enter a user name. It is used to establish logon credentials for the network share.
	Password	Enter a password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list. <div> NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account.</div>
	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]

- e. Click **Next**.

- f. On the Seed Drive Options page of the Replication Wizard, enter the information as described in the following table.

Table 92: Seed drive options

Item	Description
Maximum Size	<p>Large archives of data can be divided into multiple segments. Select the maximum size of the segment you want to reserve for creating the seed drive by doing one of the following:</p> <ul style="list-style-type: none">• Select Entire Target to reserve all available space in the path provided on the Seed Drive Location page for future use. For example, if the location is <code>D:\work\archive</code>, all of the available space on the <code>D:</code> drive is reserved if required for copying the seed drive, but is not reserved immediately after starting the copying process.• Select the text box, enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve. The default is 250MB.
Recycle action	<p>In the event the path already contains a seed drive, select one of the following options:</p> <ul style="list-style-type: none">• Do not reuse. Does not overwrite or clear any existing seed data from the location. If the location is not empty, the seed drive write will fail.• Replace this Core. Overwrites any pre-existing seed data pertaining to this Core but leaves the data for other Cores intact.• Erase completely. Clears all seed data from the directory before writing the seed drive.
Comment	Enter a comment that describes the seed drive.
Add all agents to seed drive	Select this option to replicate all protected machines on the source Core using the seed drive. This option is selected by default.
Build recovery point chains (fix orphans)	Select this option to replicate the entire recovery point chain to the seed drive. This adds some already replicated recovery points to the seed drive, preventing orphans from occurring. This option is selected by default.

- g. Do one of the following:
- If you cleared the **Add all agents to seed drive** check box, click **Next**.
 - If you selected **Add all agents to seed drive**, go to [step 10](#).
10. Click **Finish**.
11. If you created a seed drive, send it to your target Core.
The pairing of the source Core to the target Core is complete.

Unless you selected the option to initially pause replication, the replication process begins immediately.

1. If you selected the option to use a seed drive, replication produces orphaned recovery points on the target Core until the seed drive is consumed and provides the necessary base images.
2. If you specified the use of the a drive, transfer the seed drive archive file to a volume (shared folder, virtual disk, or removable storage media). Then, consume the seed drive.

Replicating to a third-party target Core

A third-party Core is a target Core that it managed and maintained by an MSP. Replicating to a Core managed by a third party does not require the customer to have access to the target Core.

The process of replicating to a third-party Core involves tasks that must be completed by the customer as well as the third party. After a customer submits a request for replication on the source Core or Cores, the MSP must complete the configuration on the target Core by reviewing the request.


i NOTE: This configuration applies to Hosted and Cloud Replication. The Rapid Recovery Core must be installed on all source Core machines. If you are configuring Rapid Recovery for Multi-Point to Point replication, you must perform this task on all source Cores.

To replicate to a target Core managed by a third party, complete the following tasks:

- [Submitting a replication request to a third-party service provider](#)
- [Reviewing a replication request from a customer](#)
- [Ignoring a replication request from a customer](#)

Submitting a replication request to a third-party service provider

If you are an end user who subscribes to a Core managed by a third party, such as an MSP, complete the steps in this procedure to submit a replication request to your third-party service provider.

1. From the icon button bar of the Rapid Recovery Core Console, click  **Replicate**. The Replication Wizard appears.

2. On the *Target Core* page of the Replication Wizard, select **I have a subscription to a third-party providing off-site backup and disaster recovery services**, and then enter the information as described in the following table.

Table 93: Third-party target Core information

Text Box	Description
Host Name	Enter the host name, IP address, or FQDN for the third-party Core machine.
Port	Enter the port number that was given to you by your third-party service provider. The default port number is 8006.

If the Core you want to add has been paired with this source Core previously, you can do the following:

- a. Select **Use an existing target Core**.
 - b. Select the target Core from the drop-down list.
 - c. Click **Next**.
 - d. Skip to [step 7](#).
3. Click **Next**.
 4. On the *Request* page of the Replication Wizard, enter the information as described in the following table.



Table 94: Third-party target Core details

Text Box	Description
Email Address	Enter the email address associated with your third-party service subscription.
Customer ID (optional)	Optionally, enter the customer ID that was assigned to you by the service provider.

5. Click **Next**.
6. On the *Protected Machines* page of the Replication Wizard, select the protected machines you want to replicate to the third-party Core.

7. If you want to perform the seeding process for the transfer of base data, complete the following steps.
- Note:** Because large amounts of data need to be copied to the portable storage device, an eSATA, USB 3.0, or other high-speed connection to the portable storage device is recommended.
- On the *Protected Machines* page of the Replication Wizard, select **Use a seed drive to perform initial transfer**.
 - If you currently have one or more protected machines replicating to a target Core, you can include these machines on the seed drive by selecting the option **Include already replicated recovery points in the seed drive**.
 - Click **Next**.
 - On the *Seed Drive Location* page of the Replication Wizard, use the **Location type** drop-down list to select from the following destination types:
 - Local
 - Network
 - Cloud
 - Enter the details for the archive as described in the following table, based on the location type you selected in [step c](#).

Table 95: Archive details

Option	Text Box	Description
Local	Output location	Enter the location for the output. It is used to define the location path where you want the seed drive archive to reside; for example, D:\work\archive.
	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, \\servername\sharename.
	User Name	Enter a user name. It is used to establish logon credentials for the network share.  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault .
Network	Password	Enter a password for the network path. It is used to establish logon credentials for the network share.
	Account	Select an account from the drop-down list.  NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account .
	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]
Cloud	Account	Select an account from the drop-down list.
Cloud	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]

- e. Click **Next**.
- f. On the *Seed Drive Options* page of the Replication Wizard, enter the information as described in the following table.

Table 96: Seed drive options

Item	Description
Maximum Size	<p>Large archives of data can be divided into multiple segments. Select the maximum amount of space you want to reserve for creating the seed drive by doing one of the following:</p> <ul style="list-style-type: none"> • Select Entire Target to reserve all available space in the path provided on the <i>Seed Drive Location</i> page (for example, if the location is D:\work\archive, all of the available space on the D: drive is reserved). • Select the blank text box, enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve.
Recycle action	<p>In the event the path already contains a seed drive, select one of the following options:</p> <ul style="list-style-type: none"> • Do not reuse. Does not overwrite or clear any existing seed data from the location. If the location is not empty, the seed drive write will fail. • Replace this Core. Overwrites any pre-existing seed data pertaining to this Core but leaves the data for other Cores intact. • Erase completely. Clears all seed data from the directory before writing the seed drive.
Comment	Enter a comment that describes the seed drive.
Add all Agents to Seed Drive	Select this option to replicate all protected machines on the source Core using the seed drive. This option is selected by default.
Build recovery points chains (fix orphans)	Select this option to replicate the entire recovery point chain to the seed drive. This option is selected by default.

- g. Do one of the following:
 - If you cleared the **Add all Agents to Seed Drive** check box, click **Next**.
 - If you selected **Add all Agents to Seed Drive**, go to [step 9](#).
 - h. On the *Machines* page of the Replication Wizard, select the protected machines you want to replicate to the target Core using the seed drive.
8. Click **Finish**.
 9. If you created a seed drive, send it as directed by your third-party service provider.

Reviewing a replication request from a customer


After an end user completes the procedure [Submitting a replication request to a third-party service provider](#), a replication request is sent from the source Core to the third-party target Core. As the third party, you can review the request, and then approve it to begin replication for your customer, or you can deny it to prevent replication from occurring.


Choose from the following options:

- [Approving a replication request](#)
- [Denying a replication request](#)

Approving a replication request

Complete the following procedure to approve a replication request on a third-party target Core.


1. On the target Core, navigate to the Rapid Recovery Core Console.
2. From the icon bar, click  (Replication).
The Replication page appears.
3. On the Replication page, click **Request (#)**.
The Pending Replication Requests section appears.
4. Under **Pending Replication Requests**, click the drop-down menu next to the request you want to review, and then click **Review**.
The Review Replication Request window appears.

 **NOTE:** The information that appears in the Source Core Identity section of this window is determined by the request completed by the customer.
5. Under Source Core Identity, do one of the following:
 - **Replace an existing replicated Core**, and then select a Core from the drop-down list.
 - Select **Create a new source Core**, and then confirm that the Core Name, customer Email Address, and Customer ID, provided are correct. Edit the information as necessary.
6. Under Agents, select the machines to which the approval applies, and then use the drop-down lists in the Repository column to select the appropriate repository for each machine.
7. Optionally, in the **Comment** text box, enter a description or message to include in the response to the customer.
8. Click **Send Response**.
Replication is accepted.

Denying a replication request

Complete the steps in the following procedure to deny a replication request sent to a third-party Core from a customer.

To deny a request without reviewing it, see [Ignoring a replication request from a customer](#).


1. On the target Core, navigate to the Rapid Recovery Core Console.
2. From the icon bar, click  (Replication).
The Replication page appears.
3. On the *Replication* page, click **Request (#)**.
The Pending Replication Requests section appears.
4. Under **Pending Replication Requests**, click the drop-down menu next to the request you want to review, and then click **Review**.
The *Review Replication Request* window appears.
5. Click **Deny**.
Replication is denied. Notification of denial appears under Alerts on the *Events* page of the source Core.

Ignoring a replication request from a customer

As a third-party service provider of a target Core, you have the option to ignore a request for replication sent from a customer. This option could be used if a request was sent by mistake or if you want to deny a request without reviewing it.

For more information about replication requests, see [Reviewing a replication request from a customer](#).


Complete the following procedure to ignore a replication request from a customer.

1. On the target Core, navigate to the Rapid Recovery Core.
2. From the icon bar, click  (Replication).
The *Replication* page appears.
3. On the *Replication* page, click **Request (#)**.
The *Pending Replication Requests* section appears.
4. Under *Pending Replication Requests*, click the drop-down menu next to the request you want to ignore, and then click **Ignore**.
5. On the *Ignoring request* dialog box, click **Yes** to confirm the command.
Notification that the request has been ignored is sent to the source Core, and the request is removed from the *Replication* page on the target Core.

Adding a machine to existing replication

After replication is established between a source and target Core, it is possible to add new protected machines to replicate to the target. Complete the steps in the following procedure on the source Core to add a new protected machine to a paired target Core for replication.

For more information about replication, see [Replication](#) and [Replicating to a self-managed target Core](#).

1. Navigate to the Rapid Recovery Core console of the source Core.
2. On the button bar, click  Replicate.
The Replication Wizard opens to the *Protected Machines* page.

3. On the *Protected Machines* page, select the protected machines you want to replicate, and then use the drop-down lists in the Repository column to select a repository for each protected machine.



4. If you want to perform the seeding process for the transfer of the base data, complete the following steps:

i **NOTE:** Because large amounts of data need to be copied to the portable storage device, an eSATA, USB 3.0, or other high-speed connection to the portable storage device is recommended.

- a. On the *Protected Machines* page of the Replication Wizard, select **Use a seed drive to perform initial transfer**.
 - If you currently have one or more protected machines replicating to a target Core, you can include these machines on the seed drive by selecting the option **Include already replicated recovery points in the seed drive**.
- b. Click **Next**.
- c. On the *Seed Drive Location* page of the wizard, use the **Location type** drop-down list to select from the following destination types:
 - Local
 - Network
 - Cloud

- d. Enter the details for the archive as described in the following table based on the location type you selected in [step c](#).

Table 97: Archive details

Option	Text Box	Description
Local	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, D:\work\archive.
Network	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, \\servername\sharename.
	User Name	Enter a user name. It is used to establish logon credentials for the network share.  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault .
	Password	Enter a password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list.  NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account .
	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]

- e. Click **Next**.

- f. On the *Seed Drive Options* page of the wizard, enter the information as described in the following table.

Table 98: Seed drive options

Item	Description
Maximum Size	<p>Large archives of data can be divided into multiple segments. Select the maximum amount of space you want to reserve for creating the seed drive by doing one of the following:</p> <ul style="list-style-type: none">• Select Entire Target to reserve all available space in the path provided on the <i>Seed Drive Location</i> page (for example, if the location is <code>D:\work\archive</code>, all of the available space on the <code>D:\</code> drive is reserved).• Select the blank text box, enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve.
Recycle action	<p>In the event the path already contains a seed drive, select one of the following options:</p> <ul style="list-style-type: none">• Do not reuse. Does not overwrite or clear any existing seed data from the location. If the location is not empty, the seed drive write will fail.• Replace this Core. Overwrites any pre-existing seed data pertaining to this Core but leaves the data for other Cores intact.• Erase completely. Clears all seed data from the directory before writing the seed drive.
Comment	Enter a comment that describes the seed drive.
Add all Agents to Seed Drive	Select this option to replicate all protected machines on the source Core using the seed drive. This option is selected by default.
Build recovery point chains (fix orphans)	Select this option to replicate the entire recovery point chain to the seed drive. This option is selected by default.

- g. Do one of the following:
- If you cleared the **Add all Agents to Seed Drive** check box, click **Next**.
 - If you selected **Add all Agents to Seed Drive**, go to [step 5](#).
- h. On the *Protected Machines* page of the wizard, select the protected machines you want to replicate to the target Core using the seed drive.
5. Click **Finish**.

Consuming the seed drive on a target Core

Complete the follow procedure to consume the data from the seed drive file on the target Core.

This procedure is only necessary if a seed drive file was created as part of [Replicating to a self-managed target Core](#) or [Replicating to a third-party target Core](#).

1. If the seed drive file was saved to a portable storage device, such as a USB drive, connect the drive to the target Core.

i **NOTE:** If the seed data is comprised of multiple segments, all segments must be available in the location type you specify in step 4. For example, all segments should be attached to the local target Core server, in the same network directory, or in the same cloud account container.


2. On the target Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication). The *Replication* page appears.
3. On the *Replication* page, under **Incoming Replication**, click the drop-down menu for the correct source Core, and then select **Consume**. The *Consume* dialog box appears.
4. In the **Location type** field, select one of the following options from the drop-down list:
 - Local
 - Network
 - Cloud
5. Enter the details for the seed drive archive file, as described in the following table based on the location type you selected in [step 4](#).

Table 99: Archive details

Option	Text Box	Description
Local	Location	Enter the path for the archive.
Network	Location	Enter the path for the archive.
	User name	Enter the user name. It is used to establish logon credentials for the network share.
	Password	Enter the password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list. i NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account .
	Container	Select a container associated with your account from the drop-down menu.
	Folder name	Enter the name of the folder in which the archived data is saved; for example, Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]

6. Click **Check File**.
The Core searches for the file.
After finding the file, the following text boxes appear in the *Consume* dialog box, pre-populated with the data gathered from [step 4](#), [step 5](#), and the file.
 - The **Date Range** displays the dates of the oldest and newest recovery points contained in the seed drive.
 - Any comments entered when the seed drive was created are automatically imported.
7. On the **Consume** dialog box, under Agents, select the machines for which you want to consume data.
8. Click **Consume**.
9. To monitor the progress of consuming data, view the [Events](#) page.

Abandoning a seed drive

If you create a seed drive with the intent of consuming it on the target Core, but later choose not to consume it, you can abandon the seed drive.



Until you abandon the seed drive or consume it, a link for the outstanding seed drive remains on the Outgoing Replication pane on the source Core.

Until you transmit information from the seed drive, orphaned recovery points (which exist on the original protected machine, but not the target Core) cannot be used to restore data.

CAUTION: If you abandon the seed drive, then the original recovery points (defined in the seed drive file) are then transmitted over the network to the target Core during the next replication job. Transmitting old recovery points over the network could slow down the network considerably, especially if there are many recovery points.

Complete the steps in the following procedure to abandon an outstanding seed drive.

NOTE: Abandoning the seed drive in the Core Console does not affect the seed drive file from its storage location.

1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click [Replication](#) (Replication). The Replication page appears.
2. On the Replication page, in the Outgoing Replication pane, click **Seed Drives (#)**.
In the Outgoing Replication pane, a section appears containing information about the outstanding seed drives.
3. Optionally, click the arrow  to expand the collapsible menu.
Information appears about outstanding seed drives, including the target Core and the date range of the recovery points included in the seed drive.
4. For the seed drive file you want to abandon, click  (More options) and then select **Abandon**.
5. In the confirmation window, confirm that you want to abandon the seed drive.
The seed drive is removed.
If no more seed drives exist on the source Core, the Seed Drives (#) link and outstanding seed drives section are removed from the Outgoing Replication pane.

Managing replication settings

Rapid Recovery Core lets you monitor, schedule, and adjust replication at the overall, Core, and protected machine levels.



You can edit the following replication settings:

- To schedule replication jobs, see [Scheduling replication](#).
- To create a seed drive of a protected machine that is already paired for replication, see [Using the Copy function to create a seed drive](#).
- To monitor the progress of a replication job, see [Viewing incoming and outgoing replication](#).
- To pause or resume a paused replication job, see [Pausing and resuming replication](#).
- To force replication of an incoming or outgoing protected machine, see [Forcing replication](#).
- To manage settings for all target Cores and replication procedures, see [Managing settings for outgoing replication](#).
- To manage settings for an individual target Core, see [Changing target Core settings](#).
- To manage priority settings for an individual protected machine being replicated to a target Core, see [Setting replication priority for a protected machine](#).

Scheduling replication

Unless you change the default behavior by setting a replication schedule, the Core starts a replication job immediately after completion of every backup snapshot, checksum check, attachability check, and the nightly jobs. You can change the replication schedule to reduce the load on the network.

Complete the steps in the following procedure to set a replication schedule for any replicated machine.

1. On the target Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication). The Replication page appears.
2. In the Outgoing Replication pane, on the Core for which you want schedule replication, click  (More options) and then select **Schedule**. The Replication Schedule for [CoreName] dialog box opens.

3. Select from one of the following options:

- **At All Times.** Replicates after every new snapshot, checksum check, and attachability check, and after the nightly jobs complete.
- **Daily (Start replication only during the specified time period).** Begins replicating only within the time range provided.

a. In the **From** text box, enter the earliest time at which replication should begin.

b. In the **To** text box, enter the latest time at which replication should begin.



NOTE: If replication is in progress when the scheduled time ends, the replication job completes after the allotted time period.

- **Custom.** Begins replicating only within the time range provided, letting you set one time range for weekdays and another time range for weekends.

a. Next to Weekdays, in the **From** text box, enter the earliest time at which replication should occur on a weekday; and then in the **To** text box, enter the latest time at which replication should occur on a weekday.



b. Next to Weekends, in the **From** text box, enter the earliest time at which replication should occur on weekends; and then in the **To** text box, enter the latest time at which replication should occur on weekends.

4. Click **Save**.

The schedule is applied to all replication to the selected target Core.


Using the Copy function to create a seed drive

If you chose not to create a seed drive when you configured replication, you can create a seed drive using the Copy function in the protected machine drop-down menu.

1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication). The *Replication* page appears.
2. On the *Replication* page, in the Outgoing Replication pane, click  to expand the Core that protects the machine for which you want to create a seed drive. The selection expands to show each protected machine in the specified Core.
3. Click in the first row of the table to select each machine for which you want to create a seed drive.
4. In the menu under the Outgoing Replication pane, click **Copy**. The Replication Wizard appears.
5. On the *Seed Drive Location* page of the wizard, use the **Location** drop-down list to select from the following destination types:
 - Local
 - Network
 - Cloud

6. Enter the details for the seed drive archive, as described in the following table, based on the location type you selected in the preceding step.


Table 100: Archive details

Option	Text Box	Description
Local	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, D:\work\archive.
Network	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, \\servername\sharename.
	User Name	Enter a user name. It is used to establish logon credentials for the network share.
	Password	Enter a password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list. <div> NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account.</div>
	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED]

7. Click **Next**.

8. On the *Seed Drive Options* page, enter the information as described in the following table.

Table 101: Seed drive options

Item	Description
Maximum Size	<p>Large archives of data can be divided into multiple segments. Select the maximum amount of space you want to reserve for creating the seed drive by doing one of the following:</p> <ul style="list-style-type: none">• Select Entire Target to reserve all available space in the path provided on the <i>Seed Drive Location</i> page (for example, if the location is D:\work\archive, all of the available space on the D: drive is reserved).• Select the blank text box, enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve.
Recycle action	<p>In the event the path already contains a seed drive, select one of the following options:</p> <ul style="list-style-type: none">• Do not reuse. Does not overwrite or clear any existing seed data from the location. If the location is not empty, the seed drive write will fail.• Replace this Core. Overwrites any pre-existing seed data pertaining to this Core but leaves the data for other Core s intact.• Erase completely. Clears all seed data from the directory before writing the seed drive.
Comment	Enter a comment that describes the seed drive.
Add all Agents to Seed Drive	Select this option to replicate all protected machines on the source Core using the seed drive. This option is selected by default.
Build recovery point chains (fix orphans)	<p>Select this option to replicate the entire recovery point chain to the seed drive. This option is selected by default.</p> <div> NOTE: Typical seeding in Rapid Recovery 5.4.x replicated only the latest recovery point to the seed drive, which reduced the amount of time and space required for creating the seed drive. Opting to build recovery point chains to the seed drive requires enough space on the seed drive to store the latest recovery points from the specified protected machines, and may take additional time to complete the task.</div>

9. Do one of the following:


- If you cleared the **Add all Agents to Seed Drive** check box, click **Next**.
- If you selected **Add all Agents to Seed Drive**, go to [step 10](#).

If you selected **Add all Agents to Seed Drive**, go to [step 10](#).

10. On the *Protected Machines* page of the wizard, select the protected machines for which you want to create a seed drive.
11. Click **Finish**.

Monitoring replication

When replication is set up, you can monitor the status of replication tasks for the source and target Cores. You can refresh status information, view replication details, and more.

1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication). The Replication page appears.

2. On this page, you can view information about and monitor the status of replication tasks as described in the following table.

Table 102: Replication tasks

Section	Description	Available Actions
Seed Drives (#)	<p>After you specify the use of a seed drive when defining replication, until you abandon or consume it, a Seed Drives (#) link appears on the Outgoing Replication pane on the source Core. The number displayed indicates how many seed drives are pending.</p> <p>i NOTE: This link does not appear unless a seed drive is pending.</p> <p>Click this link to list seed drives that have been written but not yet consumed by the target Core. Further expand the collapsible menu to show information about outstanding seed drives, including the target Core and the date range of the recovery points included in the seed drive.</p>	<p>In the drop-down menu, click Abandon to abandon or cancel the seed process.</p>
Outgoing Replication	<p>Lists all target Cores to which the source Core is replicating. It includes a state indicator, the target Core name, the number of machines being replicated, and the progress of a replication transmission.</p>	<p>On a source Core, from the ... (More) drop-down menu, you can select the following options:</p> <ul style="list-style-type: none">• Details. Lists the ID, URL, display name, state, customer ID, email address, and comments for the replicated Core.• Change Settings. Lists the display name and lets you edit the host and port for the target Core.• Schedule. Lets you set a customized schedule for replication to this target Core.




Section	Description	Available Actions
		<ul style="list-style-type: none"> • Add Machines. Lets you select a host from a drop-down list, select protected machines for replication, and create a seed drive for the new protected machine's initial transfer. You can optionally choose to include recovery points for machines already added to replication. • Delete. Lets you delete the replication relationship between source and target Cores. Doing so ceases all replication to this Core.
Incoming Replication	<p>Lists all source machines from which the target receives replicated data. It includes the remote Core name, state, machines, and progress.</p> <p>Lists all source Cores from which the target receives replicated data. The display name for the source Cores listed are populated from the value in the Replication Wizard when defining replication. It includes a state indicator, the remote Core name, and the progress of a replication transmission.</p>	<p>On a target Core, from the ... (More) drop-down menu, you can select the following options:</p> <ul style="list-style-type: none"> • Details. Lists the ID, host name, customer ID, email address, and comments for the replicated Core. • Consume. Consumes the initial data from the seed drive and saves it to the local repository. • Delete. Lets you delete the replication relationship between target and source Cores. Doing so ceases all replication from this Core.
Pending Replication Requests	<p>This information applies to managed service providers only. When a customer clicks the Requests link in the Incoming Replication pane, a summary table section appears listing the customer ID, email address, and host name for the request.</p>	<p>In the drop-down menu, click Ignore to ignore or reject the request, or Review to review the pending request.</p>

Pausing and resuming replication

You can pause replication temporarily for the source (outgoing) or target (incoming) Cores.

The option to pause replication is only available when replication is active. The option to resume replication is only available if replication is paused.


Complete the steps in the following procedure to pause or resume replication.

1. Open the Rapid Recovery Core Console, and from the icon bar, click  (Replication).
The Replication page appears.
2. To pause replication for all replicated machines, do the following:
 - a. Click the checkbox at the top of the summary table to select the source or target Core.
 - b. Click Pause from the menu preceding the summary table.
Replication for all protected machines in the selected Core is paused.
3. To pause replication for only certain machines, do the following:
 - a. Click the  arrow to the right of any Core.
The view expands to show each of the protected machines from the selected Core that are being replicated.
 - b. Click in the first column to select each machine for which you want to pause replication. Click any selection again to clear the checkbox for machines you do not want to pause.
 - c. Click Pause from the menu preceding the summary table.
Replication for the selected protected machines is paused.
4. To resume replication for all replicated machines, do the following:
 - a. Click the checkbox at the top of the summary table to select the source or target Core.
 - b. Click Resume from the menu at the top of the summary table.
Replication for all protected machines in the selected Core is resumed.
5. To resume replication for only certain machines, do the following:
 - a. Click the  arrow to the right of any Core.
The view expands to show each of the protected machines from the selected Core that are being replicated.
 - b. Click in the first column to select each machine for which you want to resume replication. Click any selection again to clear the checkbox for machines you do not want to resume.
 - c. Click Resume from the menu at the top of the summary table.
Replication for the selected protected machines is resumed.

Forcing replication

From the source Core, you can force replication at any time, instead of waiting for a replication job to queue after a specific event such as a backup or attachability check.

Complete the steps in the following procedure to force replication on either the source or the target Core.

1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication).
The *Replication* page appears.

2. Do one of the following:
 - To force replication on a source Core, from the Outgoing Replication pane, select a Core, and from the menu at the top of the summary table, click **> Force**.
 - To force replication on a target Core, from the Incoming Replication pane, select a Core, and from the menu at the top of the summary table, click **> Force**.

The *Force Replication* dialog box appears.

3. Optionally, if you want to repair any orphaned chains of recovery points, select **restore orphaned recovery point chains**.
4. To confirm, in the *Force Replication* dialog box, click **Yes**.
The dialog box closes, and replication is forced.

Managing settings for outgoing replication

The changes made to these settings affect the data transfer to all target Cores associated with this source Core.




1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication).
The Replication page appears.
2. In the Outgoing Replication pane, at the top of the summary table, click  (Settings).
The Replication Settings dialog box appears.
3. In the Replication Settings dialog box, edit the replication settings as described in the following table.

Table 103: Replication settings

Option	Description
Cache lifetime (seconds)	Specify the amount of time between each target Core status request performed by the source Core.
Volume image session timeout (minutes)	Specify the amount of time the source Core spends attempting to transfer a volume image to the target Core.
Maximum parallel streams	Specify the number of network connections permitted to be used by a single protected machine to replicate that machine's data at one time.
Maximum transfer speed (MB/s)	Specify the speed limit for transferring the replicated data.
Maximum transfer data size (GB)	Specify the maximum size in GB for transferring blocks of replicated data.
Restore Defaults	Select this option to change all replication settings to the system defaults. <div> NOTE: Take note of any customized settings before selecting this option. You will not be prompted to confirm this action</div>

4. When satisfied, click **Save** to save the replication settings and close the dialog box.

Changing target Core settings

You can change the host and port settings for individual target Cores from the source Core.




1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication).
The Replication page appears.
In the Outgoing Replication pane, the summary table includes a row for each target Core that has been configured to replicate recovery points from this source Core.
2. Click the  (Settings) drop-down menu for the target Core you want to modify, and then select **Change Settings**.
The Settings dialog box appears.
3. Edit any of the options described in the following table.

Table 104: Target Core settings

Option	Description
Host	Enter the host for the target Core.
Port	Enter a port for the target Core to use for communication with the source Core.  NOTE: The default port is 8006.

4. Click **Save**.

Setting replication priority for a protected machine



Replication priority determines which replication jobs are sent to the Core first. Prioritization is set ordinally, on a scale of 1 to 10, where a priority of 1 is the first priority, and a priority of 10 is the last priority. When you first establish replication for any machine, its priority is set to 5. You can view and change priority at the protected machine level from the source Core.


In some cases, it is possible that some replication jobs are abandoned. For example, replication jobs can be abandoned if your environment is experiencing unusually high change rates or if your network does not have enough bandwidth. This situation is particularly likely if you set schedules which limit the hours when replication occurs in your environment. For more information about setting schedules replication, see [Scheduling replication](#).

To ensure replication occurs for important machines first, set critical servers to a priority with a lower number (between 1 and 5). Set priority for less important machines to a higher number (between 6 and 10).

Setting replication priority to 4 for any protected machine assures its replication job is started before a machine with the default replication priority of 5. Replication jobs for machines with a priority of 3 are queued before 4, and so on. The lower the priority number, the sooner its replication jobs are sent. It is easy to remember that priority 1 is most important. Machines with a replication priority of 1 are the first machines queued for replication.





Complete the steps below to edit the settings that prioritize when a protected machine replicates.

1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication).
The Replication page appears.
2. In the Outgoing Replication pane, click the  arrow to the right of any source Core.
The view expands to show each of the protected machines from this source Core that are being replicated to the designated target Core.

3. Click the  (More) drop-down menu for the protected machine you want to prioritize, and then click Change Settings.
A dialog box appears.
4. Click the Priority drop-down list and select a priority, from 1 (Highest) to 10 (Lowest), based on your requirements.
5. Click Save.
The dialog box closes, and the replication priority for the selected machine updates.

Removing outgoing replication from the source Core

Complete the steps in this procedure to remove one or more protected machines from replication on the source Core.

1. On the source Core, open the Rapid Recovery Core Console, and from the icon bar, click  (Replication).
The Replication page appears.
In the Outgoing Replication pane, the summary table includes a row for each target Core that has been configured to replicate recovery points from this source Core.
2. Optionally, click the  arrow to the right of any target Core.
The view expands to show each of the protected machines from this source Core that are being replicated to the designated target Core.
3. Select the protected machines you want to remove from outgoing replication as follows:
 - To completely remove the existing replication relationship between this source Core and any target Core, click the  (More) drop-down menu for any target Core, and then select **Delete**.
 - To remove outgoing replication for a subset of machines in the specified target Core, expand the view to show all machines being replicated, and select the check box for each replicated machine that you want to remove. Clear the checkbox for any machine you want to continue replicating. Then, from the menu above the summary table, click  Delete.

A confirmation message appears asking if you want to remove the replication relationships.
4. In the resulting dialog box, click **Yes** to confirm removal.

For more information, see the following related topics:

- [Replication with Rapid Recovery](#)
- [Removing incoming replication from the target Core](#)
- [Removing outgoing replication from the source Core](#)

Removing incoming replication from the target Core

Complete the steps in this procedure to remove one or more protected machines from replication on the target Core.



NOTE: You can also remove replication of protected machines from the Outgoing Replication pane on the *Replication* page of the source Core. For more information, see [Removing outgoing replication from the source Core](#).

1. On the target Core, open the Rapid Recovery Core Console, and from the icon bar, click (Replication). The *Replication* page appears. In the Incoming Replication pane, the summary table includes a row for each source Core with protected machines that this target Core replicates.
2. Select the replicated machines to remove as follows:
 - To delete **all** machines replicated from the source Core to your target Core, in the Incoming Replication pane, select the check box for that Core.
 - To delete a smaller subset of machines from the same source Core, do the following:
 - a. Click the arrow to the right of the source Core. The view expands to show each of the machines from the selected source Core that are replicated on your target Core.
 - b. Select the check box for each replicated machine that you want to remove.
 - c. From the parent row of the selected source Core, click the (More) drop-down menu, and then select **Delete**. The *Remove Replication* dialog box appears.
3. In the *Remove Replication* dialog box, do one of the following:
 - If you want to leave the replicated recovery points on the target Core, clear the option **Delete existing recovery points**.
 - If you want to delete all replicated recovery points received from that machine as well as remove the source Core from replication, select **Delete existing recovery points**.
4. Click **Yes** to confirm deletion.



CAUTION: If you select this option, all of the recovery points replicated to this Core will be deleted.

The selected machines protected on the source Core are removed from replication on this target Core. Optionally, if you selected the option to delete recovery points, they are removed from the repository of this Core.

For more information, see the following related topics:

- [Replication with Rapid Recovery](#)
- [Removing incoming replication from the target Core](#)
- [Removing incoming replication from the target Core](#)

Recovering replicated data

Day-to-day replication functionality is maintained on the source Core, while only the target Core is capable of completing the functions necessary for disaster recovery.

For disaster recovery, the target Core can use the replicated recovery points to recover the protected machines. You can perform the following recovery options from the target Core:

- **Mount recovery points.** For more information, see [Mounting a recovery point](#).
- **Roll back to recovery points.** For more information, see [About restoring volumes from a recovery point or Restoring volumes for a Linux machine using the command line](#).
- **Perform a virtual machine (VM) export.** For more information, see [Exporting to virtual machines using Rapid Recovery](#).
- **Perform a bare metal restore (BMR).** For more information, see [Performing a bare metal restore using the Restore Machine Wizard](#).

Events

The Rapid Recovery Core includes predefined sets of events. These events can be used to notify administrators of critical issues on the Core or about issues with jobs pertaining to backups, virtual export, replication and so on.

This section describes how to view events displayed on the Rapid Recovery Core Console. You can also learn about event notification methods and configuration, including setting up email notifications. Finally, you can configure notifications to change the amount of time event logs are retained, and reduce repetitive event notification.

Topics include:


[Viewing events using tasks, alerts, and journal pages](#)

[Understanding event notifications in Rapid Recovery](#)

[Configuring event settings](#)

Viewing events using tasks, alerts, and journal pages

From the Core console, you can view events for the Core, and you can view events for a specific protected or replicated machine.

The Events pages on the Core Console display a log of all system events related to the Rapid Recovery Core. To access and view events for the Core, click  (Events).

The Events pages for a specific protected or replicated machine display a log of events related to that specific machine. To access and view events for a selected machine, click the machine name in the Protected Machines menu, and from the machine Summary page, click the Events menu.

Events pages (on the Core or a specified machine) are available in three views: Tasks, Alerts, and Journal.

Added to definition of Task.

All items shown in any category is an event. These views allow you to filter details about various subsets of events, as appropriate. The default view is to show tasks.

- A **task** is an event related to a job. A **job** is a process that the Rapid Recovery Core must perform. Each job has a current state, and a start and end time and date. Some tasks are initiated manually or scheduled by the user. Examples include forcing a snapshot, scheduling a backup, or performing a restore from a recovery point. Other tasks are automatic functions, such as running nightly jobs, or performing rollup using the default retention policy.
- An **alert** is a priority event, such as an error, warning, or important informational message. If you request notifications of any specific events, these notifications appear in the Alerts subset.
- The **journal** shows a complete list of all logged events (for the Core, or the selected machine, as appropriate). This list is more comprehensive, showing jobs, high priority events, and lower priority events. This category includes passive and non-job events (such as the Core starting successfully, or reporting status from the license portal).

Complete the steps in the following procedures to view tasks, alerts, or a journal of all events:





- [Viewing tasks](#)
- [Viewing alerts](#)
- [Viewing a journal of all logged events](#)
- [Viewing running tasks from any Core Console page](#)
- [Navigating between tasks, alerts, and the events journal](#)

Viewing tasks







A task is a job that the Rapid Recovery Core must perform, such as transferring data in a regularly scheduled backup, or performing a restore from a recovery point.

i **NOTE:** As a task is running, it is listed in the Running tasks drop-down menu at the top of the Core Console. Clicking a running task opens the Monitor Active Task dialog box. From here you can cancel one or more running tasks. For more details, see the topic [Viewing running tasks from any Core Console page](#). You can also suspend the scheduling of future tasks on the Core. This is useful, for example, when performing OS upgrades or server maintenance. For more information about this function, see the topic [Suspending or resuming scheduled tasks](#).

Complete the following steps to view tasks specifically for the Rapid Recovery Core, or to view tasks associated with a specific machine.

1. To view all tasks for the Rapid Recovery Core, from the icon bar, click  (Events). The default view displays all tasks for the Core.
If you want to view tasks for a specific protected machine, then navigate to the Summary page of the specified machine, and then click the **Events** menu. The default view displays all tasks for the selected machine.
2. Optionally, to filter the list of tasks by keyword, start date, end date, or any combination, do the following, and then press **Enter**:
 - a. To filter by keyword, enter the keyword in the **Search keyword** text box.
For example, you can filter by key words such as "rolling up," "archive," "export," or "transfer."
 - b. To filter by start date and time, enter the starting date and time using one of the following options:
 - In the **From** text box, type the date and time in format `MM/DD/YYYY HH:MM AM/PM`. For example, to search from the first day of January in 2016 at 8:00 AM, enter `1/1/2016 8:00 AM`.
 - To select the current date and time, click the  **Calendar** widget in the **From** text box and then click the current date. The current time automatically appears.
 - Click the  **Calendar** widget, select the date, then click the  **Clock** widget and select the desired time using the controls. Click away from the calendar to accept the selected changes.
 - c. To further refine the list of tasks that appears, you can also define an end date and time in the same format.
The list of tasks is immediately filtered based on the criteria you selected.

3. Optionally, you can filter the tasks appearing in the list as follows:

Option	Description
	To see only active tasks, click the Active tasks icon.
	To see only tasks that are in the queue awaiting performance, click the Queued tasks icon.
	To see only tasks that are waiting to be performed, click the Waiting tasks icon.
	To see only tasks that have been completed, click the Completed tasks icon.
	To see only tasks that have failed, click the Failed tasks icon.
	To see all events, including the service events for the Core that are not displayed by default, click the Service icon.



4. To export the list of tasks, select a format from the list and then click  **Export**. In the resulting dialog box, confirm the export and then click **OK**.
You can export using the following formats:

Table 105: Export formats

Format	Description
PDF	Portable Document Format (default export format)
HTML	Web page format
CSV	Comma-separated values
XLS	Microsoft Excel 1997 - 2003 Workbook
XLSX	Excel Workbook


The file of the type you selected is downloaded to the default location on the Core server.


5. Click the  **Job Details** icon for any task to launch a new window with task details.

Viewing running tasks from any Core Console page

To complete this procedure, there must be a task currently running on the Rapid Recovery Core.

Rapid Recovery Core offers a quick way to view tasks that are currently running on the Core from any page of the Core Console.

On the right side of the button bar is the **Running tasks** queue. In lower screen resolutions, or if the browser window is not fully expanded, the queue appears as an Events  icon. When one or more tasks are running, a number indicating the number of tasks currently running on the Core appears next to the queue, and the icon becomes animated. You can click the queue to reveal a drop-down list of running tasks and complete the actions described in the following procedure.

1. On the Core Console, while a task is running, click  **Running tasks**.
A small box appears, showing the type of tasks running and their progress, and presents options for canceling one or more of the tasks or for seeing more details.
2. To cancel one or more of the running tasks, complete one of the following options:
 - To cancel a single task, click the **X** beside the task description.
 - To cancel all running tasks, click **Cancel All**.
3. To view more details about a running task, click the task description.
The Monitor Active Task window opens and displays details such as progress and start time.

For more information, see the following related tasks:

- [Viewing tasks](#)
- [Navigating between tasks, alerts, and the events journal](#)
- [Suspending or resuming scheduled tasks](#)

Suspending or resuming scheduled tasks



To resume the Core scheduler, it must first be paused, as described in [step 3](#).

The Core tracks tasks based on a schedule. Transfers for backup, replication, and archiving can all be scheduled; reports can be generated weekly or monthly; and so on. When the time comes for a scheduled task to occur, the job is placed in a queue, and jobs are accomplished consecutively or concurrently, based on priority. This feature is known as the Core scheduler.

Users can now notify the Core to suspend the scheduler. When suspended, future jobs that otherwise would be scheduled automatically to run on the Core are held temporarily in a queue, and no new tasks are scheduled. This function is useful in situations such as performing OS updates, software installations, or maintenance on the Core server. When suspended, jobs accumulate in the queue but do not start until the Core scheduler function is explicitly resumed. In the interim, a banner is displayed across the Core Console indicating that the Core scheduler is paused.

i **NOTE:** This feature prevents tasks that would soon be scheduled from running. To view or cancel tasks that are already queued without suspending future tasks, see the topic [Viewing running tasks from any Core Console page](#).

Complete the following steps to suspend or resume the Core scheduler function.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Events).
The Tasks page is displayed.
3. To suspend the Core scheduler function, do the following:
 - a. Click  **Suspend Scheduler**.
 - b. In the dialog box, if you want to cancel all running tasks and suspend future tasks, select **Cancel Active Tasks and Suspend**.
 - c. To suspend future tasks only, click **Suspend**.

Tasks are suspended, and a banner is displayed on all pages of the Core Console notifying you that all scheduled tasks are suspended.

4. To resume suspended tasks, from the Tasks page, click  **Resume Scheduler**.

i | **NOTE:** When the scheduler is suspended, you can also click **Resume** at the top of the Core Console.

For more information, see the following related tasks:



- [Viewing running tasks from any Core Console page](#)
- [Viewing tasks](#)
- [Navigating between tasks, alerts, and the events journal](#)
- [Suspending or resuming scheduled tasks](#)
- [Viewing alerts](#)
- [Viewing a journal of all logged events](#)








Viewing alerts

An alert is a priority notification of an event. Any event for which you specifically requested notification appears in the list of alerts, along with errors, warnings, or important informational messages.

Rapid Recovery Core ships with a default set of events that are prioritized as alerts. You can customize the events which appear as alerts by editing the default notification group (or by setting up a new notification group). In addition to these events appearing on the Alerts page, you can change the methods used to notify you by changing your notification options. For more information on changing the events that appear as alerts, or changing notification options, see the topic [Understanding event notifications in Rapid Recovery](#).

Complete the following steps to view alerts specifically for the Rapid Recovery Core, or to view alerts associated with a specific machine.

1. To view alerts for the Rapid Recovery Core, from the icon bar, click  (Events). Click on the drop-down  menu to the right of the *Tasks* page title, and select **Alerts**.
If you want to view alerts for a specific protected machine, navigate to the Summary page of the specified machine, click the **Events** menu, and then click **Alerts**.
The list of events is filtered to display only important alerts for the Core or for the machine you selected.

2. Optionally, to filter the list of important alerts by start date, end date, alert message description, or any combination, do the following:
 - a. To filter by alert category (errors, informational messages, or warnings), click the  status drop-down menu, and select the status condition or conditions. Alert category filter options include  errors,  informational messages, and  warnings, or any combination of the three.
 - b. To filter by start date and time, enter the starting date and time using one of the following options:
 - In the **From** text box, type the date and time in format MM/DD/YYYY HH:MM AM/PM. For example, to search from the first day of January in 2016 at 8:00 AM, enter 1/1/2016 8:00 AM.
 - To select the current data and time, click the  **Calendar** widget in the **From** text box and then click the current date. The current time automatically appears.
 - Click the  **Calendar** widget, select the date, then click the  **Clock** widget and select the desired time using the controls. Click away from the calendar to accept the selected changes.
 - c. To filter by alert message description, enter the description in the **Search message** text box. For example, to see alerts only related to agents, enter "agent," or to see alerts related to transfers, enter "transfer;" and so on.
 - d. To further refine the list of alerts that appears, you can also define an end date and time in the same format.
The list of alerts is immediately filtered based on the criteria you selected.
3. Optionally, if you want to remove all alerts, click **Dismiss All**.

For more information, see the following related tasks:










- [Understanding event notifications in Rapid Recovery](#)
- [Configuring notification groups](#)
- [Viewing a journal of all logged events](#)
- [Navigating between tasks, alerts, and the events journal](#)
- [Viewing tasks](#)
- [Suspending or resuming scheduled tasks](#)
- [Viewing alerts](#)
- [Viewing a journal of all logged events](#)
- [Viewing running tasks from any Core Console page](#)

Viewing a journal of all logged events

The journal lists all logged events. This list is comprehensive, including both job- and non-job-related events. It includes specific events for which you requested notification. The journal also lists passive events and status events from the Core, the license portal, and so on.

i **NOTE:** If your environment is set to use repetition reduction, some repeated events may not be logged each time the event occurs. For more information about this feature, see [About repetition reduction](#).

Complete the following steps to view a journal of all events for the Rapid Recovery Core, or to view a journal of all events for a specific machine.

1. To view a journal of all events logged for the Rapid Recovery Core, from the icon bar, click  (Events). Click on the drop-down  menu to the right of the *Tasks* page title, and select **Journal**.
If you want to view a journal of all events for a specific protected machine, then navigate to the Summary page of the specified machine, click the **Events** menu, and then click **Journal**.
2. Optionally, to filter the list of all events by start date, end date, alert message description, or any combination, do the following:
 - a. To filter by event category (errors, informational messages, or warnings), click the  status drop-down menu, and select the status condition or conditions. Event category filter options include  errors,  informational messages, and  warnings, or any combination of the three.
 - b. To filter by start date and time, enter the starting date and time using one of the following options:
 - In the **From** text box, type the date and time in format MM/DD/YYYY HH:MM AM/PM. For example, to search from the first day of January in 2016 at 8:00 AM, enter 1/1/2016 8:00 AM.
 - To select the current data and time, click the  **Calendar** widget in the **From** text box and then click the current date. The current time automatically appears.
 - Click the  **Calendar** widget, select the date, then click the  **Clock** widget and select the desired time using the controls. Click away from the calendar to accept the selected changes.
 - c. To filter by alert message description, enter the description in the **Search message** text box. For example, to see alerts only related to agents, enter "agent;" to see alerts related to transfers, enter "transfer."
 - d. To further refine the list of events that appears, you can also define an end date and time in the same format.
The list of events is immediately filtered based on the criteria you selected.

For more information, see the following related tasks:


- [About repetition reduction](#)
- [Viewing running tasks from any Core Console page](#)
- [Viewing tasks](#)
- [Navigating between tasks, alerts, and the events journal](#)
- [Suspending or resuming scheduled tasks](#)
- [Viewing alerts](#)
- [Viewing a journal of all logged events](#)

Navigating between tasks, alerts, and the events journal


The events recorded in the Core logs are visible from the Core Console, and fall into three categories: tasks, alerts, and a journal of all logged events.



Regardless of whether you are viewing events for the Core for a specified machine, the default view for events is the Tasks page. By clicking the drop-down menu to the right of the page title, you can select another view of events.

Perform the following steps to view events and navigate between tasks, important alerts, and a journal of all events.

You can view events that pertain to the Core by clicking  (Events) from the icon bar.

If you first navigate to a protected or replicated machine, and then select **Events** from the menu at the top of the page, you can view events for the specified machine.

Regardless of whether you are viewing events for the Core for a specified machine, the default view for events is the Tasks page. Click on the drop-down  menu to the right of the *Tasks* page title to select another view of events.

1. If you want to view all tasks for the Rapid Recovery Core, from the icon bar, click  (Events).
The default view displays all tasks for the Core. Proceed to [step 3](#).
2. If you want to view tasks for a specific protected machine, navigate to the Summary page of the specified machine, and then click the **Events** menu.
The default view displays all tasks for the selected machine. Proceed to [step 3](#).
3. From the top left of the Tasks pane, click  (the downward-facing arrow to the right of the Tasks title).
A drop-down menu appears.
4. Select one of the following:

Option	Description
Task	A task is a job that the Rapid Recovery Core must perform, such as transferring data in a regularly scheduled backup, or performing a restore from a recovery point.
Alert	An alert is a priority notification related to a task or event, such as an error, warning, or important informational message.\
Journal	The journal shows a complete list of all logged events. This list is more comprehensive than the set included in alerts.

The selected view of events appears. For example, if you selected **Alerts**, the Alerts page appears

5. To see a different view, return to the drop-down menu to the right of the Tasks, Alerts, or Journal pane, and select the option for the view you want.
The list of events is filtered to display only the relevant set of events for the current view.

For more information, see the following related tasks:

- [Understanding event notifications in Rapid Recovery](#)
- [Configuring notification groups](#)
- [Viewing a journal of all logged events](#)
- [Navigating between tasks, alerts, and the events journal](#)
- [Viewing tasks](#)
- [Suspending or resuming scheduled tasks](#)

- [Viewing alerts](#)
- [Viewing a journal of all logged events](#)
- [Viewing running tasks from any Core Console page](#)

Understanding event notifications in Rapid Recovery

The Rapid Recovery Core tracks many events, and logs the information for diagnostic and operational purposes. You can set up notification of specific events. Rapid Recovery lets you choose the method of notification, and the duration of time for which the system should retain a record of those events. With the repetition reduction feature, you can even adjust the frequency for notifying you about the same event.

Jobs and events tracked on the Core are saved by default for 30 days. To change the retention period for tracking events, see [Configuring event retention](#).

You can be notified of events through different methods. The notification methods supported are listed in the following table:

Table 106: Supported event notification methods

Option type	Description	Default setting
Email	Notifies specified user by email, using SMTP configuration settings in the Core, and based on the email notification template.	Off
Windows Event Log	Logs events using the Windows Event Log API. This log can be read using the Windows Event Viewer or custom applications.	On
syslogd	Logs events intended to be read on a Linux logging server that also supports the syslog message protocol.	Off
Toast alert	When this notification method is selected, messages appear briefly as a pop-up in the lower-right corner of the Rapid Recovery Core.	On
SNMP trap	If you configure the Rapid Recovery Core as an SNMP agent, and this notification method is selected, events are sent to a logging server using the trap number designated in the Notification Options dialog box.	On

Notification groups let you specify the types of events for which you want to be notified, and set the notification method.

Rapid Recovery Core requires at least one notification group, and ships with a default group which is automatically applied. You can use the default settings, or you can edit them.

Optionally, you can add and configure additional notification groups. For example, you can use the default group as is, and you can set up another group that uses email notifications.

As another example, you can set up a custom notification group for one type of event (for example, Microsoft Exchange), and send all related notifications to an Exchange administrator.

For more information, see [Configuring notification groups](#).

Email notifications are disabled by default. To send notifications by email, you must set up an email server, and edit or add a notification group with the **Notify by email option** enabled. This setting requires you to enter the email address to which the notifications are sent. For more information, see [Configuring an email server](#).

If using an Exchange server, you must set up relay on the server. Otherwise, despite successful email tests, no email notifications are sent. For more information, see your Exchange Server administrator.

The Core uses an email notification template. The template includes a subject line and specific content for the message body. A default email notification template is included. The template identifies the Core and the server host, the date and time of the event, the nature of the event, and error details if relevant. Optionally, you can modify the default template, or revert any customization to restore the default. For more information, see [Configuring an email notification template](#).

You can reduce the number of events of the same type and scope that are logged and visible from the Core Console by using the repetition reduction feature. This feature is enabled by default. You can disable this feature, or you can control the span of time for which events are combined into a single occurrence in the event log. For more information, see [About repetition reduction](#).

Configuring notification groups

i **NOTE:** You must first configure Simple Mail Transfer Protocol (SMTP) server settings if you want to send notifications as email messages, as described in this procedure. For more information on setting email server configuration settings, see [Configuring an email server](#).



Notification groups let you define sets of specific events for which users are alerted, and the manner in which that notification takes place. You can configure alerts to be sent by the following methods:


- By email
- In the Windows event log
- Using syslogd
- Using toast alerts
- Using alerts
- Using SNMP trap


Rapid Recovery Core ships with a default notification group for the Core. You can edit that group to suit your needs. Optionally, you can configure more than one notification group with different notification parameters.

Notification groups can be set at the Core level, or for each specific protected machine.





Complete the steps in this procedure to configure notification groups for alerts.

1. To set notifications at the Core level, from the icon bar, click **...** (More), and then select  **Notifications**. The Notifications page appears. Skip to [step 3](#).
2. To set notifications for a specific protected machine, do the following:
 - a. From the Protected Machines menu, click the machine for which you want to specify notifications. The Summary page appears.
 - b. In the Summary page of the protected machine, from the More drop-down menu, select  **Notifications**. The Custom Notification Groups page appears.
3. If you want to add a new notification group, click **+Add Group**. Skip to [Step 5](#). The Add Notification Group dialog box appears, showing a general description area and two tabs.


- If you want to edit the default notification group or an existing notification group, in the Notification Groups pane, click the  (More) drop-down menu for the appropriate notification group, and select **Edit**. The Edit Notification Group dialog box appears, showing a general description area and two tabs.
- In the general description area, enter the basic information for the notification group, as described in the following table.

Option	Description
Name	Enter a name for the event notification group. This information is required.  CAUTION: The value you enter for the notification group name cannot be changed later.
Description	Enter a description that clarifies the purpose for the event notification group. This information is optional.

- In the **Enable Alerts** tab, configure the set of system events that result in alerts. These appear on the Alerts page when you view events in the Core Console. You can select sets of events as described in the following table:

Option	Description
All Alerts	To create alerts for all events, select
Errors	To create alerts for errors, from the Select Types menu, click Error. This is represented by a red X. 
Warning	To create alerts for errors, from the Select Types menu, click Warning. This is represented by a yellow exclamation point icon. 
Info	To create alerts for informational messages, from the Select Types menu, click Info. This is represented by a blue i. 
Restore Default	For the default Core notification group, to restore the set of events to appear as alerts to the default, from the Select Types menu, click Reset to defaults .  NOTE: This option is available when editing the default Core notification group only. It is not available for new Core notification groups or for configuration notifications for a specific protected machine.

7. To create alerts for a specific event type (error, warning, or informational message), do the following:
 - a. If the **All Alerts** option does not display alert groups, click the right angle bracket > symbol preceding the All Alerts label. The symbol changes to a downward-facing arrow, and the view expands to show groups.
 - b. Then click the right angle bracket > symbol next to any specific alert group to display related events in the group.
 - To define alerts for all events in every group, select the checkbox for **All Alerts**.
 - To define alerts for all events within any alert group, select the checkbox next to that group.
 - To select only some alert types within an alert group, expand the group and then select only those specific events for which you want to log, report, and set alerts.
8. Click the **Notification Options** tab.
9. On the Notification Options tab, specify how to handle the notification process.

Option	Description
Notify by email	<p>Designate the recipients of the email notification. You can choose to specify separate multiple email addresses as well as blind and carbon copies.</p> <p> NOTE: If using Exchange Server, SMTP relay must be set up on the server. Otherwise event notifications will not be sent to the designated email address. For more information, consult your Exchange Server administrator.</p>
Notify by Windows Event Log	Select this option if you want notifications to be reported through the Windows Event Log.
Notify by syslogd	<p>Select this option if you want notifications to be reported through syslogd. Specify the details for the syslogd in the following text boxes:</p> <ul style="list-style-type: none"> • Host: • Port:
Notify by Toast alerts	Select this option if you want notifications to appear as pop-up messages in the lower-right corner of your screen.
Notify by SNMP Trap	The Rapid Recovery Core serves as an SNMP agent, sending traps (notifications about specific events) to an SNMP manager. The result is the reporting of Core information such as alerts, repository status, and protected machines. Select this option if you want to notify Core events by SNMP trap. You must also specify a trap number, which is used by the SNMP manager.

10. Click **OK**.

If creating a new group, you will see a message indicating that the notification group name you defined cannot be changed after creating the group. Other properties within the notification group can be changed at any time.

- If you are satisfied with the group name, confirm this message and save your work.
- If you want to change the group name, click **No** to return to the Create Notification Group window, update the group name and any other notification group settings, and save your work.

The notification group appears in the summary table. You can create different notification groups using any set of parameters.

Understanding email notifications

You can set up Rapid Recovery Core to notify you of specific events by sending an email message to an email address that you specify. The events which trigger alerts, and the notification methods, are defined in the notification group.

i | **NOTE:** Notification groups must be established regardless of whether you use email as a notification method. For more information, see [Configuring notification groups](#).

Rapid Recovery uses an email notification template, which determines the information sent in each notification. The template defines the email subject line for each alert, and the content in the email message body. The template has default settings; you can use the default as-is, or you can test and make modifications to serve your needs. At any point after customizing the notification template, you can choose the Restore Defaults option to return to using the default template. For information on viewing and customizing the email template, see [Configuring an email notification template](#).

If you choose email as one of your notification options, you must first configure an email SMTP server. The Rapid Recovery Core uses the server you define to send alerts based on the parameters in the notification group.

Additionally, to receive email notifications, you must enable the **Notify by email** option within the notification group. This notification option is disabled by default. The **Notify by email** setting requires a "To" address defined at minimum. (Optionally, you can add copy and blind copy addresses if desired.)

This section includes the following topics:

- [Configuring an email server](#)
- [Configuring an email notification template](#)


The following topics are also relevant:

- [Configuring notification groups](#)
- [Configuring an email server](#)
- [Configuring an email notification template](#)



Configuring an email server

Complete the steps in this procedure to configure an email server.


i | **NOTE:** You must also configure notification group settings, including enabling the **Notify by email** option, before email alert messages are sent by the system. For more information on specifying events to receive email alerts, see [Configuring notification groups](#).

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **SMTP Server**.
 - Scroll down on the right side of the *Settings* page until you can see the **SMTP Server** heading.
3. Click on the setting you want to change.
The setting you selected becomes editable.
4. Enter the configuration information as described in the following table.

Option	Description
SMTP server	Enter the name of the email server to be used by the email notification template. The naming convention includes the host name, domain, and suffix; for example, smtp.gmail.com.
From	Enter a return email address. It is used to specify the return email address for the email notification template; for example, noreply@localhost.com.
User name	Enter a user name for the email server.
Password	Enter the password associated with the user name required to access the email server.
Port	Enter a port number. It is used to identify the port for the email server; for example, the port 587 for Gmail. The default is 25.
Timeout (seconds)	Enter an integer value to specify how long to try to connect to the email server. It is used to establish the time in seconds before a timeout occurs. The default is 60 seconds.
TLS	Select this option if the mail server uses a secure connection such as Transport Layer Security (TLS) or Secure Sockets Layer (SSL).

5. For each setting, when satisfied with your changes, click  to save the change and exit edit mode, or click  to exit edit mode without saving.

 **CAUTION: If you do not confirm each change, your settings will not change.**

6. Click  **Send Test Email** and then do the following:
 - a. In the Send Test Email dialog box, enter a destination email address for the test message and then click **Send**.
 - b. If the test message fails, exit the error dialog box and the Send Test Email dialog box, and revise your email server configuration settings. Then send the test message again.
 - c. Once the test message is successful, click **OK** to confirm the successful operation.
 - d. Check the email account to which you sent the test email message.

Configuring an email notification template

When you enable notifications of Rapid Recovery events by email, a default template is created for you by default. The SMTP email server defined for the Core uses this template to send notifications regarding Rapid Recovery events by email.

This topic describes the process of configuring the default email notification template or customizing the content. Using the Restore Default option, you can restore changes to the default notification template at any time.

CAUTION: Modify the template at your own risk. You are responsible for testing any modifications to the template. Only the default template is supported.

Complete the steps in this procedure to configure an email notification template.

NOTE: You must also configure an email server and notification group settings, including enabling the **Notify by email** option, before email alert messages are sent. For more information about configuring an email server for sending alerts, see [Configuring an email server](#). For more information on specifying events to receive email alerts, see [Configuring notification groups](#).

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click **⋮** (More), and then select **🔔 Notifications**.
The *Notifications* page appears.
3. In the Email Settings pane, click **🔗 Change**.
The *Edit Email Notification Configuration* dialog box appears.
4. Select **Enable Email Notifications**.
The email template is enabled and is visible. The values of the default email template are described in the following step.

- Review the contents in the *Edit Email Notification Configuration* dialog box and determine if the default content suits your needs.

Option	Description
Enable email notifications	<p>This setting enables or disables the email notification template.</p> <ul style="list-style-type: none"> To enable email notification, select this option. To disable email notification, clear this option.
Email Subject	<p>The contents of this text field control the subject line for email messages sent as notifications of system events. The default email subject line is as follows:</p> <pre><hostName> <level>: <name> for <agentName></pre> <p>The contents of this text area control the body for email messages sent as notifications of system events. The default email body message is as follows:</p> <pre><shortCompanyName> <coreProductName> on <hostName> has reported the <level> event "<name>" Date/Time: <localTimestamp> <message> <if(details.errorDetails)> <details.ErrorDetails.Message> <details.ErrorDetails.Details> <endif> ---</pre> <p>About this event: <description></p> <pre><coreAdminUrl></pre>
Send Test Email	Clicking this button sends a test email message to the email address you specify in the resulting Send Test Email dialog box.
Restore Defaults	Clicking this button removes any customized changes from the email template, and restores the Email Subject and Email fields with the default contents described in this table.
OK	Clicking this button confirms and saves the settings in the Edit Email Notification Configuration dialog box.
Cancel	Clicking this button cancels any changes in the Edit Email Notification Configuration dialog box.

6. If you want to customize the email template, make changes to the text or variables described in the preceding step. The variables used in the default are described in the following table.

Option	Description
hostName	The host name of the Core
details	The details object of the specific event.
agentName	The name of the protected machine associated with this event, if the event has a scope of a single protected machine.
repositoryName	The name of the repository associated with this event, if the event has repository scope.
jobSummary	The summary of the job associated with this event, if the event has job scope.
remoteSlaveCoreName	The name of the remote target Core associated with this event, if the event has target Core scope.
remoteMasterCoreName	The name of the remote source Core associated with this event, if the event has source Core scope.
productName	The name of the product, for example 'AppAssure Core' or 'Rapid Recovery Core.' This product name can be changed for branding using white labeling.
companyName	The name of the company selling the product.

7. In the **Email Subject** text box, enter a subject for the email template.
The Email Subject is used to define the subject of the email notification template, for example, <hostname> - <level>: <name>.
8. In the **Email** text box, enter the information for the body of the template which describes the event, when it occurred, and the severity.
9. Click **Send Test Email**, and then do the following:
- In the *Send Test Email* dialog box, enter a destination email address for the test message and then click **Send**.
 - If the test message fails, exit the error dialog box and the *Send Test Email* dialog box, click **OK** to save the current email template settings. Then modify your email server settings as described in the procedure [Configuring an email server](#). Ensure that you reenter the password for that email account. Save those settings and then return to this procedure.
 - Once the test message is successful, click **OK** to confirm the successful operation.
 - Check the email account to which you sent the test email message.

Once you are satisfied with the results of your tests, return to the *Edit Email Notification Configuration* dialog box, and click **OK** to close the dialog box and save your settings.

Configuring event settings

You can configure certain setting specific to events.

For example, you can set repetition reduction settings to reduce the amount of notifications you see for identical repeated events.

You can also set the amount of time, in days, that event records are retained in the database.

View the following topics to learn about configuring event settings.

- [About repetition reduction](#)
- [Configuring event retention](#)

Related topics:

- [Viewing events using tasks, alerts, and journal pages](#)
- [Understanding event notifications in Rapid Recovery](#)

About repetition reduction

The ability for administrators to receive notification upon the occurrence of certain events is critical. Nonetheless, in certain circumstances, receiving repeated notification of events that you are aware of can also be frustrating or inconvenient. Even if a notification is generated due to an environmental failure that you wish to know about immediately, it is possible for the same error condition to generate hundreds or thousands of events in the event log. To reduce repetition in the event log, and reduce the inconvenience of receiving repeated Toast alerts or e-mail notifications for the same event in the Core Console, Rapid Recovery includes a repetition reduction setting, which is enabled by default and set at 5 minutes. This setting can be set as low as 1 minute and as high as 60 minutes. It can also be disabled entirely.




To set, change, or disable repetition reduction settings, see [Configuring repetition reduction](#).


When repetition reduction is disabled, then every time an event of the same type and scope occurs, it is logged in the database. Regardless of how much time passed since that event previously occurred, each new occurrence is shown in the Alerts portion of the *Events* page.

When repetition reduction is enabled (for example, with the default time of 5 minutes), then the first time that specific event occurs, it is logged in the event database and shown in the Alerts log. If subsequently an event of the same type and scope is again logged within the threshold of time established, then the count for the event in the database increases by 1 for each repeat occurrence within that threshold. The log shows in the Alerts portion of the *Events* page. However, it displays the event only once, with the date and time of the most recent occurrence. The event log is not updated with the same event until the threshold of time from the first occurrence expires. For example, if set for 5 minutes and the event occurs again 6 minutes later, it appears in the log and you receive another notification.

Configuring repetition reduction

Complete the steps in this procedure to configure repetition reduction for events.

1. Navigate to the Rapid Recovery Core Console. From the icon bar, click  (More), and then select  **Notifications**.
The Notifications page appears.
2. In the Repetition Reduction pane, view the existing settings.
3. To enable, disable, or change the stored events threshold time, click  **Change**.
The Edit Repetition Reduction dialog box appears
4. Do one of the following:
 - To disable repetition reduction, clear the **Enable Repetition Reduction** option.
 - To enable repetition reduction, select the **Enable Repetition Reduction** option.
 - To change the time threshold (in minutes) for which repeated identical events are ignored, in the **minute(s)** text box, enter a number between 1 and 60.


 **NOTE:** The **Enable Repetition Reduction** option must be selected in order to change this value.

5. Click **OK** to save your settings and close the dialog box.


Configuring event retention

Events and jobs tracked on the Core are saved for a specified amount of time. The default setting is 30 days. This number can be set between 0 days and 3652 days (approximately 10 years).

Complete the steps in this procedure to configure retention for events.

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the *Settings* page, click **Database Connection**.
 - Scroll down on the right side of the *Settings* page until you can see the **Database Connection** heading.

The Database Connection settings appear.

3. To change the amount of days for which event information is saved to the database, click on the **Retention Period (days)** text area, enter a value between 0 and 3652, and then click  to save the change. The events retention period is adjusted as specified.

Reporting

This section provides an overview of reporting available in the Rapid Recovery Core Console.

Topics include:

- [About Rapid Recovery reports](#)
- [Generating reports from the Core Console](#)
- [Managing scheduled reports from the Core Console](#)
- [Using the Reports menu](#)
- [Using the Reports toolbar](#)
- [Understanding the Job report](#)
- [Understanding the Job Summary report](#)
- [Understanding the Failure report](#)
- [Understanding the Summary report](#)
- [Understanding the Repository report](#)
- [Understanding the Classic Summary report](#)

About Rapid Recovery reports

You can generate these reports from the Rapid Recovery Core Console. Other reports are also expected to be available from the QorePortal.

The reports available from the Core Console are described in the following table.

Table 107: Rapid Recovery reports

Report type	Description
Job report	<p>Provides a basic report on the status of successful jobs, failed jobs, and jobs with errors. Failed jobs can be further viewed in a Failure report.</p> <ul style="list-style-type: none"> By default, the report time range is for the last 31 days. This can be customized. When run from the Core, this report can specify details for one or more Cores. By default, this set of information includes jobs for all machines — database servers, protected machines, replicated machines, and recovery point-only machines — in the specified Cores. In the report parameters, you can customize the report. Use the filters to select or exclude some machines. You can also select or exclude jobs that are machine independent, in which case the report shows status for Core jobs only. <p>For more information on this report type, see Understanding the Job report.</p>
Job	Provides a more detailed report on the status of successful jobs, failed jobs, and jobs with errors,

Report type	Description
summary report	<p>showing a separate line in the report for each job type, allowing better diagnosis of potential issues.</p> <ul style="list-style-type: none"> • By default, the report time range is for the last 31 days. This can be customized. • Unlike the Job report, this report does not offer a selection of Cores as a parameter. • By default, this set of information includes jobs for all machines — database servers, protected machines, replicated machines, and recovery point-only machines — and jobs that are machine dependent, by job type. You can customize the report. Use the filters to select or exclude some machines. You can also select or exclude jobs that are machine independent. • When run from the perspective of a protected machine from the Core Console, the resulting report displays the status for jobs only for that protected machine. <p>For more information on this report type, see Understanding the Job Summary report.</p>
Failure report	<p>Provides information on failed Core jobs for the specified criteria. This report can include protected machine details or exclude them. Like the Job report, this report can also be run only from a protected machine selected in the Core. The resulting report displays detail about failed jobs only for the selected protected machine.</p> <p>For more information on this report type, see Understanding the Failure report.</p>
Summary report	<p>Provides summary information. By default, this set of information includes jobs for all machines— every protected machine, replicated machine and recovery point-only machine in the specified Cores. In the report parameters, you can customize the report. Use the filters to select or exclude some machines. You can also select or exclude jobs that are machine independent, in which case the report shows status for Core jobs only.</p> <p>This report is not available from the perspective of any single protected machine. The categories of information in this report include Core, license, and repository. The information is displayed in list, chart, and table form</p> <p>For more information on this report type, see Understanding the Summary report.</p>
Repository report	<p>This report type provides you with a report of all repositories on the selected Core or Cores. You can also select any single repository available to the Core. This report is only from the perspective of the Core.</p> <p>For more information on this report type, see Understanding the Repository report.</p>
Classic summary report	<p>This report provides a summary view of job statistics success, repository summaries by GB, snapshots success, repository usage trends, and a summary of protected machines on your Core. Report parameters include date range and the relevant Core.</p> <p>For more information on this report type, see Understanding the Classic Summary report.</p>
Scheduled report	<p>You can also schedule any of these reports from the Core Console. Scheduling a report causes the report you specify to generate repeatedly on the schedule you define.</p> <p>Optionally, you can establish email notifications each time a report is generated. For more information about scheduling, modifying, pausing, or deleting reports, see Managing scheduled reports from the Core Console.</p>

Based on the report type and the parameters that you select, you can generate a report on one or more Rapid Recovery Cores or for one or more protected machines.

Generating reports from the Core Console

You can generate reports on demand from the Core Console. The following rules apply:

- All Rapid Recovery reports can be generated from the perspective of the Core.
- Additionally, two job types (the Job report and the Failure report) can be generated from the perspective of a protected machine. For such reports, data is generated only pertaining to the selected machine.
- Failure reports contain data only if jobs on the selected Cores (or protected machines) have failed.

The method for generating on-demand reports is similar, whether the report is generated from the focus of the Core, or whether it is generated from the perspective of a protected machine. However, navigation differs slightly.

- For more information about generating Core-level reports on demand, see [Generating a Core report on demand](#).
- For more information about machine-level reports, see [Generating a protected machine report on demand](#).
- You can also schedule reports to generate on a repeated basis at the Core level only. For more information about scheduling, modifying, pausing, or deleting reports, see [Managing scheduled reports from the Core Console](#).

Generating a Core report on demand

As described in the topic [About Rapid Recovery reports](#), you can generate the full range of available reports from the Core Console.

Complete the steps in the following procedure to generate a report from the perspective of the Rapid Recovery Core.

1. Navigate to the Rapid Recovery Core Console.
2. From the icon bar, click ******* (More), and then select **📊 Reports**.
The Job Report page appears. To the right of the report name in the page title, a downward-facing arrow appears, from which you could select another report type.
If you want to generate a Job report, proceed to [step 6](#) to begin specifying your report criteria.
3. To choose another report type, click the arrow to the right of the report name to see a menu of available reports.
4. For defining scheduled reports, see [Scheduling a report](#).
5. To generate a Repository report only, skip to [step 11](#).

6. For a Job, Job Summary, Failure, or Summary report, from the **Date Range** drop-down menu, select a date range.

If you do not choose a date range, the default option (Last 31 days) is used. You can choose from the options in the following table.

Option	Description
Last 24 hours	Reports activity for the last day, relative to the time you generate the report.
Last 7 days	Reports activity for the last week, relative to the time you generate the report.
Last 31 days	Reports activity for the last 31 days, relative to the time you generate the report.
Last 90 days	Reports activity for the last 90 days, relative to the time you generate the report.
Last 365 days	Reports activity for the last year, relative to the time you generate the report.
All Time	This time period spans the lifetime of the Core.
Custom	This time period requires you to further specify start and end dates.
Month to date	Reports activity from the first day of the current calendar month to the date you generate the report.
Year to date	Reports activity from the first day of the current calendar year to the date you generate the report.

i | **NOTE:** In all cases, no report data is available before the Core software was deployed, or from before machines were protected on the Core.


7. For a Job or Failure report, from the **Target Cores** drop-down menu, select the appropriate Core or Cores for which you want to generate a report.

The default selection includes all available Cores.

8. From the **Protected Machines** drop-down menu, select the machine or machines for which you want to generate the report.

By default, this set of information includes jobs for all machines—every protected machine, replicated machine and recovery point-only machine in the specified Cores. In the report parameters, you can customize the report. Use the filters to select or exclude some machines. You can also select or exclude jobs that are machine independent, in which case the report shows status for Core jobs only.


You can choose from:

Option	Description
Select all	<p>This option selects all protected machines protected on this Core.</p> <p> NOTE: You can select all machines, and then clear some of the selections to specify a subset of all machines.</p>
Machine independent	<p>Select this option to generate a report which includes jobs from a Core perspective. Job types such as creating or deleting a repository, or creating a boot CD, are not associated with a specific machine. If deploying the Agent software to a machine that is not yet protected, this job type is also considered machine independent. These jobs do not list a protected machine in the Protected Machine column of the resulting report.</p> <p>In contrast, if you deploy the Agent software to a machine that is already protected in the Core, the protected machine name is included in the report. It is not considered machine independent.</p>
Protected machines	<p>This option lists the machines protected on this Core. You can select them all, or a subset of the protected machines.</p>
Recovery points only	<p>This option lists machines that were once protected, but still have recovery points saved in the repository.</p>
[Source cores]	<p>If your Core is a target Core, and replicates recovery points for any machines protected on a source Core, then the name of that source Core appears (in all uppercase letters). This option lists all machines protected on that source Core. You can select all machines replicated in this target Core, or you can select a subset of them.</p>
[Custom groups]	<p>If you have any custom groups created on this Core, the name of each custom group appears as an option. Each object in that custom group appears. You can select all objects in the group, or a subset of them.</p>

9. If generating a Summary report, skip to [step 12](#).
10. For a Job, Job Summary, or Failure report, from the **Job Types** drop-down menu, select the appropriate job types.


By default, this set of information includes all jobs for the selected protected machines. In the report parameters, you can customize the report. Use the filters to select or exclude every job in the Main Jobs category, and every job in the Other Jobs category. Or you can expand each of these categories when defining job parameters, and select only the job types from either category that you want to appear in the report. Click the checkbox for any job type to select or clear that type. You can select some or all jobs from either category.

You can choose from the following **other** job types:


11. For a Repository report, from the **Repositories** menu, select the repository or repositories that you want included in the report.
The default selection includes all available repositories.
12. Click **Preview** to generate the report with the specified criteria.
If the report criteria you selected is not found, the report still generates, but the report contains an empty row.
For example, if there are no errors, the contents of the Error column are null in the report.
13. Do one of the following:
 - View the generated report online.
 - Update the report dynamically by changing any of the criteria; then click Preview again.
 - Use the Reports menu to select an export format (including the default format, PDF) and click  to export the report. For more information about the Reports menu, see [Using the Reports menu](#).
 - Use the Reports toolbar to view, manipulate, or print the report. For more information about the Reports toolbar, see [Using the Reports toolbar](#).

Generating a protected machine report on demand

You can generate a Job report or a Failure report for any protected machine on demand.

 **NOTE:** Scheduled reports at the machine level are not supported.

Complete the steps in the following procedure to generate a report for a protected machine.

1. Navigate to the Rapid Recovery Core Console.
2. From the Protected Machines menu, click the protected machine for which you want to see a report.
The Summary page for the selected protected machine appears.
3. At the top of the page, from the menu options next to the protected machine name, click the downward-facing arrow  next to **Reports**, and then select a report type.
 - If you want to generate a report on all jobs pertaining to this protected machine, including failed jobs, click **Job Report**, and begin specifying your report criteria.
 - If you want to generate a list of failed jobs only pertaining to this protected machine, click **Failure Report**, and begin specifying your report criteria.

4. For a Job or Failure report, from the **Date Range** drop-down menu, select a date range.
If you do not choose a date range, the default option (**Last 31 days**) is used. You can choose from the options in the following table.

Option	Description
Last 24 hours	Reports activity for the last day, relative to the time you generate the report.
Last 7 days	Reports activity for the last week, relative to the time you generate the report.
Last 31 days	Reports activity for the last 31 days, relative to the time you generate the report.
Last 90 days	Reports activity for the last 90 days, relative to the time you generate the report.
Last 365 days	Reports activity for the last year, relative to the time you generate the report.
All Time	This time period spans the lifetime of the Core.
Custom	This time period requires you to further specify start and end dates.
Month to date	Reports activity from the first day of the current calendar month to the date you generate the report.
Year to date	Reports activity from the first day of the current calendar year to the date you generate the report.

i **NOTE:** In all cases, no report data is available before the Core software was deployed, or from before machines were protected on the Core.

5. From the **Job Types** drop-down menu, select the appropriate job types.
By default, this set of information includes all jobs for the selected protected machines. In the report parameters, you can customize the report.
- Use the filters to select or exclude every job in the **Main Jobs** category, and every job in the **Other Jobs** category.
 - Or you can expand the **Main Jobs** and **Other Jobs** categories, respectively, and select only the desired job types for that category that you want to appear in the report. Click the check box for any job type to select or clear that type. You can select some or all jobs from either category.
6. Click **Preview** to generate the report with the specified criteria.
If the report criteria you selected is not found, the report still generates, but the report contains an empty row. For example, if there are no errors, the contents of the Error column are null in the report.
7. Do one of the following:
- View the generated report online.
 - Update the report dynamically by changing any of the criteria; then click **Preview** again.
 - Use the Reports menu to select an export format and export the report. For more information about the Reports menu, see [Using the Reports menu](#).
 - Use the Reports toolbar to view, manipulate, or print the report. For more information about the Reports toolbar, see [Using the Reports toolbar](#).

Managing scheduled reports from the Core Console

You can schedule any of the reports available from the Core Console. Scheduling a report causes it to be generated repeatedly in the future. The schedule defines whether to generate the report on a daily, weekly, or monthly basis. Optionally, Rapid Recovery lets you send an email notification to one or more recipients when each report is generated. The email specifies the report type, report format, and date range, and includes the report as an attachment.



NOTE: Before you can send reports by email, you must configure an SMTP server for the Core. For more information, see [Configuring an email server](#).

Whether or not you choose to send email notifications, you can save the generated reports locally, or on a network location accessible to the Core server.

You must specify email notification and delivery, or you must specify a location to save reports. You can also choose both options.

This section includes the following topics:

- [Scheduling a report](#)
- [Modifying a report schedule](#)
- [Pausing, resuming, or deleting a scheduled report](#)

Scheduling a report

You can schedule a report available from the Core Console. The report then generates on the schedule you defined until you pause or delete the report.

You must specify email notification and delivery, or you must specify a location to save reports. You can also choose both options.

Complete the steps in this procedure to schedule a report.

1. Navigate to the Rapid Recovery Core Console.
2. From the icon bar, click **...** (More), and then select **Reports**.
The Job Report page appears. A downward-facing arrow appears to the right of the current report name.
3. Click the arrow to the right of the report name, and from the drop-down menu, select **Scheduled Reports**.
The Scheduled Reports page appears.
4. To schedule a report to generate on a repeated basis, click **+Add**.
The Set Reporting Schedule Wizard appears.

5. On the Configuration page of the wizard, enter the details for the report you want to schedule, and then click **Next**. The configuration options are described in the following table.

Table 108: Scheduled report configuration options

Machine	Available Reports
Name	Type the display name you want to assign to this particular schedule. The default name is Schedule report 1. Limit your name to 64 or fewer characters. Do not use prohibited characters or prohibited phrases .
Report format	Select a report output format. If you do not select a value, the default format (pdf) is used.
Report type	Select the type of report you want to generate on a repeated basis.
Labels	Select the labels you want to appear on your scheduled report. At least one label is required. The Custom Groups feature allows you to group Core objects in one logical container, for which you define a label. Using the Labels parameter in the Set Reporting Schedule Wizard, you can select a custom group for which scheduled reports are run. If no custom labels exist, the available options in the Labels drop-down menu include Select All and Protected Machines. If custom groups appear, each group appears as an option. You can select or clear any of the options to include or exclude those objects in the scheduled report.
Protected machine	Select one or more protected machines to be included in the report. This option is not available for the Repository report.
Job Types	Select the job types you want to appear in the report. By default, this set of information includes jobs for all machines—every protected machine, replicated machine and recovery point-only machine in the specified Cores. In the report parameters, you can customize the report. Use the filters to select or exclude some machines. You can also select or exclude jobs that are machine independent, in which case the report shows status for Core jobs only. The Job Types parameter is not available for the Core Summary and Repository scheduled report types.

6. On the Destination page of the wizard, select a destination for the reports you want to schedule. You must choose one of the following, and may select both. When satisfied, click **Next**.
 - In the **Send to email addresses** text box, enter one or more valid email addresses to notify users by email message when a scheduled report is generated.


 **NOTE:** If you do not specify email notifications and delivery, then you must specify a storage location.
 - Select **Save as file** to save the generated report files to a location you specify, and in the **Location type** drop-down menu, select a local, network, or cloud storage option. Then, in the **Location** text box, specify additional location information as described in the following table.

Table 109: Location options for scheduled reports

Location type	Location type description	Location
Local	Select location type Local to save generated reports in a local path accessible to the Core.	Specify the path in the Location text box. Type a location accessible to the Core locally. For example, to store reports in the Reports folder on the D drive, enter <code>D:\Reports\</code> .
Network	Select location type Network to save generated reports in a path accessible to the Core on the network. Specify the path in the Location text box.	Specify the path in the Location text box. Type a location accessible to the Core from the network. Use format <code>\\servername\sharename</code> . For example, to store reports on the Data server in the shared folder called Reports, enter <code>\\Data\Reports\</code> . Specify network credentials in the User name and Password text boxes.
Cloud	Select location type Cloud to save generated reports in a Cloud storage account configured in the Core. The storage account must already be defined before performing this step. For information on setting up a Cloud storage account to work with the Core, see Cloud accounts .	From the Account text box, select the appropriate Cloud storage account to use to store generated reports. From the Container text box, specify an appropriate container in the storage account. From the Folder Name text box, specify a folder into which to store future generated reports.

7. When satisfied with your Destination options, click **Next**.

8. On the Schedule page of the wizard, from the **Send data** menu, select an option to determine how frequently to generate the report that you specified. You can generate reports daily, weekly, or monthly. Each option has its own parameters, as described in the following table.

Table 110: Frequency options for generating scheduled reports



How frequent	Frequency details	Frequency parameters
Daily	Generates and saves or sends the specified report once daily at the specified time. Default time for this action is 12:00 AM (based on the time on the Core server).	To change the default time that the report generates, in the time text text box, type a new value or use the controls to change the hour, minutes, and AM or PM.
Weekly	Generates and saves or sends the specified report once weekly at the specified time of the specified day. Default time for this action is 12:00 AM on Sunday (based on the time on the Core server).	To change the default day that the report generates, from the day of week menu, select a day of the week. To change the default time that the report generates, in the time text text box, type a new value or use the controls to change the hour, minutes, and AM or PM.
Monthly	Generates and saves or sends the specified report once monthly on the specified date and time of day. Default date for this action is the first of each month at 12:00 AM (based on the time on the Core server).	To change the default date that the report generates, from the day of month menu, select a date. To change the default time that the report generates, in the time text text box, type a new value or use the controls to change the hour, minutes, and AM or PM.


9. Optionally, on the Schedule page of the wizard, if you want to prevent the scheduled report from generating until you resume paused reports, select **Initially pause reporting**.
If you want this report to generate as scheduled, clear this option.
10. When satisfied with the schedule, click **Finish** to exit the wizard and save your work.
The new report schedule appears in the Summary Reports summary table.

Modifying a report schedule

Once a report is scheduled, you can modify any of its parameters or details. You can edit report configuration information (report name, output format, report type, included repositories. You can also change email notification options, and the destination to save the generated report. Finally, you can also change the schedule of the report.

Complete the steps in this procedure to modify parameters for a scheduled report.

1. Navigate to the Rapid Recovery Core Console.
2. From the icon bar, click  (More), and then select  Reports.
The *Job Report* page appears. A downward-facing arrow appears to the right of the current report name.

3. Click the arrow to the right of the report name, and from the drop-down menu, select **Scheduled Reports**. The *Scheduled Reports* page appears.
4. In the Scheduled Reports summary table, from the row of the report you want to modify, click  (More options) and then select **Edit**. The *Set Reporting Schedule Wizard* appears.
5. Navigate through the pages of this wizard, changing any parameters necessary. For information on any of the parameters in this wizard, see the topic [Scheduling a report](#).
6. On the *Schedule* page of the wizard, click **Finish** to close the wizard and save your changes. The wizard closes, and the report schedule is modified.

Pausing, resuming, or deleting a scheduled report



Once a report is scheduled, it generates on the schedule defined. If you want to temporarily stop the generation of a scheduled report, then you can pause the schedule.

If a scheduled report is paused, and you wish to resume the generation of the report, then you can resume the report as described in this procedure.

If you are currently generating a scheduled report, and no longer need to generate that report, you can delete it.

To determine if any scheduled report is paused, check the status column in the scheduled reports summary table. A green sphere indicates an active scheduled report; a yellow sphere indicates a paused schedule; and a red sphere indicates an error.

Complete the steps in this procedure to pause, resume, or delete a schedule for a report.

1. Navigate to the Rapid Recovery Core Console.
2. From the icon bar, click  (More), and then select  **Reports**. The *Job Report* page appears. A downward facing arrow appears to the right of the current report name.
3. Click the arrow to the right of the report name, and from the drop-down menu, select **Scheduled Reports**. The *Scheduled Reports* page appears.
4. In the Scheduled Reports summary table, view the status of all scheduled reports, using the colored indicators.
5. For each report you want to pause or resume, select the check box in the first column.
6. From the Scheduled Reports options above the summary table, do one of the following:
 - To pause the generation of the selected reports, click **Pause**.
 - To resume generation of scheduled reports that have been paused, click **Resume**.
 - To delete the selected schedules for existing scheduled reports, click **Delete**.
Deleting a scheduled report only prevents the generation of future reports. If previous scheduled reports have been saved, they are not removed.

Using the Reports menu




The Reports menu appears at the top of the page when viewing Reports. This menu includes a report title, which is also a drop-down menu that lets you see which report types are available. Below this menu are one or more filters

that help you to define your report criteria.

The specific filters available depend on the report type. For information on the parameters that apply to each report type, see the topic for understanding that report type.

On the right side of the reports menu, some controls appear. These controls, described in the following table, help you generate and export the report.

Table 111: Reports menu controls

UI Element	Description
 Preview button	Click the preview button to generate a report based on the selected report type and the report parameters specified in the filters.
 Export format drop-down menu	The Export format drop-down menu lets you select a report output format. If you do not select a value, the default format (pdf) is used.
 Download button	The Download button exports the generated report in the format type selected in the Export format menu

Reports include units of measure which make it easier to determine if a column is represented in GB, TB, or in seconds.

If you are not satisfied with the appearance of a generated or exported report, you can change the font used in the reports. For more information, see [Managing report settings](#).

Once a report is generated, you can use the reports toolbar, as described in the topic [Using the Reports toolbar](#).

For more information, see the following topics:



- [Managing report settings](#)
- [Understanding the Job report](#)
- [Understanding the Failure report](#)
- [Understanding the Summary report](#)
- [Understanding the Repository report](#)




















Using the Reports toolbar

After you generate it from the Reports menu, the report appears below a Reports toolbar. The toolbar can help you manipulate report output, including saving and printing the reports.

On the left of the toolbar, there is a Toggle sidebar option. This tool expands or contracts the sidebar, giving access to a few more display options. To the right of the toolbar, the Tools option expands a drop-down menu providing report navigation controls. The elements of the Reports toolbar are described in the following table.

Table 112: Reports toolbar icons

Icon	Description
	Toggle sidebar. All report pages are displayed as thumbnails. Other options in the sidebar are not supported.
	Sidebar: Show thumbnails. This is the default view for all pages of a

Icon	Description
	generated report.
	Sidebar: Show document outline. This feature is not supported.
	Sidebar: Show attachments. There are no attachments for reports. This feature is not supported.
	Find. Allows you to search text within the generated report. Includes options to highlight all text that matches the criteria you enter, and also to match or ignore case.
	Previous page. Move the report view to the previous page.
	Next page. Progress to the next page in the report view.
	Enter page number. Click in the page number text field, enter a valid page number, and press Enter to progress to that page in the report view.
	Zoom out. Lets you zoom out the view of the generated report. Each successive click zooms out further, to a minimum of 25%.
	Zoom in. Lets you zoom in the view of the generated report. Each successive click zooms in further, to a maximum of 1000%.
	Automatic Zoom. Lets you control the zoom view of the generated report, including viewing by actual size, fit page, full width, or by percentage, including 50%, 75%, 100%, 125%, 150%, 200%, 300%, or 400%.
	Open file. Lets you navigate your file system to locate and open a saved report.
	Print. Lets you print the generated report.
	Tools. The Tools drop-down menu expands or contracts when you click this icon. The Tools options are described below.
	Tools: Go to first page. Navigates you to the first page of the generated report.
	Tools: Go to last page. Navigates you to the last page of the generated report.
	Tools: Rotate clockwise. This option rotates the canvas of the generated report in a clockwise direction.
	Tools: Rotate counterclockwise. This option rotates the canvas of the generated report in a counterclockwise direction.
	Tools: Hand tool. When you select this tool, it Lets you move the report by clicking and dragging across the screen.
	Tools: Document properties. Provides information about the document properties of the generated report. Click Close to close this window.

For information about generating a report, see [Generating reports from the Core Console](#).

Understanding the Job report

The Job report is available for the Rapid Recovery and for machines protected on the Core. This report provides you with a method to view the status of jobs performed by a selected Core or a protected machine. Rows or columns of data that appear in the report with no data indicate that the tested parameter was null. For example, if a column (such as Errors) appears with no information, then no errors are occurred for the selected record. If the report generates a blank row, the job for the selected record reflects machine-independent activity.

For information on how to generate a Job report from the Core, see [Generating a Core report on demand](#). For information on how to generate a Job report for a protected machine, see [Generating a protected machine report on demand](#).

When you generate a Job report, report details include the following:

- Selection criteria for the report
- A summary table showing a row for each job in the date range you specified. In addition to listing the appropriate Core, protected machine, and job type, each row includes:
 - A summary of the job
 - The job status
 - Any errors related to the job
 - The start and end dates for the job
 - The job duration in seconds
 - The total work in MB

If information is not relevant for a specific category, that cell appears with no information in the report. For example, if the Core for a specified protected machine has no errors, the Error column is blank for that row in the report.

Understanding the Job Summary report

The Job Summary report is available when reporting from the Core perspective only; this report is not available from reports for a protected machine. This report has a single summary, showing summary information about all jobs performed on the Core, including a count of failed, passed, and canceled jobs. It shows more detail than the Job report, because it specifies each job as a separate line in the report.

For information on how to generate a Job Summary report, see [Generating reports from the Core Console](#).

Report parameters for this report type include:

- Date range
- Protected machines
- Job types

When you generate a Job Summary report, report details include selection criteria for the report, as well as information about protected machines, volumes, and job types.

Core information

The Core portion of the Summary Report includes data regarding the Rapid Recovery Core being reported. This information includes:

- The number of machines protected in the Rapid Recovery Core
- The number of machines with failed jobs

Protected machines summary

The Protected machines portion of the Summary report includes data for all machines protected by the selected Rapid Recovery Core, and the volumes on those machines.

The chart shows a line for each job type for each machine, and includes the ratio of successful jobs (of any type), number of jobs passed, number of jobs failed, and canceled jobs. (Canceled jobs are not considered for these statistics.)

Understanding the Failure report

The Failure report is a subset of the Job report and is available for the Rapid Recovery Core and for machines protected on the Core. A Failure report includes only the canceled and failed jobs listed in the Job report, and compiles them into a single report that can be printed and exported. If the report generates with a blank row, there are no errors within the date range specified in your report criteria.

NOTE: Results for target Cores and protected machines parameters appear for the Core-level report only.

For information on how to generate a Job report from the Core, see [Generating a Core report on demand](#). For information on how to generate a Job report for a protected machine, see [Generating a protected machine report on demand](#).

When you generate a Failure report, a summary table appears, showing a row for each job in the date range you specified. In addition to listing the appropriate Core, protected machine, and job type, each row includes:

- A summary of the job
- The job status
- Any errors related to the job
- The start and end dates for the job
- The job duration in seconds
- The total work in MB

Understanding the Summary report

The Summary report is available for one or more Cores. This report is not available from reports for a protected machine. The Summary report includes information about the repositories on the selected Rapid Recovery Core and about the machines protected by that Core. The information appears as two summaries within one report.

For information on how to generate a Summary report, see [Generating reports from the Core Console](#).

Report parameters for this report type include:

- Date range
- Protected machines

When you generate a Summary report, report details include selection criteria for the report, as well as information about repositories and protected machines.

Core information

The Core portion of the Summary Report includes data regarding the Rapid Recovery Core being reported. This information includes:

- The license key (identifier)
- The current version of the Rapid Recovery Core software

Repositories summary

The Repositories portion of the Summary Report includes data for the repositories located on the selected Rapid Recovery Core. This information includes:

- The number of repositories in the Rapid Recovery Core
- A summary of repositories on the Core.

Protected machines summary

The Protected machines portion of the Summary report includes data for all machines protected by the selected Rapid Recovery Core or Cores. This includes a chart and a summary table.

The chart shows protected machines by the ratio of successful jobs (of any type), compared to failed jobs. (Canceled jobs are not considered for these statistics.)

The X or horizontal axis shows the number of protected machines. The Y or vertical axis shows tiers of success. Specifically, the Y axis shows, by protected machine, how many had:

- No jobs performed
- Less than 50% success rate
- 50% or more success rate
- 100% success rate

Below the chart, information appears about protected machines. This information includes:

- The amount of protected machines
- The number of protected machines with failed jobs

- A summary table, by protected machine, which shows:
 - Protected machine name
 - Volumes protected by the machine
 - Protected space, in GB (total and current)
 - Daily change rate (average and median)
 - Job statistics (success, completed, failed, canceled)
 - If encryption was applied
- The Core version

Understanding the Repository report

The Repository report includes information about the repositories on the selected Rapid Recovery Core and about the machines protected by that Core. The information appears as two summaries within one report.

For information on how to generate a Repository report from the Core, see [Generating a Core report on demand](#).

Report parameters for this report type include only repositories.

When you generate a Repository report, report details for each repository includes a summary list of repositories on the Core.

Understanding the Classic Summary report

The Classic Summary report provides a summary over the selected period for a variety of useful metrics for your Core.

The top portion includes three pie charts:

- The first shows job statistics (successful, canceled, and failed jobs).
- The second pie chart shows space consumed per repository in GB.
- The third pie chart displays snapshot statistics (successful, canceled, and failed jobs).

Following the pie charts in the report is a trend chart showing repository usage.

Finally, the report includes a summary of protected machines.

For information about how to generate a Classic Summary report, see [Generating reports from the Core Console](#).

i | **NOTE:** In release 6.2x, this report was known as the Core Nostalgia Report.

Report parameters for this report type include:

- Date range
- The Core on which you want to report

VM export

This section describes how to export a recovery point to create a virtual machine.


Topics include:

- [Exporting to virtual machines using Rapid Recovery](#)
- [Exporting data to an ESXi virtual machine](#)
- [Exporting data to a VMware Workstation virtual machine](#)
- [Exporting data to a Hyper-V virtual machine](#)
- [Exporting data to a VirtualBox virtual machine](#)
- [Exporting data to an Azure virtual machine](#)
- [Managing exports](#)

Exporting to virtual machines using Rapid Recovery

From the Rapid Recovery Core, you can export a recovery point of a Windows or Linux machine from a repository to a virtual machine (VM). If the original machine protected on the Core fails, you can boot up the virtual machine to quickly replace it temporarily, allowing you time to recover the original protected machine without substantial downtime. This virtual export process results in a VM with all of the backup information from a recovery point, as well as the operating system and settings for the protected machine. The VM becomes a bootable clone of the protected machine.

i **NOTE:** The recovery point used must be part of a complete recovery point chain. For more information about recovery point chains, see the topic [Recovery point chains and orphans](#).

You can perform a virtual export from the Virtual Standby page in the Core Console, or by selecting **VM Export** from the  **Restore** drop-down menu on the button bar.

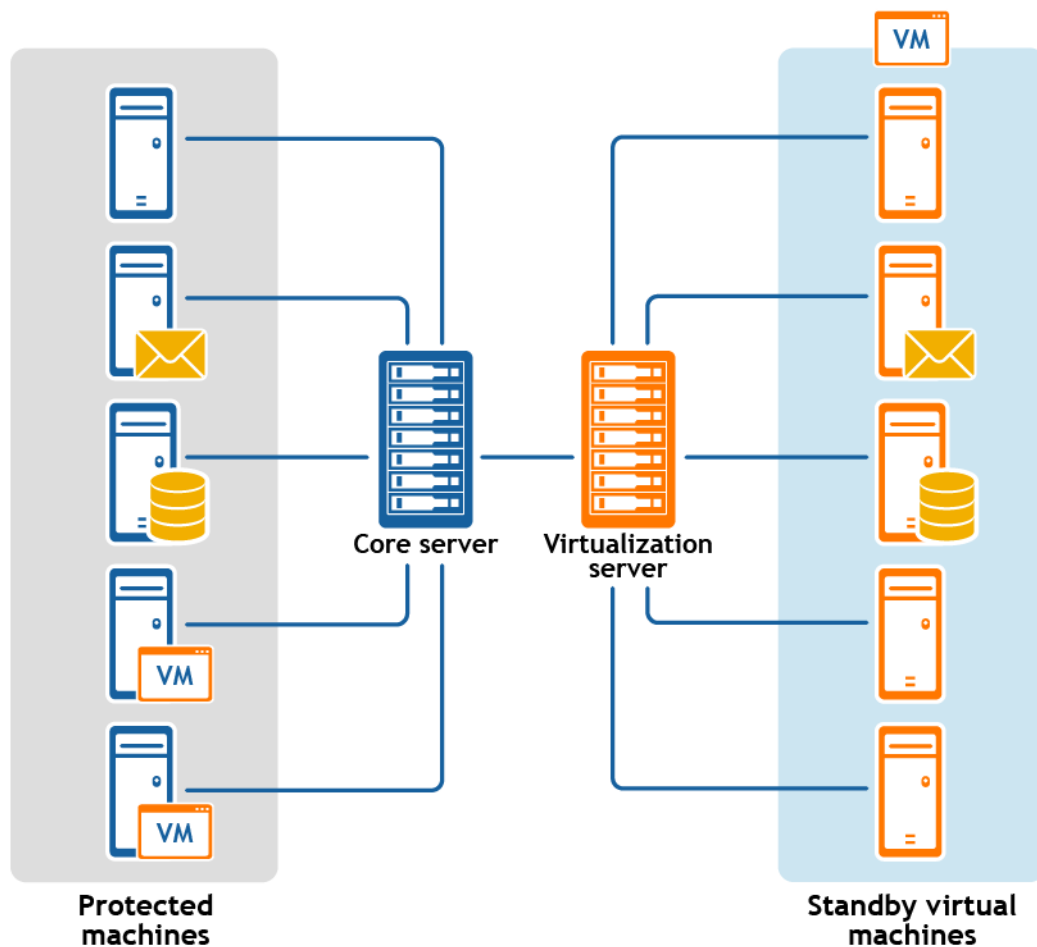
When you perform a virtual export from Rapid Recovery Core, you have two choices:

- You can perform a **one-time virtual export**, which creates a bootable VM representing a single snapshot in time from the information in the selected recovery point. The export job is queued immediately, and when it completes, the cloned VM exports to the location you specified. The configuration information used for a one-time export is not saved.
- You can set up **continual export**. This process creates a bootable VM for the protected machine you specify, saving the VM in a location you designate. The configuration information for performing that virtual export is saved in the Virtual Standby page in the Core Console. Subsequently, each time a new snapshot of the protected machine is captured, the Core queues a new virtual export job, and the bootable VM is refreshed with the updated information. Because this creates a high-availability resource for data recovery, this feature is also called **virtual standby**.

In between the time a virtual export job queues and is completed, the job is listed on the Export Queue pane of the Virtual Standby page in the Core Console.

The following diagram shows a typical deployment for exporting data to a virtual machine.

Virtual standby deployment



NOTE: In a continual export configuration involving replication set up between two Cores (source and target), you can export from either Core. However, you can only perform virtual export from the target Core after the initial replication is complete.

Subsequently, each time a new snapshot of the protected machine is captured, replication from a source Core queues a new virtual export job after each snapshot is captured. Replication from a target Core queues a new virtual export job after replication job.

Compatible VM hypervisors include vCenter/ESXi, VMware Workstation, Hyper-V, Oracle VM VirtualBox, and Azure. For information about supported versions of these hypervisors, see the topic "Hypervisor requirements" in the *Rapid Recovery System Requirements Guide*.

For ESXi, VMware Workstation, or Hyper-V, the virtual machine version must be a licensed version of these virtual machines, not the trial or free versions. Exporting to Azure requires you to have an account on Azure, with other prerequisites.

i **NOTE:** Working with Azure involves aspects unique to that cloud service provider. Like all other Azure features in Rapid Recovery, virtual export now uses the Azure Resource Management (ARM) deployment method. Azure setup steps and prerequisites to performing virtual export from the Rapid Recovery Core Console have changed accordingly. For details on Azure prerequisites prior to export, see [Before virtual export to Azure](#).

Related Topics

- [Exporting data to an ESXi virtual machine](#)
- [Exporting data to a VMware Workstation virtual machine](#)
- [Exporting data to a Hyper-V virtual machine](#)
- [Exporting data to a VirtualBox virtual machine](#)
- [Exporting data to an Azure virtual machine](#)
- [Managing exports](#)

Exporting data to an ESXi virtual machine

In Rapid Recovery, you can export data to ESXi by performing a one-time export, or by establishing a continual export (for virtual standby). Complete the steps in the following procedures for the appropriate type of export.



Topics include:

[Performing a one-time ESXi export](#)

[Setting up continual export to ESXi](#)


Performing a one-time ESXi export

Complete the steps in this procedure to perform a one-time export to ESXi.

1. From the Rapid Recovery Core Console button bar, click the  **Restore** drop-down menu, and then select  **VM Export**.
The Virtual Machine Export Wizard appears.
2. In the wizard, select **One-time Export**, and then click **Next**.
3. On the *Machines* page, select the protected machine that you want to export, and then click **Next**.
4. On the *Recovery Points* page, select the recovery point that you want to use for the export, and then click **Next**.

5. On the *Destination* page, in the **Export a virtual machine to** drop-down menu, select **vCenter/ESXi**, and then enter the parameters for accessing the virtual machine as described in the following table.

Table 113: Virtual machine parameters

Options	Description
Host name	Enter a name for the host machine.
Port	Enter the port for the host machine. The default is 443.
User name	Enter the user name for logging on to the host machine.  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault .
Password	Enter the password for logging on to the host machine.

6. Click **Next**.

7. On the *Virtual Machine Options* page, enter the information described in the following table, and then Click **Next**.

Table 114: Virtual machine options

Option	Description
Resource pool	Select a resource pool from the drop-down list.
VM configuration location	Select a data store from the drop-down list.
VM name	Enter a name for the virtual machine you want to export. The VM name that automatically appears by default is the name of the machine from which the recovery point originated.
Amount of RAM	Specify the memory usage for the virtual machine by clicking one of the following: <ul style="list-style-type: none">• Use the same amount of RAM as source machine• Use a specific amount of RAM, and then specify the amount in MB The minimum amount is 1024 MB and the maximum allowed by the application is 1,035,264MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.
Number of processors	Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.
Cores per processor	Enter the number of cores to use for each processor. The minimum is 1.
Disk provisioning	Select the type of disk provisioning from the following options: <ul style="list-style-type: none">• Thin. Thin provisioning creates a virtual disk the size of the used space on the original volumes, rather than the entire volume size. For example, if the original volume is 1 TB, but contains only 2 GB of used space, Rapid Recovery creates a virtual disk of 2 GB.• Thick. Thick provisioning creates a new disk or volume that is the same size as the original volume from the protected server, even if only a portion of the original volume is being used. For example, if the volume is 1 TB large but contains 2 GB of used space, Rapid Recovery creates a virtual disk of 1 TB.
Disk mapping	Specify the type of disk mapping as appropriate. You can choose from: <ul style="list-style-type: none">• Automatic. Using this option, the VM is exported to any available datastore with sufficient space.

Option	Description
	<ul style="list-style-type: none"> • Manual. Select this option to manually specify the datastore onto which to export the VM. • With VM. Select this option to export all virtual disks to the same datastore as the VM configuration.
Version	Select the version of the virtual machine.
Network adapters	<p>Optionally, specify network adapter options for the VM to be exported.</p> <p>i NOTE: This option appears when exporting to a hypervisor of the same type (in this case, when exporting to ESXi).</p> <p>You can choose from:</p> <ul style="list-style-type: none"> • Same as source machine. If you select this option, the VM include the same amount of network adapters as on the source, and they are assigned to the default network on the host. • Specific. Select this option to see the <i>Network Adapters</i> page of the Virtual Machine Export wizard. Here you can add one or more network adapters to the VM and assign a specific network to each adapter.
Secure boot	<p>Optionally, choose the secure boot option for the exported VM.</p> <p>i NOTE: This option appears only when secure boot is configured for the source protected machine.</p>
Restore all configuration data	<p>If you want to recover all VM configurations for volumes being recovered, select this option. If you want to restore data only and not the VM configurations, clear this option.</p> <p>This option appears only when performing virtual export of an agentless VM and after a resource pool is selected.</p> <p>For more information on the VM configuration backup and restore feature of Rapid Recovery, see VMware VM configuration backup and restore.</p>

- Optionally, on the *Network Adapters* page of the wizard, if you want to add a virtual network adapter to the VM to be exported, click **+ Add**, and then from the **Available Networks** drop-down menu, select the name of a network available on the hypervisor host.
- Optionally, to add additional network adapters, click **+ Add**, and then from the **Available Networks** drop-down menu, select the name of another network available on the hypervisor host.
- Click **Next**.

11. On the *Volumes* page, select the volumes from the source recovery point that you want to export to the VM, and then click **Finish** to complete the wizard and start the export.

i **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

12. On the *Summary* page, click **Finish** to complete the wizard and start the export.

i **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

Setting up continual export to ESXi

Complete the steps in this procedure to set up continual export to an ESXi virtual machine (VM) using Rapid Recovery. This process is also known as setting up virtual standby.




1. In the Rapid Recovery Core Console, do one of the following:
 - From the Core Console, in the button bar, click the  **Restore** drop-down menu, and then select  **VM Export**.
 - a. In the Virtual Machine Export Wizard, select **Continual export (virtual standby)**.
 - b. Click **Next**.
 - From the Core Console, in the icon bar, click  (Virtual Standby).
 - On the *Virtual Standby* page, click **Add** to launch the Virtual Machine Export Wizard.
2. On the *Machines* page of the Virtual Machine Export Wizard, select the protected machine that you want to export, and then click **Next**.
3. On the *Destination* page of the Export Wizard, in the **Export a machine to** drop-down menu, select **vCenter/ESXi**.
4. Enter the information for accessing the virtual machine as described in the following table, and then click **Next**.

Table 115: ESXi credentials

Option	Description
Host name	Enter a name for the host machine.
Port	Enter the port for the host machine. The default is 443.
User name	Enter the credentials for the host machine. <div>i NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault.</div>
Password	Enter the credentials for the host machine.

5. On the *Virtual Machine Options* page, enter the information described in the following table.

Table 116: Virtual machine options

Option	Description
Resource pool	Select a resource pool from the drop-down list.
VM configuration location	Select a data store from the drop-down list.
VM name	Enter a name for the virtual machine you want to export. The VM name that automatically appears by default is the name of the machine from which the recovery point originated.
Amount of RAM	Specify the memory usage for the virtual machine by clicking one of the following: <ul style="list-style-type: none">• Use the same amount of RAM as source machine• Use a specific amount of RAM, and then specify the amount in MB The minimum amount is 1024 MB and the maximum allowed by the application is 1,035,264MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.
Number of processors	Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.
Cores per processor	Enter the number of cores to use for each processor. The minimum is 1.
Disk provisioning	Select the type of disk provisioning from the following options: <ul style="list-style-type: none">• Thin. Thin provisioning creates a virtual disk the size of the used space on the original volumes, rather than the entire volume size. For example, if the original volume is 1 TB, but contains only 2 GB of used space, Rapid Recovery creates a virtual disk of 2 GB.• Thick. Thick provisioning creates a new disk or volume that is the same size as the original volume from the protected server, even if only a portion of the original volume is being used. For example, if the volume is 1 TB large but contains 2 GB of used space, Rapid Recovery creates a virtual disk of 1 TB.
Disk mapping	Specify the type of disk mapping as appropriate. You can choose from: <ul style="list-style-type: none">• Automatic. Using this option, the VM is exported to any available datastore with sufficient space.• Manual. Select this option to manually specify the datastore onto which to export the VM.• With VM. Select this option to export all virtual disks to the same datastore as the VM configuration.
Version	Select the version of the virtual machine.
Network adapters	Optionally, specify network adapter options for the VM to be exported.

Option	Description
	<p>i NOTE: This option appears when exporting to a hypervisor of the same type (in this case, when exporting to ESXi).</p> <p>You can choose from:</p> <ul style="list-style-type: none"> • Same as source machine. If you select this option, the VM include the same amount of network adapters as on the source, and they are assigned to the default network on the host. • Specific. Select this option to see the <i>Network Adapters</i> page of the Virtual Machine Export wizard. Here you can add one or more network adapters to the VM and assign a specific network to each adapter.
Secure boot	<p>Optionally, choose the secure boot option for the exported VM.</p> <p>i NOTE: This option appears only when secure boot is configured for the source protected machine.</p>
Restore all configuration data	<p>If you want to recover all VM configurations for volumes being recovered, select this option. If you want to restore data only and not the VM configurations, clear this option.</p> <p>This option appears only when performing virtual export of an agentless VM and after a resource pool is selected.</p> <p>For more information on the VM configuration backup and restore feature of Rapid Recovery, see VMware VM configuration backup and restore.</p>
Perform initial one-time export	<p>Select this option to queue the export job immediately. Clear this option if you want the Core to wait until the next forced or scheduled backup snapshot.</p>

6. Click **Next**.
7. Optionally, on the *Network Adapters* page of the wizard, if you want to add a virtual network adapter to the VM to be exported, click **+ Add**, and then from the **Available Networks** drop-down menu, select the name of a network available on the hypervisor host.
8. Optionally, to add additional network adapters, click **+ Add**, and then from the **Available Networks** drop-down menu, select the name of another network available on the hypervisor host.
9. Click **Next**.
10. On the *Volumes* page, select the volumes to export (for example, **C:** and **D:**), and then click **Finish** to close the wizard and start the export.

i **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

Exporting data to a VMware Workstation virtual machine

In Rapid Recovery, you can export data to VMware Workstation by performing a one-time export or by establishing a continual export (for virtual standby). Complete the steps in the following procedures for the appropriate type of

export.

Topics include:

[Performing a one-time VMware Workstation export](#)

[Setting up continual export to VMware Workstation](#)

Performing a one-time VMware Workstation export

Complete the steps in this procedure to perform a one-time export to VMware Workstation.




1. In the Rapid Recovery Core Console, in the button bar, click the  **Restore** drop-down menu, and then click  **VM Export**.
The *Virtual Machine Export Wizard* appears.
2. In the wizard, from the Select VM Export Type page, select **One-time Export** and then click **Next**.
3. On the *Machines* page, select the protected machine that you want to export, and then click **Next**.
4. On the *Recovery Points* page, scroll through the list of recovery points if necessary, and select the recovery point that you want to use for the export. Then click **Next**.
5. On the Destination page, in the **Export a virtual machine to** drop-down menu, select **VMware Workstation**.
6. To export the VM to a local drive, in the **VM location** field, specify the path of the local folder in which to create the exported virtual machine. For example, enter `E:\VirtualExports\`. Then proceed to step 9.
7. To export the VM to a network shared directory, enter the required information as described in the following table.


Table 117: Network shared location parameters

Option	Description
VM location	Specify the path of the network shared directory in which to create the exported virtual machine. For example, enter <code>\\ServerName\MySharedDirectory\VirtualExports\</code> .
User name	Enter the user name for an account that is registered on the target machine. The account must have read and write permissions to the network share.  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault .
Password	Enter the password for the specified user account.


8. Click **Next**.

9. On the Virtual Machine Options page, enter the settings for the new virtual machine, as described in the following table.

Table 118: Virtual machine parameters




Option	Description
VM name	Enter a name for the virtual machine being created; for example, VM-0A1B2C3D4.  NOTE: The default name is the name of the source machine.
Version	Specify the version of VMware Workstation for the virtual machine.
Amount of RAM	Specify the amount of RAM, in megabytes, for the exported virtual machine to use. The minimum amount is 1024 MB and the maximum allowed by the application is 65536 MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.
Number of processors	Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.
Cores per processor	Enter the number of cores to use for each processor. The minimum is 1.

10. On the Volumes page, select the volumes to export (for example, C:\ and D:\), and then click **Finish** to close the wizard and start the export.

 **NOTE:** You can monitor the status and progress of the export by viewing the Virtual Standby or Events pages.


Setting up continual export to VMware Workstation

Complete the steps in this procedure to set up continual export to a VMware Workstation virtual machine (VM) using Rapid Recovery. This is also known as setting up virtual standby.

1. In the Rapid Recovery Core Console, do one of the following:
 - From the Core Console, in the button bar, click the  **Restore** drop-down menu, and then select  **VM Export**.
 - a. In the Virtual Machine Export Wizard, select **Continual export (virtual standby)**.
 - b. Click **Next**.
 - From the Core Console, in the icon bar, click  (Virtual Standby).
 - On the *Virtual Standby* page, click **Add** to launch the Virtual Machine Export Wizard.
2. On the *Machines* page of the Virtual Machine Export Wizard, select the protected machine that you want to export, and then click **Next**.
3. On the *Destination* page, in the **Export a virtual machine to** drop-down menu, select **VMware Workstation**.


4. To set up continual export to a local drive, in the **VM location** field, specify the path of the local folder in which to create the exported virtual machine. For example, enter `E:\VirtualExports\`. Then proceed to step 9.
5. To set up continual export to a network shared directory, enter the required information as described in the following table.

Table 119: Network shared location parameters

Option	Description
VM location	Specify the path of the network shared directory in which to create the exported virtual machine. For example, enter <code>\\ServerName\MySharedDirectory\VirtualExports\</code> .
User name	Enter the user name for an account that is registered on the target machine. The account must have read and write permissions to the network share. <div>  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault. </div>
Password	Enter the password for the specified user account.

6. Click **Next**.
7. On the *Virtual Machine Options* page, enter the settings for the new virtual machine, as described in the following table.

Table 120: Virtual machine parameters

Option	Description
VM name	Enter a name for the virtual machine being created; for example, <code>VM-0A1B2C3D4</code> . <div>  NOTE: The default name is the name of the source machine. </div>
Version	Specify the version of VMware Workstation for the virtual machine.
Amount of RAM	Specify the amount of RAM, in megabytes, for the exported virtual machine to use. The minimum amount is 1024 MB and the maximum allowed by the application is 65536 MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.
Number of processors	Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.
Cores per processor	Enter the number of cores to use for each processor. The minimum is 1.
Perform initial one-time export	Optionally, select to perform the virtual export immediately instead of after the next scheduled snapshot.

8. Click **Next**.

9. On the *Volumes* page, select the volumes to export (for example, *C:* and *D:*), and then click **Finish** to close the wizard. If you selected **Perform initial one-time export**, the export job is immediately queued.

i | **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

Exporting data to a Hyper-V virtual machine

In Rapid Recovery, you can export data to Hyper-V Export by performing a one-time export, or by establishing a continual export (for Virtual Standby).

Rapid Recovery supports first-generation Hyper-V export to the following hosts:

- Windows 8
- Windows 8.1
- Windows 10 x64
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2016

Rapid Recovery supports second-generation Hyper-V export to the following hosts:

- Windows 8.1
- Windows 10 x64
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2016

i | **NOTE:** Not all protected machines can be exported to Hyper-V second generation hosts.

Only protected machines with the following Unified Extensible Firmware Interface (UEFI) operating systems support virtual export to Hyper-V second generation hosts:

- Windows 8 (UEFI)
- Windows 8.1 (UEFI)
- Windows 10 (UEFI)
- Windows Server 2012 (UEFI)
- Windows Server 2012 R2 (UEFI)
- Windows Server 2016 (UEFI)

i | **NOTE:** Hyper-V export to second-generation VM can fail if the Hyper-V host does not have enough RAM allocated to perform the export. In such cases, you will see an error message on the VM

Complete the steps in the following procedures for the appropriate type of export.

Topics include:

[Performing a one-time Hyper-V export](#)

[Setting up continual export to Hyper-V](#)

Performing a one-time Hyper-V export

Complete the steps in this procedure to perform a one-time export to Hyper-V.




1. In the Rapid Recovery Core Console, in the button bar, click the  **Restore** drop-down menu, and then click  **VM Export**.
The Virtual Machine Export Wizard appears.
2. In the wizard, from the *Select VM Export Type* page, select **One-time Export** and then click **Next**.
3. On the *Machines* page, select the protected machine that you want to export, and then click **Next**.
4. On the *Recovery Points* page, scroll through the list of recovery points if necessary, and select the recovery point that you want to use for the export. Then click **Next**.
5. On the *Destination* page, in the **Export a virtual machine to** drop-down menu, select **Hyper-V**.
6. To export to a local machine with the Hyper-V role assigned, click **Use local machine**. Then proceed to step 8.
7. To export a virtual machine to a remote Hyper-V server, click **Remote host**, and then enter the information for the remote host as described in the following table.




Table 121: Remote host information

Text Box	Description
Hyper-V host name	Enter an IP address or host name for the Hyper-V server.
Port	Enter a port number for the machine. It represents the port through which the Core communicates with this machine.
User name	Enter the user name for the user with administrative privileges for the workstation with the Hyper-V server. It is used to specify the logon credentials for the virtual machine.  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault .
Password	Enter the password for the user account with administrative privileges on the workstation with Hyper-V server. It is used to specify the logon credentials for the virtual machine.

8. Click **Next**.

9. On the *Virtual Machine Options* page, enter the settings for the new virtual machine, as described in the following table.

Table 122: Virtual machine parameters

Option	Description
VM location	<p>Specify the path of the local folder in which to create the exported virtual machine. For example, enter <code>E:\VirtualExports</code>.</p> <p>You can specify a directory name that does not yet exist; this process will create the directory.</p> <div> CAUTION: Do not attempt to manually create a directory.</div> <div> NOTE: Export to shared folders (for example, to <code>\\data\share</code>) is not permitted.</div>
VM name	<p>Enter a name for the virtual machine you want to export; for example, <code>VM-0A1B2C3D4</code>. The VM name that automatically appears by default is the name of the machine from which the recovery point originated. The name entered in this text box appears after export in the list of virtual machines in the Hyper-V Manager console.</p>
Amount of RAM	<p>Specify the amount of RAM for the exported virtual machine to use by selecting one of the following:</p> <ul style="list-style-type: none">• Same as source machine. Select this option to specify that the amount of RAM allocated to the exported VM is identical to the RAM allocated to the source machine.• Specific. Select this option, and then enter the amount of RAM in megabytes for the exported virtual machine to use. <p>The minimum amount is 1024 MB and the maximum allowed by the application is 65536 MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.</p>
Number of processors	<p>Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.</p>
Cores per processor	<p>Enter the number of cores to use for each processor. The minimum is 1.</p>
Generation	<p>Specify the Hyper-V generation type.</p>
Disk format	<p>You can select from:</p> <ul style="list-style-type: none">• VHDX• VHD <div> NOTE: Hyper-V Export supports VHDX disk formats if the target machine is running Windows 8 (Windows Server 2012) or later. If the VHDX is not supported for your environment, the option is disabled.</div>

Option	Description
Network adapters	<p>Optionally, specify network adapter options for the VM to be exported.</p> <p>i NOTE: This option appears when exporting to a hypervisor of the same type (in this case, when exporting to Hyper-V).</p> <p>You can choose from:</p> <ul style="list-style-type: none"> • Same as source machine. If you select this option, the VM include the same amount of network adapters as on the source, and they are assigned to the default network on the host. • Specific. Select this option to see the <i>Network Adapters</i> page of the Virtual Machine Export wizard. Here you can add one or more network adapters to the VM and assign a specific network to each adapter.
Secure boot	<p>Optionally, choose the secure boot option for the exported VM.</p> <p>i NOTE: This option appears only when secure boot is configured for the source protected machine.</p>
Perform initial one-time export	<p>Select this option to queue the export job immediately. Clear this option if you want the Core to wait until the next forced or scheduled backup snapshot.</p>

- Optionally, on the *Network Adapters* page of the wizard, if you want to add a virtual network adapter to the VM to be exported, then from the **Available Networks** drop-down menu, select the name of a network available on the hypervisor host.
- Optionally, to add additional network adapters, click **+ Add**, and then from the **Available Networks** drop-down menu, select the name of a network available on the hypervisor host.
- Click **Next**.
- On the *Volumes* page, select the volume(s) to export; for example, `C:\`.

i **NOTE:** If the selected volumes are larger than the appropriate maximum allocations supported by the application as indicated below, or exceed the amount of space available, you will receive an error.

 - For VHDX disk format, your selected volumes should be no larger than 64 TB.
 - For VHD disk format, your selected volumes should be no larger than 2040 GB.
- On the *Volumes* page, click **Finish** to complete the wizard and to start the export.

i **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

Setting up continual export to Hyper-V

Complete the steps in this procedure to set up continual export to a Hyper-V virtual machine (VM) using Rapid Recovery. This process is also known as setting up virtual standby.





1. In the Rapid Recovery Core Console, do one of the following:
 - From the Core Console, in the button bar, click the  **Restore** drop-down menu, and then select  **VM Export**.
 - a. In the Virtual Machine Export Wizard, select **Continual export (virtual standby)**.
 - b. Click **Next**.
 - From the Core Console, in the icon bar, click  (Virtual Standby).
 - On the *Virtual Standby* page, click **Add** to launch the Virtual Machine Export Wizard.
2. On the *Machines* page of the Virtual Machine Export Wizard, select the protected machine that you want to export, and then click **Next**.
3. On the *Destination* page, in the **Export a virtual machine to** drop-down menu, select **Hyper-V**.
4. To export to a local machine with the Hyper-V role assigned, click **Use local machine**. Proceed to step 8.
5. To export a virtual machine to a remote Hyper-V server, click **Remote host**, and then enter the information for the remote host as described in the following table.




Table 123: Remote host information

Text Box	Description
Hyper-V host name	Enter an IP address or host name for the Hyper-V server.
Port	Enter a port number for the machine. It represents the port through which the Core communicates with this machine.
User name	<p>Enter the user name of a user with administrative privileges on the protected machine you are exporting. This is also the user name you must use when logging into the exported VM.</p> <p> NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault.</p>
Password	Enter the password associated with the user name above.

6. Click **Next**.

7. On the *Virtual Machine Options* page, enter the settings for the new virtual machine, as described in the following table.

Table 124: Virtual machine parameters

Option	Description
VM location	<p>Specify the path of the local folder in which to create the exported virtual machine. For example, enter <code>E:\VirtualExports</code>.</p> <p>You can specify a directory name that does not yet exist; this process will create the directory.</p> <div> CAUTION: Do not attempt to manually create a directory.</div> <div> NOTE: Export to shared folders (for example, to <code>\\data\share</code>) is not permitted.</div>
VM name	<p>Enter a name for the virtual machine you want to export; for example, <code>VM-0A1B2C3D4</code>. The VM name that automatically appears by default is the name of the machine from which the recovery point originated. The name entered in this text box appears after export in the list of virtual machines in the Hyper-V Manager console.</p>
Amount of RAM	<p>Specify the amount of RAM for the exported virtual machine to use by selecting one of the following:</p> <ul style="list-style-type: none">• Same as source machine. Select this option to specify that the amount of RAM allocated to the exported VM is identical to the RAM allocated to the source machine.• Specific. Select this option, and then enter the amount of RAM (in MB) that you want to specify for the exported VM. <p>The minimum amount is 1024 MB and the maximum allowed by the application is 65536 MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.</p>
Number of processors	<p>Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.</p>
Cores per processor	<p>Enter the number of cores to use for each processor. The minimum is 1.</p>
Generation	<p>Specify the Hyper-V generation type.</p>
Disk format	<p>You can select from:</p> <ul style="list-style-type: none">• VHDX• VHD <div> NOTE: Hyper-V Export supports VHDX disk formats if the target machine is running Windows 8 (Windows Server 2012) or later. If the VHDX is not supported for your environment, the option is disabled.</div>

Option	Description
Network adapters	<p>Optionally, specify network adapter options for the VM to be exported.</p> <p>i NOTE: This option appears when exporting to a hypervisor of the same type (in this case, when exporting to Hyper-V).</p> <p>You can choose from:</p> <ul style="list-style-type: none"> • Same as source machine. If you select this option, the VM include the same amount of network adapters as on the source, and they are assigned to the default network on the host. • Specific. Select this option to see the <i>Network Adapters</i> page of the Virtual Machine Export wizard. Here you can add one or more network adapters to the VM and assign a specific network to each adapter.
Secure boot	<p>Optionally, choose the secure boot option for the exported VM.</p> <p>i NOTE: This option appears only when secure boot is configured for the source protected machine.</p>
Perform initial one-time export	<p>Select this option to queue the export job immediately. Clear this option if you want the Core to wait until the next forced or scheduled backup snapshot.</p>

8. Click **Next**.
9. Optionally, on the *Network Adapters* page of the wizard, if you want to add a virtual network adapter to the VM to be exported, then from the **Available Networks** drop-down menu, select the name of a network available on the hypervisor host.
10. Optionally, to add additional network adapters, click **+ Add**, and then from the **Available Networks** drop-down menu, select the name of a network available on the hypervisor host.
11. Click **Next**.
12. On the *Volumes* page, select the volume(s) to export; for example, `C:\`.

i **NOTE:** If the selected volumes are larger than the appropriate maximum allocations supported by the application as indicated below, or exceed the amount of space available, you will receive an error.

 - For VHDX disk format, your selected volumes should be no larger than 64 TB.
 - For VHD disk format, your selected volumes should be no larger than 2040 GB.
13. Select **Perform initial one-time export** to perform the virtual export immediately instead of after the next scheduled snapshot.
14. On the *Volumes* page, select the volumes to export (for example, `C:\` and `D:\`).

i **NOTE:** If the selected volumes are larger than the appropriate maximum allocations supported by the application as indicated below, or exceed the amount of space available, you will receive an error.

 - For VHDX disk format, your selected volumes should be no larger than 64 TB.
 - For VHD disk format, your selected volumes should be no larger than 2040 GB.

15. On the *Volumes* page, click **Finish** to complete the wizard and to start the export. On the *Virtual Standby* page, the continual export requirements you just confirmed appear in the Virtual Standby pane.

i | **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

Exporting data to a VirtualBox virtual machine

In Rapid Recovery, you can export data to Oracle VM VirtualBox by performing a one-time export, or by establishing a continual export (for virtual standby).

Complete the steps in the following procedures for the appropriate type of export.

i | **NOTE:** To perform this type of export, you should have VirtualBox installed on the Core machine.

Topics include:

[Performing a one-time VirtualBox export](#)



[Setting up continual export to VirtualBox](#)

Performing a one-time VirtualBox export

To export to Oracle VM VirtualBox on a remote Windows host, you should have VirtualBox installed on the Core machine.

i | **NOTE:** For exporting to VirtualBox on a remote Linux host, VirtualBox is not required on the Core machine.

Complete the steps in this procedure to perform a one-time export to VirtualBox.

1. In the Rapid Recovery Core Console, click the  **Restore** drop-down menu on the button bar, and then select  **VM Export**. The Virtual Machine Export Wizard appears.
2. In the wizard, select **One-time Export**, and then click **Next**.
3. On the *Machines* page, select the protected machine that you want to export, and then click **Next**.
4. On the *Recovery Points* page, select the recovery point that you want to use for the export, and then click **Next**.
5. On the *Destination* page, in the **Export a virtual machine to** drop-down menu, select **VirtualBox**.

6. To export a Windows VM to a local directory or to a network shared directory, do the following:
- Select **Use Windows machine**.
 - To export locally, in the **Target path** text box, type a local path; for example, type `E:\VirtualExports`.
 - To export the VM to a network shared directory, enter the required information as described in the following table.


Table 125: Network shared location parameters

Option	Description
Target path	Specify the path of the network shared directory in which to create the exported virtual machine. For example, enter <code>\\ServerName\MySharedDirectory\VirtualExports\</code> .
User name	<div>Enter the user name for an account that is registered on the target machine. The account must have read and write permissions to the network share.</div> <div><div>i</div><div>NOTE: Credentials are required to address the case in which multiple user accounts exist on the virtual machine. When the user is authenticated, only this user will see this Virtual Machine in VirtualBox manager. If an account is not specified, then the Virtual Machine will be registered for all existing users on the Windows machine with VirtualBox.</div></div> <div><div>i</div><div>NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault.</div></div>
Password	Enter the password for the specified user account.


- Click **Next**.

7. On the Virtual Machine Options page, enter the settings for the new virtual machine, as described in the following table.

Table 126: Virtual machine parameters


Option	Description
VM name	Enter a name for the virtual machine being created; for example, VM-0A1B2C3D4.  NOTE: The default name is the name of the source machine.
Amount of RAM	Specify the amount of RAM for the exported virtual machine to use by selecting one of the following: <ul style="list-style-type: none">• Same as source machine. Select this option to specify that the amount of RAM allocated to the exported VM is identical to the RAM allocated to the source machine.• Specific. Select this option, and then enter the amount of RAM in megabytes for the exported virtual machine to use. The minimum amount is 1024 MB and the maximum allowed by the application is 65536 MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.
Number of processors	Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.
Cores per processor	Enter the number of cores to use for each processor. The minimum is 1.

8. Click **Next**.
9. On the Volumes page, select the volumes to export (for example, `C:\` and `D:\`), and then click **Finish** to close the wizard and start the export.

 **NOTE:** You can monitor the status and progress of the export by viewing the Virtual Standby or Events pages.

Setting up continual export to VirtualBox

To export to Oracle VM VirtualBox on a remote Windows host, you should have VirtualBox installed on the Core machine.

 **NOTE:** For exporting to VirtualBox on a remote Linux host, VirtualBox is not required on the Core machine.

Complete the steps in this procedure to perform a continuous export to a VirtualBox virtual machine (VM) using Rapid Recovery. This process is also known as setting up virtual standby.






1. In the Rapid Recovery Core Console, do one of the following:
 - From the Core Console, in the button bar, click the  **Restore** drop-down menu, and then select  **VM Export**.
 - a. In the Virtual Machine Export Wizard, select **Continual export (virtual standby)**.
 - b. Click **Next**.
 - From the Core Console, in the icon bar, click  (Virtual Standby).
 - On the Virtual Standby page, click **Add** to launch the Virtual Machine Export Wizard.
2. On the Machines page of the Virtual Machine Export Wizard, select the protected machine that you want to export, and then click **Next**.
3. On the Destination page of the wizard, in the **Export a virtual machine to** drop-down menu, select **VirtualBox**.
4. On the Destination page of the wizard, to set up continual virtual export of a Windows VM to a local directory or to a network shared directory, do the following:
 - Select **Use Windows machine**.
 - To export locally, in the **Target path** text box, type a local path; for example, type `E:\VirtualExports`.
 - To export the VM to a network shared directory, enter the required information as described in the following table.



Table 127: Network shared location parameters

Option	Description
Target path	Specify the path of the network shared directory in which to create the exported virtual machine. For example, enter <code>\\ServerName\MySharedDirectory\VirtualExports\</code> .
User name	<p>Enter the user name for an account that is registered on the target machine. The account must have read and write permissions to the network share.</p> <div>  NOTE: Credentials are required to address the case in which multiple user accounts exist on the virtual machine. When the user is authenticated, only this user will see this Virtual Machine in VirtualBox manager. If an account is not specified, then the Virtual Machine will be registered for all existing users on the Windows machine with VirtualBox. </div> <div>  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault. </div>
Password	Enter the password for the specified user account.

- To set up continual virtual export to a remote Linux machine, do the following:
- Click **Next**.
- Proceed to step 6.

5. On the Destination page of the wizard, to set up continual virtual export to a remote Linux machine, select **Remote Linux machine**, and then enter information about the virtual machine as described in the following table.


Table 128: Remote Linux machine settings

Option	Description
VirtualBox host name	Enter an IP address or host name for the VirtualBox server. This field represents the IP address or host name of the remote VirtualBox server.
Port	Enter a port number for the machine. This number represents the port through which the Core communicates with this machine. For Linux machines, the default is port 22.
Target path	Specify a target path to create the virtual machine. For example, enter <code>/home/username/VirtualExport/VMname</code> .  NOTE: It is recommended that you create a root folder from root so that the virtual machine runs from root. If you do not use root, you will need to create the destination folder manually on the target machine prior to setting up the export. You will also need to manually attach or load the virtual machine after the export.
User name	User name of the account on the target machine, for example, root.  NOTE: Optionally, you can select credentials for a user in the Credentials Vault, or you can save credentials you add here to the vault. For more information, see Credentials Vault .
Password	Password for the user account on the target machine.


6. Click **Next**.

- On the Virtual Machine Options page, enter the settings for the new virtual machine, as described in the following table.

Table 129: Virtual machine parameters

Option	Description
VM name	Enter a name for the virtual machine being created; for example, VM-0A1B2C3D4.  NOTE: The default name is the name of the source machine.
Amount of RAM	Specify the amount of RAM for the exported virtual machine to use by selecting one of the following: <ul style="list-style-type: none">Same as source machine. Select this option to specify that the amount of RAM allocated to the exported VM is identical to the RAM allocated to the source machine.Specific. Select this option, and then enter the amount of RAM in megabytes for the exported virtual machine to use. The minimum amount is 1024 MB and the maximum allowed by the application is 65536 MB. The maximum amount of memory usage is limited by the amount of RAM available to the host machine.
Number of processors	Enter the number of virtual CPUs you want for the exported virtual machine. The minimum is 1.
Cores per processor	Enter the number of cores to use for each processor. The minimum is 1.
Perform initial one-time export	Select this option to queue the export job immediately. Clear this option if you want the Core to wait until the next forced or scheduled backup snapshot.

- Click **Next**.
- On the Volumes page, select the volumes to export (for example, `C:\` and `D:\`), and then click **Finish** to close the wizard and start the export.

 **NOTE:** You can monitor the status and progress of the export by viewing the Virtual Standby or Events pages.

Exporting data to an Azure virtual machine

In Rapid Recovery, you can export data to Azure by performing a one-time export, or by establishing a continual export (for virtual standby). One-time export to the Azure platform includes deployment in the workflow. If using a continual export, you can also later deploy the exported files to a bootable VM.

Azure supports virtual export of machines with 64-bit operating systems, including those that use BIOS firmware and those that use UEFI, such as Generation 2 virtual machines. For specific information about operating systems that support export to Azure, see the operating system installation and compatibility matrix in the appropriate edition of the *Rapid Recovery System Requirements Guide*.

- i** | **NOTE:** While Rapid Recovery supports the export of Generation 2 virtual machines to Azure, there are limitations to the size and types of disks that Azure supports. For more information, see <https://docs.microsoft.com/en-us/azure/virtual-machines/generation-2#generation-2-vm-sizes>.
- i** | **NOTE:** If using virtual export to Azure in Rapid Recovery 6.2.x or earlier, you will see a compatibility warning when you upgrade. Since Microsoft no longer supports its classic Azure Service Management model, users upgrading to Rapid Recovery as of this release must re-create virtual standby to Azure using the updated Rapid Recovery Core Console. When you create a new continual export to Azure, the virtual standby definition uses the Azure Resource Manager model.

Working with Microsoft Azure

Microsoft Azure is a subscription-based cloud computing platform. The following information is provided to Rapid Recovery customers to facilitate using Azure with our product.

Azure login requires JavaScript. You may need to enable JavaScript or otherwise adjust security settings in the browser accordingly. For more information, consult your systems administrator.

- [Azure interface disclaimer](#)
- [Country codes used on the Azure website](#)
- [Before virtual export to Azure](#)
- [Creating an Azure storage account](#)
- [Creating a container in an Azure storage account](#)
- [Creating an Azure Active Directory web application](#)
- [Obtaining Azure subscription information](#)
- [Obtaining the directory ID for your Azure web application](#)
- [Obtaining a secret key for your Azure web application](#)
- [Microsoft Azure documentation](#)

Azure interface disclaimer

! | **CAUTION:** The Microsoft Azure interface is subject to change.

The information provided in this document relating to steps required in Azure were current as of the date of publication. This information is provided as a service to our customers to assist them with Azure prerequisites.

However, when working with Azure, be aware that specific steps, URLs or even the Azure interface may change at any time, which is beyond our control.

If you are having difficulty performing any steps related to your Azure subscription, please seek the advice of a Microsoft Azure representative.

Country codes used on the Azure website

The Azure website uses language and country codes for its web addresses, which affect display of the content in the appropriate language. The typical URL construction uses the format: `https://[Microsoft or Azure domain]/[country-code]/[destination]/`, in which the country code controls the language display and the remainder of the URL specifies the content.


For example, when viewing the documentation center for US English, the URL is <https://docs.microsoft.com/en-us/azure/>. If viewing the same page for Spanish (Spain), the correct URL is <https://docs.microsoft.com/es-es/azure/>.

The URLs for Azure used throughout this document include the country code for English in the United States. For other languages, URLs may differ based on the settings on your computer, and the languages and country codes Microsoft supports.

If you are browsing in a language other than US English, or if your machine settings are configured for a different language, the language and country code portion of the various URLs cited in this guide may differ accordingly.

Before virtual export to Azure

Before you can perform virtual export to Azure, you need to meet all of the following prerequisites.

- You must have a protected machine with at least one recovery point in a Rapid Recovery Core that you want to export as a VM to Azure.
 - Remote access must be enabled on the protected machine for the deployed VM to boot successfully.
 - You must have administrative access to an account on Azure.
 - You must have an Azure Active Directory (AD) web application. Each web application must have the following characteristics:
 - **A secret key.** Immediately save the secret key to a secure location, as it cannot be recovered or viewed later.
 - **A valid URL.** The URL that you assign to your web application must contain valid URL syntax. However, from a Rapid Recovery perspective, this URL is not required to be active.
 - **Appropriate permissions.** An Azure administrator for your subscription must grant the web application the appropriate Identity and Access Management (IAM) permissions (specifically, the **Owner** role).
-  **NOTE:** For details on creating and setting up your AD web application, see [Creating an Azure Active Directory web application](#).
- Your Azure subscription must be associated with a resource group created using Azure Resource Manager. If using an older resource group created using Azure Service Manager, create a new resource group, since Microsoft no longer supports ASM objects.
 - Within Azure, you must create (or have access to) a virtual network that is associated with your Azure subscription, location, and resource group. Since this controller is required when deploying a VM to Azure, it is a prerequisite for performing one-time virtual export, or for deploying a continual export (virtual standby) to an active VM in Azure.
 - You must have access to specific information for your Azure subscription, as described in [Required Azure subscription information](#) later in this topic.

About Azure storage containers

Before performing virtual export to Azure, you must have an Azure storage account. For information about creating a storage account, see [Creating an Azure storage account](#).

When you complete the one-time VM export, the necessary files are exported to the specified storage account in format `<storageAccountName>/<exportContainer>/<export_folder>`. The default name for the export container is "export" and the default name for the export folder is the name of the protected VM.

Since one-time virtual export automatically includes deployment of the exported VM, the export files are then copied into a deployment folder (of the same name) within a deployment container (named "deploy" by default). The VM is then deployed from this second location.

If you select **Show advanced options** in the *Storage* page of the export wizard, you can select different existing container names, or you can enter new names for the export container and folder name and the deployment container.

Required Azure subscription information

When adding an Azure cloud account to your Rapid Recovery Core, or when performing virtual export to Azure, you must be able to identify specific information related to your Azure account. You are required to have details about your account or subscription (including account name and ID), region or location, your Azure web application and its appropriate role (Owner) and properties (directory or tenant ID, secret key, and virtual network). You must identify your resource group, and you must know the name of your storage accounts and their containers and sub-folders. Some of this information is described using more than one term.

The following matrix can help guide you to the Azure information you may be asked to identify, and how to find it.

Rapid Recovery text box label	Corresponding Azure item label	Description
Cloud account name	Subscription name (or Subscription)	Using the Cloud account name drop-down menu in the Virtual Machine Export Wizard is optional. The first time you perform virtual export to Azure, no information is available from this menu. After you successfully enter all credentials for an Azure subscription into the Rapid Recovery Core Console, the information is cached. Subsequently, instead of entering all credential information, you can select the appropriate subscription from the Cloud account name drop-down menu.
Region	Region	Each Azure portal is associated with a geographic region. Choose the region your portal is accessed from. Options include: <ul style="list-style-type: none">• Azure Global Cloud• Azure China Cloud• Azure German Cloud• Azure US Government Cloud
Application ID	Application ID	Each Azure AD web application created is assigned an application ID. For information about creating an Azure AD web application and its associated secret key, see Creating an Azure Active Directory web application .
Secret key	Keys	Each web application must have one or more secret keys that you can use to authenticate using Azure APIs.

Rapid Recovery text box label	Corresponding Azure item label	Description
		<p>! CAUTION: When creating any secret key, immediately record the description and key value in a secure location for the long term. If you do not retain the secret key for your Azure AD web application when you create it, it cannot be recovered.</p> <p>For information about creating an Azure AD web application and its associated secret key, see Creating an Azure Active Directory web application.</p>
Tenant ID	Directory ID	This is the directory ID for the AD web application that connects the Core to your Azure subscription.
Subscription ID	Subscription ID	This ID is associated with your Azure subscription and your unique subscription name.
Storage account name	Resource group	The resource group is a container for resources that share a common life cycle. Using resource groups, you can deploy, manage, and monitor all the services for your solution as a group.
Export container	Storage container	<p>An export container is a child object of the Azure storage container (resource group). When deploying a VM to Azure, this information is stored in a storage container within its parent resource group.</p> <p>Optionally, before exporting, you can create an appropriate storage container from within Azure into which the exported data is stored.</p> <p>You can also create the export storage container dynamically. In Rapid Recovery 6.8, the default name provided is "export."</p> <p>Storage requirements grow as your protected machine protects more data. In your Azure account, the container you specify must be associated with a storage location with sufficient space to accommodate the VM.</p> <p>For more information, see the topic Creating a container in an Azure storage account.</p>
Export folder name	None	By default, this folder is named after the VM you want to export.
Deployment container	None	This is the container name into which the VM is deployed. The default name provided is "deploy."
Resource group		Specify the resource group to be used.
Virtual network		Specify the virtual network to be associated with the selected resource group.

Rapid Recovery text box label	Corresponding Azure item label	Description
		You must create this in Azure before deploying a VM to Azure. Thus, it is required before performing a one-time virtual export, or before deploying an existing continual export to a VM.

For more information

- For more information on working with Azure, see [Working with Microsoft Azure](#).
- For a detailed description of aspects of virtual export that are unique to Azure, see [Exporting and deploying VMs for Azure](#).
- For more information on Azure configurations and pricing, see the [virtual machines pricing](#) page on the Azure website.
- For links to other useful Azure-related references on Microsoft websites, see [Microsoft Azure documentation](#).

Related topics:

For procedures related to exporting or deploying VMs on Azure, see the following topics:

- [Before virtual export to Azure](#)
- [Performing a one-time Azure export](#)
- [Setting up continual export to Azure](#)
- [Deploying a virtual machine in Azure](#)

About Azure storage accounts

Users are advised to research features of Azure before using them with Rapid Recovery. Proper research enables you to balance your needs, preferences, and costs.

Take the example of an Azure storage account, which contains data objects: Binary Large Objects (or “blobs”), files, queues, tables, and disks. When creating an Azure storage account, consider the following aspects:

- What type of objects are you storing?
- Is the retention length expected to be brief or long?
- Do you intend to access data frequently or rarely?
- How quickly do you need access to the information stored there (immediately, minutes, hours)?

When performing virtual export to Azure, you must select or create a storage account that supports the relevant type of blobs. For export, Rapid Recovery uses page blobs, which are a collection of 512-byte pages optimized for random read and write operations. Azure now supports a maximum page blob size of 8TB. Eventually, this Azure restriction is likely to be increased in the future. Accordingly, as of release 6.4, Rapid Recovery Core has doubled the supported maximum data disk size for virtual export and deploy to Azure from 4TB to 8TB.

Azure storage offers different access tiers, which affect cost, restrict data types, affect speed of access, and apply to the frequency of use. Select the Azure storage account kind that best reflects your needs.

For Azure storage, there are 3 account kinds, which is relevant when determining which blobs you want to store and how quickly or often you need to access them. These are shown in the following table:

Kind	Description	Recommendations
Storage (general purpose v1), or GPV1	Legacy storage type. Supports page blobs, required for virtual export. Does not have an access tier.	Can be used for virtual export or archiving to Azure. Microsoft documentation suggests using GPV2 instead when possible.
Storage V2 (general purpose v2), or GPV2	Contemporary and default storage account type. Supports page blobs, required for virtual export. Lets you select an access tier when creating the storage account. Incorporates all of the functionality of GPV1 and BlobStorage accounts.	Recommended by Microsoft. Hot access tier has higher storage costs, but the lowest access costs. Use GPV2 with hot access tier for continual export to Azure; consider GPV2 with cool access tier for one-time export.
BlobStorage	Legacy account kind. Supports only block and append blobs, not page blobs, and thus does not support virtual export. Lets you select an access tier when creating the storage account.	Can be used for archiving but not for virtual export. Microsoft documentation suggests using GPV2 instead when possible.


NOTE: You can upgrade a GPV1 or BlobStorage account to a GPV2 account with no downtime and without the need to copy data. For more information, see [Azure article Upgrade to a General-purpose v2 storage account](#).

Creating an Azure storage account



You must have administrative access to an account on Azure.

NOTE: The Microsoft Azure interface is subject to change. For more information, see [Azure interface disclaimer](#).


To perform certain functions (including virtual export from Rapid Recovery Core to an Azure account), you must first create a storage account on Azure.

NOTE: When creating a storage account, many of the options have additional information available by clicking the  information icon. Place your cursor over the icon for more information.

Complete the steps in this procedure to create an Azure storage account.


1. Open the Microsoft Azure dashboard.
2. From the left navigation area, click  **All resources**.
3. From the All resources pane, click **+ Add**.
4. In the **New** blade, click  **Storage Account**.

5. On the *Create storage account* page, enter the parameters for your Azure storage account as described in the following table:

Option	Description
Subscription	<i>Required.</i> Select the subscription into which you want this storage account to be created.
Resource group	<i>Required.</i> Select the resource group in which you want this storage account to be created.  NOTE: For successful virtual export, this owning resource group must be created using Azure Resource Manager.
Storage account name	<i>Required.</i> Provide a clear unique description for your storage account, using between 3 and 24 characters consisting only of lower-case letters and numbers. For example, companyabcstorage1 .
Location	<i>Required.</i> Choose the appropriate Azure region.
Performance	Select Standard or Premium .
Account kind	Select an account kind. The default, StorageV2 (general purpose v2), is recommended for most scenarios.
Replication	Choose a replication strategy for your storage account. The default is Read-access geo-redundant storage (RA-GRS).
Access tier (default)	Choose from Cool or Hot , based on frequency of access.

6. Optionally, to change the requirement for secure transfer or to associate with a virtual network, click **Next: Advanced**. Otherwise, skip to step 10.
7. On the *Advanced* page, in the Security area, to disable or enable secure transfer, click the appropriate selection. Secure transfer is enabled by default.

8. Optionally, on the *Advanced* page, in the Virtual Networks area, to select an existing virtual network associated with your subscription, location, and resource group, from the **Virtual network** drop-down menu, select a network and skip to step 9. To create a new virtual network associated with this subscription, location, and resource group, do the following:
 - a. Under the **Virtual network** option, click **Create new**.
 - b. On the *Create virtual network* page, in the **Name** text field, enter a name for your new virtual network.


 **NOTE:** Restrict the name to letters, numbers, underscores, periods, or hyphens. See additional name restrictions in the Azure UI.
 - c. In the Address Space area, select the default address range to restrict access to IP addresses between 10.1.0.0 and 10.1.255.255. To restrict to a narrower set
 - d. In the Subnets area, select a subnet name contained by the address range specified in the previous step. For example, select **default**.
 - e. When satisfied with your virtual network settings, click **Create**.
The Advanced page closes and the virtual network you created is selected.
9. On the *Advanced* page, in the Virtual Networks area, if you selected an existing virtual network, from the **Subnets** drop-down menu, select an appropriate subnet for your virtual network.
10. When satisfied with your storage account details, click **Review + create** to validate the settings.
11. If validation errors are found, visit the page with the error, update the parameters you defined, and review again.
12. When satisfied, click **Create**.
You see a message indicating that the deployment of your storage account has started. If toast alerts are active, you see another message when the deployment is completed.

Creating a container in an Azure storage account




The following items are prerequisites:

- You must have administrative access to an account on Azure.
- You must have a storage account defined within your Azure account.


When you perform virtual export, the information is stored in a container within an Azure storage account. You can define the container from your Azure account before performing virtual export, using the procedure below.

 **NOTE:** If you do not define containers in advance, you can choose default containers (named **export** and **deploy**, respectively)

Complete the steps in this procedure to create a container in an Azure storage account.

1. Open the Microsoft Azure dashboard.
2. From the left navigation area, click  **All resources**.
3. From the All resources pane, click the name of the  storage account in which you want to store data from your Rapid Recovery virtual exports.
4. In the Settings pane, under *Services*, click  **Blobs**.
5. From the top of the Blob service pane, in the header, click **+ Container**.

6. From the New container pane, in the **Name** field, type the name for your new container.

 **NOTE:** Type a name between 3 and 63 characters, using only lowercase letters, numbers, and hyphens.

7. From the New container pane, from the **Public access level** drop-down menu, select an access level to define whether the container can be accessed publicly. Use the following as guidance.

Option	Description
Private (no anonymous access)	This option restricts the container to the account owner.
Blob (anonymous read access for blobs only)	This option allows public read access for Binary Large Objects (Blobs).
Container (anonymous read access for containers and blobs)	This option allows public read and list access to the entire container.

For example, select **Container (anonymous read access for containers and blobs)**.

- Click **OK**.
If Toast alerts are active, you should see a message indicating that the container was successfully created.
The Blob service page refreshes, with the new container name displayed in the list.




Creating an Azure Active Directory web application



Perform these steps before attempting virtual export to Azure.






You must use an Azure Active Directory (AD) web application to serve as a connection between your Rapid Recovery Core and your Azure subscription. After creating the web application, record its application ID, and create a secret key associated with the application.

You should also gather the tenant ID. Finally, associate the appropriate privileges to your web application.

Complete the steps in this procedure to create an Azure AD web application with the appropriate keys and privileges.




1. From the Azure navigation menu, click  **Azure Active Directory** and select  **App registrations**.
2. Click **+ New application registration**.
3. On the *Create* page, in the **Name** field, provide a name for your application. Your name must have at least 4 characters.
4. From the **Application type** drop-down field, select **Web app / API**.
5. In the **Sign-on URL** field, enter the URL where a user can sign into Azure and use the app. This value can later be changed, but it must be a valid URL, for example: `http://YourAppLogin.com` or `https://YourSecureAppLogin.net`.
6. When satisfied click **Create**.
The details pane for your web application appears.
7. From the details pane for your web application, copy the Application ID to an easily accessible location (for example, to a Notepad document on your Core server).
8. From the details page for your web application, click  **Settings**.

9. From the *Settings* pane, click  **Keys**.
10. From the *Keys* pane, do the following:
 - a. In the **Description** field, enter a text description to describe the secret key.
 - b. From the **Expires** drop-down menu, select a duration for this secret key, for example, **2 years**.
 - c. From the top of the *Settings* pane, click  **Save**.

CAUTION: Immediately record the secret key description and value in a secure location for the long term. If you do not retain the secret key for your Azure AD web application when you create it, it cannot be recovered.
11. Now obtain the Directory ID for the AD web application (described in the Rapid Recovery Core Console as the Tenant ID) by doing the following:
 - a. From the Azure navigation menu, click  **Azure Active Directory**.
 - b. From the *Properties* pane, scroll down if necessary and click  **Properties**.
 - c. From the *Properties Details* pane, copy the **Directory ID** value to an easily accessible location.
12. Finally, as an Azure user with administrative privileges, add the Owner role to your web application by doing the following:
 - a. From the Azure navigation menu, click  **All services**.
 - b. From the General category, click  **Subscriptions** and then click on your Azure subscription.
 - c. In the Subscription blade, click  **Access control (IAM)**, and then click **+ Add**.
The *Add permissions* dialog box appears.
 - d. From the **Role** drop-down menu, select **Owner**.
 - e. From the **Assign access to** drop-down menu, select **Azure AD user, group, or application**.
 - f. From the **Select** drop-down menu, search for and select the name of your AD web application, and then click **Save**.


Obtaining the application ID for an Azure web application

Complete the steps in this procedure to obtain the application ID for an existing Azure AD web application.

1. From the Azure navigation menu, click  **Azure Active Directory**.
The *Overview* page appears.
2. From the *Overview* page, click  **App registrations**.
The *Application registrations* page appears.
3. Review the list of application registrations, and click the appropriate application.
The details page for the specified application appears.
4. Click  **Click to Copy** to copy the application ID to your clipboard. Retain this information for connecting your Azure account to your Rapid Recovery Core.



Obtaining Azure subscription information

This step requires you to have an Azure subscription. You must have the *Owner* role.

 **NOTE:** If you are not the owner, contact your Azure administrator and request ownership.

When you have an Azure subscription, certain information associated with your subscription is required to perform virtual export to Azure.

Complete the steps in this procedure to obtain information about your Azure subscription.






1. Open the Microsoft Azure dashboard.
2. From the Azure left navigation area, click  **All services**.
3. From the All services pane, click  **+ Subscriptions**.
4. From the *Subscriptions* pane, locate the correct subscription. Note the subscription name, and ID, and verify that your role is *Owner*.

Obtaining the directory ID for your Azure web application

As a prerequisite for this task, you must create an Active Directory (AD) web application in your Azure subscription. For more information, see [Creating an Azure Active Directory web application](#).

When associating an Azure cloud account with the Rapid Recovery Core, one parameter you must provide is the tenant ID. Azure calls this parameter the directory ID. The directory ID is a property of the Active Directory (AD) web application you create within your Azure subscription.

Complete the steps in this procedure to obtain the directory ID for your AD web application from your Azure account.






1. From the Azure navigation menu, click  **Azure Active Directory**.
2. From the *Properties* pane, scroll down if necessary and click  **Properties**.
3. From the *Properties Details* pane, copy the **Directory ID** value to an easily accessible location. This parameter is referred to in the Rapid Recovery Core Console as the tenant ID.
4. Finally, as an Azure user with administrative privileges, add the Owner role to your web application by doing the following:
 - a. From the Azure navigation menu, click  **All services**.
 - b. From the General category, click  **Subscriptions** and then click on your Azure subscription.
 - c. In the Subscription blade, click  **Access control (IAM)**, and then click **+ Add**.
The *Add permissions* dialog box appears.
 - d. From the **Role** drop-down menu, select **Owner**.
 - e. From the **Assign access to** drop-down menu, select **Azure AD user, group, or application**.
 - f. From the **Select** drop-down menu, search for and select the name of your AD web application, and then click **Save**.

Obtaining a secret key for your Azure web application

As a prerequisite for this task, you must create an Active Directory (AD) web application in your Azure subscription. You must have a secret key to connect to and use your Azure Active Directory (AD) web application with Rapid Recovery. Typically, a key is created as part of the process for defining the web application. When creating the secret key, immediately save the secret key to a secure location, as it cannot be recovered or viewed later.

i **NOTE:** The steps for defining your secret key after creating the web application is described in the topic [Creating an Azure Active Directory web application](#)

If you do not record the secret key value, you can create a new secret key for your web application in Azure. Complete the steps in this procedure to create a new secret key associated with an existing Azure AD web application.

1. From the Azure navigation menu, click  **Azure Active Directory** and select  **App registrations**.
2. Click the name of the appropriate AD web application in your Azure account. Details for the selected web application appear.
3. From the details page for your web application, click  **Settings**.
4. From the *Settings* pane, click  **Keys**.
5. From the *Keys* pane, do the following:
 - a. In the **Key description** text area, enter a text description to help identify the purpose of the secret key. For example, to describe this key as the one you will use with Rapid Recovery Core, type **RRCore-key**.
 - b. From the **Expires** drop-down menu, select a duration for this secret key, for example, **2 years**.
 - c. From the top of the *Settings* pane, click  **Save**.

! **CAUTION:** Immediately record the secret key description and value in a secure location for the long term. If you do not retain the secret key for your Azure AD web application when you create it, it cannot be recovered.

Microsoft Azure documentation

Microsoft has substantial documentation on using Azure available in its documentation center.

For information on creating an Azure subscription or user account, selecting Azure resources for VMs you create on Azure, adding a storage account to your VMs, and more, see the Microsoft documentation at <https://azure.microsoft.com/en-us/documentation>.

For example, for information on provisioning or managing Windows VMs, see <https://azure.microsoft.com/en-us/documentation/services/virtual-machines/windows/>.

For online videos about using Azure, see <http://azure.microsoft.com/en-us/get-started/>.

Relevant Microsoft links

Some relevant articles on Microsoft websites are listed below:

- [Azure login page \(US\)](#)
- [Microsoft Azure home page](#)

- [Microsoft documentation center](#)
- [Windows virtual machines documentation](#)
- [Subscriptions, licenses, accounts, and tenants for Microsoft's cloud offerings](#)
- [Videos: Get started with Azure](#)
- [Windows Virtual machines pricing](#)
- [Azure regions](#)
- [Azure storage account overview](#)
- [Understanding Block Blobs, Append Blobs, and Page Blobs](#)
- [Azure Blob storage: hot, cool, and archive access tiers](#)
- [Upgrade to a General-purpose v2 storage account](#)
- [Create a storage account](#)
- [Attach a managed data disk to a Windows VM by using the Azure portal](#)
- [What is Azure Import/Export service? \[Describes using the service to transfer data to BLOB storage\]](#)
- [What disk types are available in Azure?](#)
- [Import/Export Pricing](#)
- [What are Availability Zones in Azure?](#)

Third-party related links

- [Storage: Import/Export Hard Disk Drives to Windows Azure \(blog post\)](#)
- [Azure VM comparison](#). This third-party website offers comparisons on Azure VM costs. If using PowerShell or Command Line commands to perform Azure tasks at the command line, you can obtain the VM name parameter from the first column in the this table.

Exporting and deploying VMs for Azure

Unlike virtual export for other platforms, virtual export for Azure is divided between two processes: exporting, and deploying.

Be advised that Microsoft Azure customers are responsible for their own fees. Some aspects of our integration with Azure are designed with this fee structure in mind. For example, Microsoft charges fees when you deploy a VM on Azure, and when data is transmitted from Azure to another source.

i | **NOTE:** Since Microsoft can change prerequisites, requirements, costs, and so on, always verify such information with Azure. For more information, see the Azure website or contact an Azure representative.

To avoid incurring unnecessary charges, virtual export to Azure consists of two separate processes.

The process of **exporting** extracts the necessary set of files from Rapid Recovery, validates them, and uploads them to the specified container in Azure. These files include:

- One virtual hard disk (VHD) file for each volume in the recovery point.
- One XML file, which contains metadata information about each disk (a list of files present on each disk and a flag indicating if a volume is a system disk).
- One VHD file containing the backup snapshot.

Other than costs for the required storage, exporting by itself does not incur any Azure fees.

The **deployment** process combines these files into a bootable virtual machine. Deployment directly uses Azure cloud REST APIs. The original set of files placed on Azure during the export process is read-only in Azure, and consumes space but does not otherwise incur Azure charges. When you deploy these files, a duplicate copy of them is created, stored in a separate container you define, and combined into a working virtual machine. From an Azure account perspective, after you deploy, you are then charged fees for the VM on its servers. Since the deployed VM is a copy of the exported files, the deployment process also doubles the amount of storage space used in Azure for that virtual export.

For a one-time virtual export, there is no mechanism for deploying as a separate process. Thus, for the export to be useful, you should deploy to Azure when you create the virtual machine on demand. As a result, one-time exports to Azure have an immediate cost associated with the VM you deploy.

When establishing virtual standby for a protected machine on Azure, to avoid use of extra storage space and VM charges, you can simply define the export process. The result is an initial virtual export to Azure which is continually updated. Each time a snapshot is captured on the Core, the exported files are refreshed in your Azure account with updated information. Before the virtual export can be used as a bootable VM, you must deploy it, which triggers VM costs on Azure. If you do not need to convert the exported files for a protected machine to a bootable VM, no VM costs are incurred in your Azure account.

For information about performing a one-time export to Azure, including deployment, see the topic [Performing a one-time Azure export](#).

For information about setting up continual export to Azure, excluding deployment, see the topic [Setting up continual export to Azure](#).

For information about deploying the most recent exported files to create a bootable virtual standby VM in Azure, see the topic [Deploying a virtual machine in Azure](#).

Performing a one-time Azure export

For prerequisites, see [Before virtual export to Azure](#).

Before you can perform a one-time Azure export, you need the following:

- You must have a protected machine with at least one recovery point in a Rapid Recovery Core that you want to export to Azure.
- Remote access must be enabled on the protected machine for the deployed VM to boot successfully.
- You must have administrative access to an account on Azure.
- Optionally, before exporting, you can create an appropriate container within your Azure storage account into which the exported data is stored. For more information, see the topic [Creating a container in an Azure storage account](#).

As described in the topic [Exporting and deploying VMs for Azure](#), virtual export to Azure consists of two processes: exporting and deploying.

The process of **exporting** extracts the necessary set of files from Rapid Recovery, validates them, and uploads them to the specified container in Azure. These files include:

- One virtual hard disk (VHD) file for each volume in the recovery point
- One XML file, which contains metadata information about each disk (a list of files present on each disk and a flag indicating if a volume is a system disk)
- One VHD file containing the backup snapshot



The **deployment** process combines these files into a bootable virtual machine. Deployment uses direct Azure cloud REST APIs. The original set of files placed on Azure during the export process is read-only in Azure, and consumes space but does not invoke additional Azure charges. When you deploy these files, a duplicate copy of them is

created, stored in a separate container you define, and combined into a working virtual machine. From an Azure account perspective, after you deploy you are then charged fees for the VM on its servers. The deployment process also doubles the amount of storage space used in Azure for that virtual export.

For a one-time virtual export, there is no mechanism for deploying as a separate process; thus, for the export to be useful, you should deploy when you create the virtual machine on demand.

When establishing virtual standby for a protected machine on Azure, to avoid use of extra storage space and VM charges, you can export, and continually update the recovery point automatically, without the need to deploy. You can then deploy in Azure only when and if you need to use the VM. For information about deploying a virtual standby in Azure to a functioning VM, see the topic [Deploying a virtual machine in Azure](#).

Complete the steps in this procedure to perform a one-time export to Azure on demand, including deploying to a VM.

1. In the Rapid Recovery Core Console, in the button bar, click the  **Restore** drop-down menu, and then click  **VM Export**.
The *Virtual Machine Export Wizard* appears.
2. In the wizard, from the *Select VM Export Type* page, select **One-time export** and then select **Next**.
3. On the *Machines* page, click the row in the Machines grid representing the protected machine that you want to export, and then click **Next**.
4. On the *Recovery Points* page, scroll through the list of recovery points if necessary, and click to select the recovery point that you want to use for the export. Then click **Next**.
5. On the *Destination* page, in the **Export a virtual machine to** drop-down menu, select **Azure**.

6. Enter the parameters for accessing the virtual machine as described in the following table, and then click **Next**.

Table 130: Azure credentials

Options	Description
Cloud account name	<p>Entering this parameter is optional.</p> <p>Each time you successfully connect to a cloud account, Rapid Recovery Core caches your credentials so you can use them again.</p> <p>Do one of the following:</p> <ul style="list-style-type: none">• If you previously connected this Core with the Azure account you want to export to, then optionally, from the Cloud account name drop-down menu, select the appropriate Azure subscription name, and then click Next. Skip to step 5.• If no value appears in the drop-down menu, or if you want to export to a different cloud account, provide the credential information requested for your Azure account, starting with Region.
Region	<p>Each Azure portal is associated with a geographic region. Choose the region your portal is accessed from. Options include:</p> <ul style="list-style-type: none">• Azure Global Cloud• Azure China Cloud• Azure German Cloud• Azure US Government Cloud
Application ID	<p>Provide the application ID or client ID associated with your Azure Active Directory (AD) web application.</p> <p>i NOTE: If you do not have an Azure AD web application to serve as a connection between your Rapid Recovery Core and your Azure subscription, create one using the steps described in Creating an Azure Active Directory web application. Step 7 describes how to obtain the application ID. Step 10 describes how to create a secret key. Step 11 describes how to obtain the tenant ID or directory ID. Ensure you perform step 12 to provide this application with the <i>Owner</i> role. If you already have an AD web application to use for the Core and you need the application ID, see Obtaining the application ID for an Azure web application.</p>
Secret key	<p>Each web application must have one or more secret keys that you can use to authenticate using Azure APIs. Provide the secret key (or authentication key) associated with your web application.</p> <p>i NOTE: Immediately after creating your secret key in Azure, be sure to copy the key value and store it in a secure location. Once that page is closed, you will no longer be able to view or obtain the secret key. You can, however, create a new secret key.</p> <p>For information about creating a secret key, see Obtaining a secret key for your Azure web application.</p>

Options	Description
Tenant ID	Provide the tenant ID, also known as the directory ID. For information about obtaining this ID, see Obtaining the directory ID for your Azure web application .
Subscription ID	This is the ID associated with your Azure subscription. For information about obtaining this ID, see Obtaining Azure subscription information .

- On the Storage page, from the **Storage account name** drop-down menu, select the existing Azure storage account name in which you want to store your exported Azure VM.
- Optionally, if you want to specify the container and folder names for exporting or deploying, select **Show advanced options**, and proceed to the next step. Otherwise, skip to step 11.

i **NOTE:** When you complete the one-time VM export, the necessary files are exported to an export folder within an export container in the specified storage account. These files are then copied into a deployment folder within a deployment container, and the VM is then deployed from this second location.

9. On the Storage page, with advanced storage options displayed, enter information as described in the following table:

Table 131: Continual export to Azure advanced storage options

Option	Description
Export container	<p>Do one of the following:</p> <ul style="list-style-type: none">When using Rapid Recovery for virtual export, each Azure storage account has a default export container named <code>export</code>. If you want to use the default export container, leave that name in this text box.If you want to specify a different container to hold the exported VM, in the Export container text box, type a new name following Azure's required naming conventions. <p>i NOTE: The container name must be between 3 and 63 characters, starting with a letter or number, and must consist of lower-case letters, numbers, and single hyphens only.</p>
Export folder name	<p>Do one of the following:</p> <ul style="list-style-type: none">By default, the export folder is named after the protected machine from which this recovery point was captured. If you want to use the default export folder name, leave that name in this text box.If you want to specify a different folder name to hold exported VM components, in the Export folder name text box, type a new name. <p>i NOTE: The export folder name does not have the same naming or character restrictions. Nonetheless, best practice is to name the folder with upper and lower-case letters, numbers, hyphens or underscores. Quest recommends that you avoid using prohibited characters or prohibited phrases.</p> <p>When the VM is deployed, a copy of the export folder becomes the deployment folder, containing the same name.</p>
Deployment container	<p>Do one of the following:</p> <ul style="list-style-type: none">When your VM is deployed, Rapid Recovery provides a default deployment container named <code>deploy</code>. If you want to use the default deploy container, leave that name in this text box.If you want to specify a different container to hold the deployed VM, in the Deployment container text box, type a new name following Azure's required naming conventions. <p>i NOTE: The container name must be between 3 and 63 characters, starting with a letter or number, and must consist of lower-case letters, numbers, and single hyphens only.</p>
Resource group	Select an Azure resource group created using Azure Resource Manager.

Option	Description
	<p>i NOTE: If you only have older resource groups created using Azure Service Manager, stop here and create a new resource group in Azure, since Microsoft no longer supports ASM objects. Then return to this step of the procedure and continue.</p>

10. When satisfied with your advanced storage options, click **Next**.

i **NOTE:** If any custom field values do not pass validation, place your cursor over each highlighted text box to see the restrictions. Change the custom values to comply with those rules and then click **Next**.

11. On the Virtual Machine Options page, enter the information described in the following table.

Table 132: Virtual machine options

Option	Description
Virtual machine name	<p>Enter a name for the virtual machine.</p> <p>i NOTE: Type a name between 3 and 15 characters, using only lowercase letters, numbers, and hyphens.</p>
Virtual machine size	<p>From the drop-down menu, select an appropriate VM size.</p> <p>i NOTE: For more information on Azure configurations and pricing, see the virtual machines pricing page on the Azure website. For links to other useful references on Microsoft websites, see Microsoft Azure documentation.</p>
Virtual network	<p>Select a virtual network created in Azure.</p> <p>You can associate a container created in Azure with a virtual network, as described in the topic Creating an Azure storage account.</p> <p>i NOTE: If you do not yet have a virtual network, stop here and create a new virtual network in Azure. Then return to this step of the procedure and continue.</p>

12. Click **Next**.

13. On the *Volumes* page, select the volumes to export (for example, the system reserved volume and `C:\`), and then click **Finish** to close the wizard and start the export.




i **NOTE:** You can monitor the status and progress of the export by viewing the export queue on the *Virtual Standby* page, or on the *Events* page.

Setting up continual export to Azure


For prerequisites, see [Before virtual export to Azure](#).

Complete the steps in this procedure to perform a continual virtual export of the selected machine to a specified container in an Azure cloud account using Rapid Recovery. This process is also known as setting up virtual standby.

i **NOTE:** This process does not include deploying the exported files to create a bootable VM. For steps on deploying, see [Deploying a virtual machine in Azure](#).



1. In the Rapid Recovery Core Console, do one of the following:
 - From the Core Console, in the button bar, click the  **Restore**  drop-down menu, and then select  **VM Export**.

The Virtual Machine Export Wizard appears.

 - a. In the wizard, select **Continual export (virtual standby)**.
 - b. Click **Next**.
 - From the Core Console, in the icon bar, click  (Virtual Standby).
 - On the Virtual Standby page, click **+ Add** to launch the Virtual Machine Export Wizard.
2. On the Machines page of the wizard, select the protected machine that you want to export, and then click **Next**.
3. On the Destination page, from the **Export a virtual machine to** drop-down menu, select **Azure**.

4. Enter the parameters for accessing the virtual machine as described in the following table, and then click **Next**.

Table 133: Azure credentials

Options	Description
Cloud account name	<p>Entering this parameter is optional.</p> <p>Each time you successfully connect to a cloud account, Rapid Recovery Core caches your credentials so you can use them again.</p> <p>Do one of the following:</p> <ul style="list-style-type: none">• If you previously connected this Core with the Azure account you want to export to, then optionally, from the Cloud account name drop-down menu, select the appropriate Azure subscription name, and then click Next. Skip to step 5.• If no value appears in the drop-down menu, or if you want to export to a different cloud account, provide the credential information requested for your Azure account, starting with Region.
Region	<p>Each Azure portal is associated with a geographic region. Choose the region your portal is accessed from. Options include:</p> <ul style="list-style-type: none">• Azure Global Cloud• Azure China Cloud• Azure German Cloud• Azure US Government Cloud
Application ID	<p>Provide the application ID or client ID associated with your Azure Active Directory (AD) web application.</p> <div> NOTE: If you do not have an Azure AD web application to serve as a connection between your Rapid Recovery Core and your Azure subscription, create one using the steps described in Creating an Azure Active Directory web application. Step 7 of that procedure describes how to obtain the application ID. Step 10 describes how to create a secret key. Step 11 describes how to obtain the tenant ID or directory ID. Ensure you perform step 12 to provide the web application with the <i>Owner</i> role.</div> <p>If you already have an AD web application to use for the Core and you need the application ID, see Obtaining the application ID for an Azure web application.</p>
Secret key	<p>Each web application must have one or more secret keys that you can use to authenticate using Azure APIs. Provide the secret key (or authentication key) associated with your web application.</p> <div> NOTE: Immediately after creating your secret key in Azure, be sure to copy the key value and store it in a secure location. Once that page is closed, you will no longer be able to view or obtain the secret key. You can, however, create a new secret key.</div> <p>For information about creating a secret key, see Obtaining a secret key for your Azure web application.</p>

Options	Description
Tenant ID	Provide the tenant ID, also known as the directory ID. For information about obtaining this ID, see Obtaining the directory ID for your Azure web application .
Subscription ID	This is the ID associated with your Azure subscription. For information about obtaining this ID, see Obtaining Azure subscription information .

- On the Storage page, from the **Storage account name** drop-down menu, select the existing Azure storage account name in which you want to store your exported Azure VM.
- Optionally, if you want to specify the export container or provide an export folder name different than the name of your protected machine, select **Show advanced options**, and proceed to the next step. Otherwise, skip to step 11.
- On the Storage page, with advanced storage options displayed, enter information as described in the following table:

Table 134: One-time export to Azure advanced storage options

Option	Description
Export container	<p>Do one of the following:</p> <ul style="list-style-type: none"> When using Rapid Recovery for virtual export, each Azure storage account has a default export container named <code>export</code>. If you want to use the default export container, leave that name in this text box. If you want to specify a different container to hold the exported VM, in the Export container text box, type a new name following Azure's required naming conventions. <p>NOTE: The container name must be between 3 and 63 characters, starting with a letter or number, and must consist of lower-case letters, numbers, and single hyphens only.</p>
Export folder name	<p>Do one of the following:</p> <ul style="list-style-type: none"> By default, the export folder is named after the protected machine from which this recovery point was captured. If you want to use the default export folder name, leave that name in this text box. If you want to specify a different folder name to hold exported VM components, in the Export folder name text box, type a new name. <p>NOTE: The export folder name does not have the same naming or character restrictions. Nonetheless, best practice is to name the folder with upper and lower-case letters, numbers, hyphens or underscores. Quest recommends that you avoid using prohibited characters or prohibited phrases.</p> <p>When the VM is deployed, a copy of the export folder becomes the deployment folder, containing the same name.</p>

- When satisfied with your advanced storage options, click **Next**.

i **NOTE:** If any custom field values do not pass validation, place your cursor over each highlighted text box to see the restrictions. Change the custom values to comply with those rules, and then click **Next**.

- If you see an export warning indicating that the export will replace files from the previous export, if this is acceptable, confirm to close the dialog box and continue.
- On the *Virtual Machine Options* page, if you want to queue the export job immediately, select **Perform initial one-time export**. Clear this option if you want the Core to wait until the next forced or scheduled backup snapshot.
- When satisfied with your virtual machine options, click **Next**.
- On the Volumes page, select the volumes to export (for example, the system reserved volume and C:\), and then click **Finish** to close the wizard and start the export.

i **NOTE:** You can monitor the status and progress of the export by viewing the *Virtual Standby* or *Events* pages.

The virtual standby parameters you define in this procedure cause the export of the files necessary to create a VM in your Azure account. After every snapshot (forced or scheduled), these files are updated in Azure with any new backup information. Before you can boot these files as a VM, you must deploy the VM on Azure. For steps on deploying, see [Deploying a virtual machine in Azure](#).

Deploying a virtual machine in Azure



Before you can deploy a VM in Azure, you must have a protected machine on a Rapid Recovery Core with at least one recovery point, and you must set up continual export (virtual standby) in the Core Console.

On the Azure side, you must have already created a virtual network.

For more information about setting up continual export, see [Setting up continual export to Azure](#). This process also requires you to have an Azure account with sufficient storage associated with your Core.


When you set up virtual standby for a protected machine to Azure, the latest backup information is continually exported from the Core to your Azure account after every backup snapshot. This process overwrites the previous set of export files stored on the Azure VM in the export folder location with updated backup information. Before you can boot the virtual export as a VM (for example, if your original protected machine has failed), you must first select the appropriate virtual standby machine in the Core Console, and deploy it. This process generates a bootable VM in the deployment container within Azure.


Complete the steps in this procedure to deploy your most recent virtual standby export files to a bootable VM in Azure.


- From the Rapid Recovery Core Console, in the icon bar, click  (Virtual Standby).
- In the Virtual Standby pane, identify the machine in your Rapid Recovery Core that is set up for continual export to Azure.
- From the row representing the virtual standby machine you want to deploy, click  (More options) and then select **Deploy Virtual Machine**.


The Deploy to Azure Wizard appears on the *Destination* page. Several parameters are automatically populated with information.


4. On the *Destination* page, do one of the following:
 - If you want to use a deployment container that already exists in your Azure account, then from the **Deployment container** drop-down menu, select the appropriate container name.
 - If you want to create a new deployment container in your Azure account, then in the **Deployment container** text box, type the name for your new deployment container.
5. From the **Resource group** drop-down menu, select the Azure resource group with which you want to associate the deployed virtual machine.
6. Click **Next**.
The *Virtual Machine Options* page appears.
7. On the *Virtual Machine Options* page, in the **Virtual machine name** text box, enter a name for the virtual machine.


 **NOTE:** Type a name between 3 and 15 characters, using only letters, numbers, and hyphens. The name must start with a letter and cannot end with a hyphen.
8. From the **Virtual machine size** drop-down menu, select the appropriate size for the new VM you want to create within Azure. The size includes a number of processor cores, the desired amount of virtual memory, and the required number of data disks.

 **NOTE:** For more information on Azure configurations and pricing, see the [virtual machines pricing](#) page on the Azure website. For links to other useful references on Microsoft websites, see [Microsoft Azure documentation](#).
9. From the **Virtual network** drop-down menu, select the appropriate virtual network controller to associate with your new VM.
10. Click **Next**.
The *Disks* page appears. If multiple disk volumes exist in the source recovery point, each disk appears in a separate row on the Disks table.
11. From the *Disks* page, select the disks you want to export to your new VM.

 **NOTE:** Your VM must include a system disk. Accordingly, the system disk is automatically selected and cannot be excluded from the new VM.
12. When satisfied, click **Finish** to close the wizard and start the deployment.
The Deploy to Azure Wizard closes and a Deploy job is queued. If resources are available, the deployment begins immediately.

 **NOTE:** If Toast alerts are enabled, you can open the *Monitor Active Task* dialog box to view the progress. Alternatively, you can monitor the progress of the deployment by viewing tasks on the *Events* page.

Once the deployment completes, in your Azure account, you can see the new VM in Azure's  **Virtual machines** view.

-  **CAUTION:** Once the VM is available, you are also paying fees. To avoid ongoing charges from Microsoft, delete the deployed VM when it is not needed. You can always deploy a VM from the latest set of virtual export files by repeating this procedure.

Managing exports

If your Core has continual export set up, the configuration parameters for each virtual export appear as a row on the Virtual Standby page. From here you can view the status of established continual exports, and manage your virtual standby machines. You can add a virtual standby, force export, pause or resume virtual standby, or remove the requirements for continual export from your Core Console.

When a one-time export takes place, the job is listed in the export queue on the Virtual Standby page. During this time, you can pause, resume, or cancel the one-time export operation.

Virtual export to a virtual standby VM does not occur if the VM is powered on.

Complete the steps in this procedure to manage virtual exports.



1. On the Rapid Recovery Core Console, in the icon bar, click  (Virtual Standby).
The Virtual Standby page appears. Here you can view two tables of saved export settings. They include the information described in the following table.

Table 135: Virtual standby information

Column	Description
Select item	For each row in the summary table, you can select the check box to perform actions from the list of menu options preceding the table.
Status indicator	<p>Colored spheres in the Status column show the status of virtual standby. If you hover the cursor over the colored circle, the status condition is displayed.</p> <ul style="list-style-type: none">• Green. Virtual standby is successfully configured, is active, and is not paused. The next export is performed immediately following completion of the next snapshot.• Yellow. Virtual standby pauses, but the parameters are still defined and saved in the Core. However, after a new transfer, the export job will not start automatically and there will be no new exports for this protected machine until the status changes.
Machine Name	The name of the source machine.
VM Status	This column shows for each virtual standby definition whether continual export has been initiated.
Destination	The virtual machine and path to which data is being exported.
Export Type	This column shows the type of virtual machine platform for the export, such as vCenter/ESXi, VMware Workstation, Hyper-V, VirtualBox, or Azure.
Hypervisor Status	This column displays the availability of the hypervisor host.
Last Export	<p>This column shows the date and time of the last export.</p> <p>If an export has just been added but has not completed, a message displays stating the export has not yet been performed. If an export has failed or was canceled, a corresponding message also displays.</p>
Settings	<p>The  (More options) drop-down menu lets you perform the following functions:</p> <ul style="list-style-type: none">• Edit. Lets you edit the virtual standby settings.• Force. Forces a virtual export.• Pause. Pauses virtual export. Only available when status is active.• Resume. Resumes virtual export. Only available when status is paused.• Remove. Removes the requirement for continual export. Does not remove the exported VM most recently updated.



Column	Description
	<ul style="list-style-type: none"> • Start VM. Starts an already-exported virtual machine. <div>  NOTE: New data cannot be written to the virtual standby machine when the VM is started. </div> <ul style="list-style-type: none"> • Stop VM. Stops an already-exported virtual machine. • Network Adapters. Lets you add or modify virtual network adapters. • Deploy Virtual Machine. For Azure continual export only, this option converts the exported files in your Azure account to a bootable VM.

Table 136: Export queue information

Column	Description
Select item	<p>For each row in the summary table, you can select the check box to perform actions from the list of menu options preceding the table. These options include:</p> <ul style="list-style-type: none"> • Cancel. Cancel the current one-time virtual export. • Settings. Lets you update the maximum concurrent exports setting.
Status indicator	Shows as a percentage the status of the current export. When no one-time exports are queued, this column has no value.
Machine Name	The name of the source machine.
Destination	The virtual machine and path to which data is being exported.
Export Type	This column shows the type of virtual machine platform for the export, such as vCenter/ESXi, VMware, Hyper-V, VirtualBox, or Azure.
Schedule Type	Click on  to see the schedule type. This shows the type of export as either One-time or Continuous.

- To manage saved export settings, select an export, and then click one of the following:
 - **Edit.** Opens the Virtual Machine Export Wizard to the VM Options page. Here you can change the location of the exported VM, change the version of the VM type, or specify RAM or processors for the export. To immediately start the VM export, select **Perform initial one-time export**.
 - **Force.** Forces a new export. This option could be helpful when virtual standby is paused and then resumed, which means the export job will restart only after a new transfer. If you do not want to wait for the new transfer, you could force an export.
 - **Pause.** Pauses an active export.
 - **Resume.** Resumes the requirement for continue export at the next scheduled or forced snapshot.
- To remove an export from the system, select the export, and then click **Remove**. The export configuration is permanently removed from the system. Removing the virtual standby configuration does not remove any virtual machine exported as a result of the configuration.

4. To deploy a VM to Azure, select **Deploy Virtual Machine** and complete details in the Deploy to Azure Wizard.
Data from the most recent virtual export saved to your Azure account is deployed within your associated Azure account as a bootable VM.
5. To manage the number of exports that run at the same time, do the following:
Under Export Queue, click **Settings**.
In the Maximum Concurrent Exports dialog box, enter the number of exports you want to run simultaneously.
The default number is 5.
Click **Save**.
6. To cancel a one-time or continual export currently listed in the Export Queue, select the export, and then click **Cancel**.
7. To add a new virtual standby export, you can click **Add** to launch the Export Wizard. Completing the resulting wizard results in a continual export for the selected protected machine. For further information about setting up virtual standby for a specific virtual machine, see one of the following topics:
 - [Setting up continual export to ESXi](#)
 - [Setting up continual export to VMware Workstation](#)
 - [Setting up continual export to Hyper-V](#)
 - [Setting up continual export to VirtualBox](#)
 - [Setting up continual export to Azure](#)

Restoring data

This section describes how to restore data from recovery points saved to your repository using Rapid Recovery Core.

Topics include:

- [About restoring data with Rapid Recovery](#)
- [Understanding Live Recovery](#)
- [Restoring data from recovery points](#)
- [VMware VM configuration backup and restore](#)
- [About the file search and restore feature](#)
- [About restoring volumes from a recovery point](#)
- [Restoring clusters and cluster nodes](#)
- [Restoring from an attached archive](#)
- [Mail Restore in Rapid Recovery](#)

About restoring data with Rapid Recovery

The Rapid Recovery Core can instantly restore data or recover machines to physical or virtual machines from recovery points. The recovery points contain agent volume snapshots captured at the block level. These snapshots are application aware, meaning that all open transactions and rolling transaction logs are completed and caches are flushed to disk before creating the snapshot. Using application-aware snapshots in tandem with Verified Recovery enables the Core to perform several types of recoveries, including:

- Recovery of files and folders
- Recovery of data volumes, using Live Recovery
- Recovery of data volumes for Microsoft Exchange Server and Microsoft SQL Server, using Live Recovery
- Bare metal restore, using Universal Recovery
- Bare metal restore to dissimilar hardware, using Universal Recovery
- One-time on demand and continual export to virtual machines

i **NOTE:** When you restore data or perform virtual export, the recovery point used must be part of a complete recovery point chain. For more information about recovery point chains, see the topic [Recovery point chains and orphans](#).

Understanding Live Recovery

Live Recovery is a feature of restoring data in Rapid Recovery Core. If your protected machine experiences data failure of a non-system Windows volume, you can restore data from a recovery point on the Rapid Recovery Core. Selecting Live Recovery in the Restore Machine Wizard allows users to immediately continue business operations with near-zero downtime. Live Recovery during restore gives you immediate access to data, even while Rapid Recovery continues to restore data in the background. This feature allows near-zero recovery-time, even if the restore involves terabytes of data.

Rapid Recovery Core uses unique block-based backup and recovery technology that allows full user access to target servers during the recovery process. Requested blocks are restored on-demand for seamless recovery.

Live Recovery applies to physical and virtual machines protected by Rapid Recovery Agent, with the following exclusions:

- Live Recovery is accessible to Windows volumes, excluding system volumes. The C : \ drive and the system-reserved partition cannot be restored using Live Recovery.
- Live Recovery is accessible to Windows-based volumes using the Rapid Recovery Agent. Agentless volumes or Linux volumes cannot take advantage of Live Recovery.

Live Recovery lets you instantly restore physical or virtual servers directly from the backup file. When a non-system volume is being restored, Rapid Recovery presents the volume metadata to the operating system instantly, making that data available on demand. For example, if the database volume of Microsoft Exchange is corrupt, Live Recovery can restore the volume, database, and Exchange services in minutes.

This feature provides the fastest method of recovering large quantities of data with minimal downtime. Users can immediately continue business operations.

Once Live Recovery begins, the restored volume and its contents become instantly available. Rapid Recovery Core continues to restore the data in the background, even though the volume, its data, applications and services are already back in production. If specific data is requested, the background process prioritizes the restoration of this data immediately. This powerful functionality allows even the most stringent service-level agreement to be met.

Once you start Live Recovery, metadata (directory structure, security descriptors, NTFS file attributes, free space map, and so on) of the target volume is quickly restored on the protected machine. Thereafter, the volume and its contents become available to the system. The Rapid Recovery Agent begins restoring data blocks from the Rapid Recovery Core server, writing the blocks to the target volume.

Requests for data that has not yet been restored are immediately answered, with the requesting program or system unaware that the blocks were just restored.

Restoring data from recovery points

Rapid Recovery protects your data on Windows and Linux machines. Backups of protected machines are saved to the repository associated with your Rapid Recovery Core as recovery points. From these recovery points, you can restore your data using one of the following methods.

- From the Rapid Recovery Core Console, you can restore entire volumes from a recovery point of a non-system volume, replacing the volumes on the destination machine. This process uses the Restore Machine Wizard. For more information, see [About restoring volumes from a recovery point](#).
- You can also restore entire volumes on Linux machines from recovery points using the command line from the protected Linux machine. For more information on using the command line `local_mount` utility, see [Restoring volumes for a Linux machine using the command line](#).

When restoring data for agentlessly protected machines only, the *Volume Mapping* page of the Restore Machine Wizard includes the option **Restore all configuration data**. This option is associated with the VM configuration and restore feature. For more information, see [VMware VM configuration backup and restore](#).

You cannot restore a volume that contains the operating system directly from a recovery point, because the machine to which you are restoring is using the operating system and drivers that are included in the restore process. If you want to restore from a recovery point to a system volume (for example, the C drive of the protected machine), you must perform a Bare Metal Restore (BMR). This involves creating a bootable image from the recovery point, which includes operating system and configuration files as well as data. You then start the target machine from that bootable image to complete the restore. The boot image differs if the machine you want to restore uses a Windows operating system or a Linux operating system. If you want to restore from a recovery point to a system volume on a Windows machine, see [Performing a bare metal restore using the Restore Machine Wizard](#). If you want to restore from a recovery point of a system volume on a Linux machine, see [Performing a bare metal restore for Linux machines](#).

If you have a software RAID on a Linux machine protected by Rapid Recovery Agent release 6.2 or later, you can restore the software RAID from a recovery point.

NOTE: Since this feature was introduced in release 6.2, it is not compatible for snapshots taken on earlier Agent versions. If you upgrade Rapid Recovery Agent to release 6.2 or later and then capture snapshots in your Rapid Recovery Core, you will be able to restore the software RAID from the new snapshots.

Finally, in contrast to restoring entire volumes, you can mount a recovery point from a Windows machine, and browse through individual folders and files to recover only a specific set of files. For more information, see [Restoring a directory or file using Windows Explorer](#). If you need to perform this while preserving original file permissions (for example, when restoring a user's folder on a file server), see [Restoring a directory or file and preserving permissions using Windows Explorer](#).

The topics in this section describe information about restoring data on physical machines. For more information on exporting protected data from a recovery point to a virtual machine, see [VM export](#).

NOTE: When recovering data on Windows machines, if the volume that you are restoring has Windows data deduplication enabled, you will need to make sure that deduplication is also enabled on the Core server.

VMware VM configuration backup and restore

Rapid Recovery Core release 6.8 introduces a new feature, the ability to back up and restore VMware VM configurations, including the option to include VM configurations during virtual export to VMware/ESXi virtual machines.

Backup. Rapid Recovery Core release 6.3 and later automatically saves agentlessly protected ESXi virtual machine configurations in each volume image when snapshots are captured. VMware virtual machine configurations are stored in `.vmx` files (and related BIOS settings are stored in `.nvram` files). The relevant files are saved in the custom metadata for each relevant VM volume, and includes hypervisor version information to ensure compatibility.

Restore. Optionally, when restoring data from a recovery point of an agentlessly protected ESXi machine, you can choose whether to include in the VM all VM configurations and data, or only the data. This choice is presented in the UI through the **Restore all configuration data** check box. This option appears only for VMware machines protected agentlessly (replacing the **Show advanced options** check box that is relevant only for machines protected by Rapid Recovery Agent). When the option is selected, all VM configurations for volumes being recovered are restored. When the option is cleared, only data (and not VM configurations) are restored for those volumes.

Virtual export. Optionally, when performing virtual export from a recovery point of an agentlessly protected ESXi machine to VMware/ESXi, you can choose whether to export all VM configurations and data, or export only the data. This choice is presented in the UI through the **Restore all configuration data** check box. This option appears only for agentlessly protected ESXi machines. When the option is selected, all VM configurations for volumes being exported to a VM are included in the exported VM. When the option is cleared, only data (and not VM configurations) are included in the exported VM.

Based on the restore or virtual export type, The **Restore all configuration data** option is selected by default in the following situations:

- When restoring data or performing virtual export from a recovery point to the same agentless virtual machine.
- When performing virtual export to a different server. There is no backward compatibility between hypervisor versions.

Otherwise, the **Restore all configuration data** option is not selected by default, although you can change the default option by selecting or clearing this setting.

About the file search and restore feature

The Rapid Recovery file search and restore feature lets you find one or more files in the recovery points of a protected machine. You can then restore one or more of the results to a local disk.

Searching guidelines

On the File Search page of the Core Console, you can search for a file from a set of recovery points from the machine that you select. The search criteria are divided into two groups: basic and advanced.

The basic group includes the following parameters:

- The protected machine whose recovery points you want to search.
- A time range that limits the search to only recovery points that were created between the start time and end time.
- The name or mask of the file that you want to find. You can use the "?" wildcard to replace any single character and the "*" to replace zero or multiple characters; however, more specific filenames produce more specific results.
- A list of paths to directories in which to search.



NOTE: All basic criteria is required. If no directory is provided, Rapid Recovery searches all volumes of the specified protected machine.

The **More Options** button reveals the advanced group, which includes the following parameters:

- The option to search recursively in subdirectories of the search location or only in the specified location.
- The ability to run an algorithm that increases the speed of searches on NTFS volumes.
- The ability to limit the number of search results to a more manageable sum.

i **NOTE:** Specific search criteria produce faster and more accurate your search results. Including subdirectories (for example, `C:\work\documents\accounting` instead of `C:`) reduces the amount of time it takes to complete the search, as does providing restrictive file masks (for example, `invoice*.pdf` instead of `in*.*`).

Because the feature continues to search through recovery points and locations even after the requested file is found, you can pause or stop a search before it completes. You can run multiple searches simultaneously, but you cannot begin them at the same time. For example, to find another file, you can begin a second search while the first search is still in progress. However, you can only search one protected machine at a time.

i **NOTE:** In the previous example, pausing the first search makes more memory available for the second search, which helps the second search finish faster. Running multiple searches at one time is memory intensive and increases the amount of time it takes to complete a search.

Each search appears as a tab on the page. When you are finished searching, you can close the tabs individually or all at once.

Restoring guidelines

After you find the file, you can restore it directly from the File Search page.



The file search and restore feature limits restoring capabilities to only locations on the Core. You cannot restore a file to a protected machine.

Finding and restoring a file

When you want to restore a file instead of a volume, you can use Rapid Recovery to find that file among the recovery points for your protected machine. Search criteria, such as date range and directory, let you narrow the search to a small group of relevant recovery points.

i **NOTE:** Specific search criteria produce faster and more accurate your search results, and consume less memory. Including subdirectories (for example, `C:\work\documents\accounting` instead of `C:`) reduces the amount of time it takes to complete the search, as does providing restrictive file masks (for example, `invoice*.pdf` instead of `in*.*`).

After you find the file, you can then restore it directly from the list of search results.

1. From the Rapid Recovery Core Console icon bar, click the  More drop-down menu, and select  File Search.
The File Search page opens.

- On the File Search page, to search for a file within the recovery points of a specific protected machine, complete the information described in the following table.

Table 137: File search criteria


Text Box	Description
Machine	<p>Select the protected machine that you want to search from the drop-down list.</p> <p>NOTE: You can search through the recovery points of only one protected machine at a time.</p>
Recovery points date range	<p>Specify the date and time of the oldest recovery point and the newest recovery point that you want to search. Only the recovery points created within this span of time are searched.</p> <p>NOTE: The default time span is the previous month. For example, if conducting the search on August 22, 2018 at 2:04 PM, the default date range is 7/22/2018 2:04 PM to 8/22/2018 2:04 PM.</p>
Filename (can use * and ? wildcards)	<p>Enter the name of the file or a file mask for the file that you want to find and restore. Wildcards may be used as substitutes for unknown characters.</p> <p>NOTE: You can use the "?" wildcard to replace any single character and the "*" wildcard to replace zero or multiple characters.</p>
Directories to search	<p>List one or more directories on the protected machine to limited the search to only those locations.</p> <p>NOTE: If no directory is provided, Rapid Recovery searches all volumes of the specified protected machine.</p>

- Optionally, click **More Options**, and then complete the information described in the following table.

Table 138: More file search options


Text Box	Description
Include subdirectories	Searches all the subdirectories of the directories listed in step 3 . Enabled by default.
Use fast search algorithm for NTFS volumes	<p>Searches NTFS volumes without mounting them by parsing file system data structures, which is faster and consumes less memory while searching. Enabled by default.</p> <p>NOTE: If you encounter an issue during a search of an NTFS volume, attempt the search again without this option selected.</p>
Limit search results to	Enter the maximum number of results that you want to appear in the results. The default is 1000.


- Click **Start Search**
The search begins. Each search appears as a tab under Search Results. You can use the buttons for each tab to pause or stop a search, or you can click the **X** on the tab to delete the search. Multiple searches can run simultaneously.
- From the search results, select the file that you want to restore.


6. Click  **Restore**.
The Restore Files dialog box opens.
7. For **Location**, enter a destination path for the restored file on the machine on which the Core is installed and running.
8. Click **Restore**.
The file you selected is restored to the specified destination path with the original directory tree in which the file appeared on the protected machine.

About restoring volumes from a recovery point

You can restore the volumes on a protected machine from the recovery points stored in the Rapid Recovery Core using the Restore Machine Wizard.

 **NOTE:** In earlier releases, the restore process was referred to as performing a rollback.

 **NOTE:** Rapid Recovery supports the protection and recovery of machines configured with EISA partitions.

For Windows or Linux machines, you can begin a restore from any location on the Rapid Recovery Core Console by clicking the  **Restore** icon in the Rapid Recovery button bar. When you start a restore in this manner, you must specify which of the machines protected on the Core you want to restore, and then drill down to the volume you want to restore.

Or you can go to **Recovery Points** page for a specific machine, click the drop-down menu for a specific recovery point, and then select **Restore**. If you begin a restore in this manner, then follow start with [step 5](#) in the topic [Restoring volumes from a recovery point](#).


You can also restore from a recovery point on a Linux machine from the command line. For more information, see the topic [Restoring volumes for a Linux machine using the command line](#).

If you want to restore from a recovery point to a system volume, or restore from a recovery point using a boot CD, you must perform a Bare Metal Restore (BMR). For information about BMR, see [Bare metal restore](#). You can access BMR functions for both Windows and Linux machines using the Restore Machines Wizard, accessible from the button bar of the Core Console. Ensure you read [Prerequisites for performing a bare metal restore for Windows or Linux machines](#) before attempting the process. For specific instructions, see the procedure [Performing a bare metal restore using the Restore Machine Wizard](#).

Restoring volumes from a recovery point

To restore volumes from a recovery point, your machine must be protected on the Core at the volume level, and the Core must contain recovery points from which you perform the restore operation.

Complete the following procedure to restore volumes from a recovery point.

1. To restore a volume on a protected machine, navigate to the Rapid Recovery Core Console and click  **Restore** from the Rapid Recovery button bar.
The Restore Machine Wizard appears.

2. From the *Machines* page, select the protected machine for which you want to restore data, and then click **Next**.

The *Recovery Points* page appears.

3. From the list of recovery points, search for the snapshot you want to restore to the protected machine. If necessary, use the buttons at the bottom of the page to display additional pages of recovery points. Optionally, to limit the number of recovery points showing in the *Recovery Points* page of the wizard, you can filter by volumes (if defined) or by creation date of the recovery point.

4. Click any recovery point to select it, and then click **Next**.

The *Destination* page appears.

5. On the *Destination* page, choose the machine to which you want to restore data as follows:

- To restore data from the selected recovery point to the same machine, and if the volumes you want to restore do not include the system volume, then select **Recover to a protected machine**, verify that the destination machine is selected, and then click **Next**.
The *Volume Mapping* page appears. Proceed to [step 9](#).
- To restore data from the selected recovery point to a different protected machine (for example, replace the contents of Machine2 with data from Machine1), then select **Recover to a protected machine**, select the destination machine from the list, and then click **Next**.
The *Volume Mapping* page appears. Proceed to [step 9](#).
- If you want to restore from the selected recovery point to the same machine or a different machine using a boot CD, this process is considered a bare metal restore (BMR). For information about BMR, see [Bare metal restore](#).

i **NOTE:** Performing a BMR has specific requirements, based on the operating system of the machine you want to restore. To understand these prerequisites, see [Prerequisites for performing a bare metal restore for Windows or Linux machines](#).

If the volumes you want to restore include the system volume, then select **Recover to any target machine using a Boot CD**. This option prompts you to create a boot CD.

- To continue and create the boot CD with information from the selected recovery point using the Restore Machine Wizard, click **Next** and proceed to [Performing a bare metal restore using the Restore Machine Wizard](#).
- If you have already created the boot CD and the target machine has been started using the boot CD, then proceed to [step 8](#) of the topic [Performing a bare metal restore using the Restore Machine Wizard](#).

If you want to restore from a recovery point to a system volume (for example, the C: / drive of the agent machine named Machine1), this process is also considered a BMR. Select **Recover to any target machine using a Boot CD**. This option prompts you to create a boot CD.

- To continue and create the boot CD with information from the selected recovery point using the Restore Machine Wizard, click **Next** and proceed to [Performing a bare metal restore using the Restore Machine Wizard](#).
 - If you have already created the boot CD, then proceed to [step 6](#).
6. Start the machine you want to restore to using the boot CD. For more information, for BMR on a Windows machine, see [Loading the boot CD and starting the target machine](#) and for BMR on a Linux machine, see [Loading the Live DVD and starting the target machine](#).

7. On the Core server, in the *Destination* page of the Restore Machine Wizard, select **I already have a Boot CD running on the target machine**, and then enter the information about the machine to which you want to connect described in the following table.

Table 139: Machine information

Text Box	Description
IP Address	The IP address of the machine to which you want to restore. This is identical to the IP address displayed in the URC.
Authentication Key	The specific password to connect to the selected server. This is identical to the Authentication Key displayed in the URC.

8. Click **Next**.

If the connection information you entered matches the URC, and if the Core and the target server can identify each other properly on the network, then the volumes for the selected recovery point are loaded. The *Disk Mapping* page appears.

To complete your BMR from the Restore Machine Wizard, proceed to [step 9](#) of the topic [Performing a bare metal restore using the Restore Machine Wizard](#).

i **NOTE:** Rapid Recovery supports FAT32 and ReFS partitions. Only full restore and BMR are supported as a driver limitation exists with ReFS. Restore is implemented in user mode, VM export, and so on. If a Core is protecting at least one agent volume that contains the ReFS file system, it should be installed on Windows 10, Windows Server 2012, Windows Server 2016, or the most recent version of Windows Server 2019, which provide native support of ReFS.

Otherwise, functionality is limited and operations that involve such things as mounting a volume image do not work. The Rapid Recovery Core Console presents applicable error messages in these occurrences. Bare metal restore of Storage Spaces disks configuration (a feature of Windows 8.1) is also not supported in this release. For details, see the *Rapid Recovery Installation and Upgrade Guide*.

9. On the *Volume Mapping* page, for each volume in the recovery point that you want to restore, select the appropriate destination volume. If you do not want to restore a volume, in the Destination Volumes column, select **Do not restore**.

i **NOTE:** You must restore at least one volume.

10. If restoring a machine protected by Rapid Recovery Agent, skip to step 14.
11. If restoring an agentlessly protected machine, review the **Restore all configuration data** option and do one of the following:
 - If you want to restore backed-up VM configurations from the recovery point, select this option.
 - If you want to restore data only and not VM configurations, clear this option.
For more information, including an explanation of when this option is selected or cleared by default, see [VMware VM configuration backup and restore](#).
12. Click **Next**.
13. For agentlessly protected machines, if a *Warnings* page appears, read the message, make changes if necessary, and then click **Next**. Skip to step 18.

14. Select **Show advanced options** and then do the following:
 - For restoring to Windows machines, if you want to use Live Recovery, select **Live Recovery**. Using the Live Recovery instant recovery technology in Rapid Recovery, you can instantly recover or restore data to your physical machines or to virtual machines from stored recovery points of Windows machines, which includes Microsoft Windows Storage Spaces. Live Recovery is not available for Linux machines or VMs using agentless protection.
 - If you want to force the selected volumes to dismount before the restore begins, select **Force Dismount**.

CAUTION: If you do not force a dismount before restoring data, the restore may fail with an error stating that the volume is in use.

15. Click **Next**.
16. On the *Dismount Databases* page, if the volumes you want to restore contain Oracle, SQL or Microsoft Exchange databases, you are prompted to dismount them. If you want to remount these databases automatically after the restore is complete, select **Automatically remount all databases after the recovery point is restored**. Otherwise, clear this option.

NOTE: The appropriate VSS writer captures snapshots of database in backup mode. If you do choose not to remount all databases automatically (the default option), then after you restore, you must manually start the databases.
17. Click **Next**.

The *Warning* page may appear and prompt you to close all programs on the volumes that you want to restore. If it does, click **Next** again.
18. On the *Summary* page, select the option **IMPORTANT! I understand that this operation will overwrite selected volumes with the data from the selected recovery point** to acknowledge that you understand the consequences of a volume restore.

CAUTION: This option emphasizes the consequence that any data that was saved on the selected volume after the date and time of the selected recovery point is lost upon restore.
19. Click **Finish**.

Restoring a directory or file using Windows Explorer

You can use Windows Explorer to copy and paste directories and files from a mounted recovery point to any Windows machine. This can be helpful when you want to distribute only a portion of a recovery point to your users. When you copy directories and files, the access permissions of the user who is performing the copy operation are used and applied to the pasted directories and files. If you want to restore directories and files to your users while preserving original file permissions (for example, when restoring a user's folder on a file server), see [Restoring a directory or file and preserving permissions using Windows Explorer](#).

1. Mount the recovery point that contains the data you want to restore. For details, see [Mounting a recovery point](#).
2. In Windows Explorer, navigate to the mounted recovery point and select the directories and files that you want to restore. Right-click and select **Copy**.

3. In Windows Explorer, navigate to the machine location to where you want to restore the data. Right-click and select **Paste**.

Restoring a directory or file and preserving permissions using Windows Explorer

You can use Windows Explorer to copy and paste directories and files from a mounted recovery point to any Windows machine while preserving file access permissions.

For example, if you need to restore a folder accessed only by specific users on a file server, you can use the Copy and Paste with Permissions commands to ensure that the restored files retain the permissions that restrict access. In this way, you can avoid having to manually apply permissions to the restored directories and files.

Some files have file access restrictions that require administrative privileges. Especially for Windows Server 2012 and later operating systems, the user attempting the restore must have the correct NTFS permissions for restoring with permissions to be successful. For example, to copy full NTFS permissions from a mount point, the user must have administrative privileges (with full NTFS permissions).

i **NOTE:** The Paste with Permissions command is installed with Rapid Recovery Core and Agent software. It is not available in the Local Mount Utility.

1. Mount the recovery point that contains the data you want to restore. For details, see [Mounting a recovery point](#).
2. In Windows Explorer, navigate to the mounted recovery point and select the directories and files that you want to restore. Right-click and select **Copy**.
3. In Windows Explorer, navigate to the machine location to where you want to restore the data. Right-click and select **Paste with Permissions**.

i **NOTE:** In this step, if the Paste with Permissions command is disabled on the right-click menu, then Windows Explorer is not aware of the files that you want to copy. Repeat [step 2](#) to enable the Paste with Permissions command on the right-click menu.

Restoring clusters and cluster nodes

A restore is the process of restoring the volumes on a machine from recovery points. For a server cluster, you perform a restore at the node, or machine, level. This section provides guidelines for performing a restore for cluster volumes.

Performing a restore for CCR and DAG (Exchange) clusters

Complete the steps in this procedure to perform a restore for CCR and DAG (Exchange) clusters.

1. Turn off all nodes except one.
2. Perform a restore using the standard Rapid Recovery procedure for the machine as described in [About restoring volumes from a recovery point](#) and [Restoring volumes for a Linux machine using the command line](#).
3. When the restore is finished, mount all databases for the cluster volumes.
4. Turn on all other nodes.
5. For Exchange, navigate to the Exchange Management Console, and, for each database, perform the Update Database Copy operation.

Performing a restore for SCC (Exchange, SQL) clusters

Complete the steps in this procedure to perform a restore for SCC (Exchange, SQL) clusters.

1. Turn off all nodes except one.
2. Perform a restore using the standard Rapid Recovery procedure for the machine as described in [About restoring volumes from a recovery point](#) and [Restoring volumes for a Linux machine using the command line](#).
3. After the restore is finished, mount all databases from the cluster volumes.
4. Turn on all other nodes one-by-one.

i **NOTE:** You do not need to roll back the quorum disk. It can be regenerated automatically or by using cluster service functionality.

Restoring from an attached archive

There are two ways you can restore data from an archive: You can use an archive as a source for a bare metal restore (BMR); or you can attach an archive, mount a recovery point from the archive, and then restore the archived data.

When you attach an archive, it appears under Attached Archives on the Archives page of the Core Console, while the contents of the archive become accessible from the left navigation area. The contents appear under the name of the archive. Machines that were archived appear as recovery-points-only machines so that you can access the recovery points in the same way that you would for a currently protected machine: by mounting a recovery point, locating the item that you want to recover, and using Windows Explorer to copy and paste the item to your destination.

There are advantages to restoring from an attached archive rather than importing an archive to a repository.

- Restoring from an attached archive saves the time you may spend importing an entire archive to a repository.
- Also, when you import an archive, the archived recovery points are added to the repository. Because these archived recovery points are likely the oldest items in the repository, they may be rolled up according to your retention policy during the next nightly job. (Although, this action does not delete them from the archive; you could re-import them the next day.)

- Lastly, the Core remembers the attachment association with archives, even after you detach an archive, making it easier and faster to attach the archive again later.
You can remove the association by deleting the attachment.

To restore data from an attached archive, complete the following steps using the related links:

i **NOTE:** The procedure for restoring from an attached archive assumes that you already have an archive of rolled-up recovery points.


1. Attach the archive.
2. Mount the recovery point that contains the data that you want to recover.
3. Restore data using any of the following methods:
 - Restore data, such as file or folder, from the recovery point.
 - Restore the entire recovery point.
 - Export the recovery point to a virtual machine.

For more information, see the following related topics:

- [Understanding archives](#)
- [Attaching an archive](#)
- [Importing an archive](#)
- [Mounting a recovery point](#)
- [Restoring volumes from a recovery point](#)
- [Exporting to virtual machines using Rapid Recovery](#)
- [Performing a bare metal restore using the Restore Machine Wizard](#)

Mail Restore in Rapid Recovery

The Mail Restore feature in Rapid Recovery lets you restore a mailbox, folder, or item — such as a message, calendar event, or contact — from the data store of a protected Exchange Server machine. You can restore your selection to a recovery folder, to the original source, or to one or more PST files.

You can access the Mail Restore page from the  (More) menu of the Rapid Recovery Core Console. From there, you can complete the following actions:

- Open an Exchange database
- Restore an item from the open database
- Close the database
- Search for an item in the open database

For more information, see [Opening an Exchange database in Rapid Recovery Core](#) and [Restoring a mail item in Rapid Recovery](#).

Mail Restore prerequisites

Before you can restore mail items, you must meet the following prerequisites:

- Outlook 2007 or later is installed on the Core machine.
- There is at least one profile configured in Microsoft Outlook.
- The Outlook profile has full-control permissions, including Send As and Receive As permissions. For more information, see [Granting the required permissions in Microsoft Exchange Server](#).
- The Outlook Cached Exchange Mode option under the Outlook profile associated with Rapid Recovery is disabled.
- The Core machine is in the same domain as the Exchange database.
- The required Exchange database is open and you are on the Mail Restore page of the Rapid Recovery Core Console. For more information, see [Opening an Exchange database in Rapid Recovery Core](#).

Without the proper permissions and an instance of Outlook installed on the Core machine, it is not possible to perform Exchange-item recovery, even if the Exchange server is protected by a Rapid Recovery Core.

Granting the required permissions in Microsoft Exchange Server

Certain permissions must be set on Microsoft Exchange Server to complete a recovery with Rapid Recovery. For example, add the Administrator role to the mailbox, and then grant full access permission to the mailboxes you want to restore. The procedures for these permissions are specific to the version of Exchange installed on the server.

Rapid Recovery requires that you have full access permissions in Exchange to complete a recovery. For Exchange 2016, 2013, and 2010, you need to add a mailbox as a member of a role group and then grant full access permissions for the mailbox.

For more information about setting permissions for Exchange, refer to the following topics found on the Microsoft Web site <http://www.technet.com>, which contains a knowledge base of helpful procedures and topics that pertain to Microsoft Exchange:

NOTE: Use the drop-down list below the topic title to select your version of Exchange, when applicable.

- **Add-MailboxPermission.** All Exchange versions. Shows how to use the Add-MailboxPermission cmdlet to grant Full Access permissions for a mailbox.

NOTE: You must have the proper role permissions to complete these procedures. Refer to the topic, Role Management Permissions on <http://www.technet.com> for the required permissions.

- **Permissions.** Exchange 2016, 2013, and 2010 only. Provides an overview of permissions topics for the selected version of Exchange.
- **Manage Role Group Members.** Exchange 2016 and 2013 only. Provides instructions for using the Exchange Administration Center (EAC) to add members to a role group, and using the Shell to add members to a role group.
- **Add Members to a Role Group.** Exchange 2010 only. Provides instructions for using the Exchange Administration Center (EAC) to add members to a role group, and using the Shell to add members to a role group.

- **Manage Full Access Permissions.** Exchange 2010 and 2007 only. Provides instructions for how to use the Exchange Management Console (EMC) or the Shell to manage Full Access permissions for a mailbox.
i **NOTE:** You must have the proper role permissions to complete these procedures. Refer to the topic, Role Management Permissions, and then Role Groups on <http://www.technet.com> to determine which permissions are required.
- **Allow Mailbox Access.** Exchange 2010 and 2007 only. Explains how to use the Exchange Management Console to grant the Full Access permission for a mailbox.
i **NOTE:** Any administrator delegated with Exchange administrator permissions must be a member of the local administrators group. Microsoft does not recommend delegating local administrator permissions to Exchange Recipient or Exchange View Only administrator roles.
- **How to Add a User or Group to an Administrator Role.** Exchange 2007 only. Explains how to use the Exchange Management Console (EMC) to add a user or group to an administrator role.

For more information, see the following topics in the *Mailbox Restore for Exchange 6.3 User Guide*.

- Mailbox Restore system requirements
- Microsoft Outlook criteria

Opening an Exchange database in Rapid Recovery Core

Before you begin this task, assure that the following prerequisites are complete:

- Outlook 2007 or later is installed on the Core machine.
- The Core machine is in the same domain as the Exchange database.
- The proper permissions are set in Exchange. For more information, see [Granting the required permissions in Microsoft Exchange Server](#).

Rapid Recovery lets you restore mail items without leaving the interface. The mail items exist in the Exchange database within the recovery point of a protected Exchange Server machine, which you can open using the Open Exchange Databases Wizard.

1. From the Rapid Recovery Core Console, click the **...** (More) menu, and then click **Mail Restore**.
2. On the Mail Restore page, to access the Exchange database where the mail item is stored, click **Open Database**.
The Open Exchange Databases Wizard opens.
3. On the Location page of the wizard, you can open a database from a protected machine or from a local path such as your current machine or a file share. Do one of the following:
 - *Open from protected machine:* Click **Next**, and then continue to the next step in this task.
 - *Open from local path:* Enter the following information for the location, and then click **Finish**:
 - Database file path
 - Logs path
 - System path

4. On the Machines page, select the protected machine that houses the Exchange database, and then click **Next**.
5. On the Recovery Points page, select the recovery point for the point in time from which you want to open the database, and then click **Next**.
6. On the Databases page, select the Exchange database that you want to open, and then click **Finish**. Rapid Recovery opens the selected database and displays it on the Mail Restore page, with the mailboxes and folders listed on the left in an expandable navigation tree. Items within folders display on the right.

i | **NOTE:** The amount of time it takes for Rapid Recovery to open the Exchange database depends on the size of the database.

To restore an item from the open database, see [Restoring a mail item in Rapid Recovery](#).

Restoring a mail item in Rapid Recovery

Before you begin this task, assure that you have met the prerequisites for completing a mail restore. For more information, see [Mail Restore prerequisites](#).

The Rapid Recovery Mail Restore feature lets you restore a mailbox, folder, or item — such as a message, calendar event, or contact — from the data store of a protected Exchange Server machine. You can restore your selection to a recovery folder, to the original source, or to one or more PST files. To restore a mail item from the Rapid Recovery Core Console, complete the following steps.

1. From the open Exchange database on the Mail Restore page, select the item that you want to recover and then, on the Mail Restore actions bar, click **Restore**. The Email Restore Wizard opens.
2. On the Restore Session page, complete one of the following options, and then click **Next**.
 - If you are restoring mail items for the first time, enter a display name and your Outlook credentials for the restore session. You can then select this session for a future restore.
 - If you have previously created restore sessions, select one of the following options:
 - Select **Use an existing restore session**, and then select a session from the drop-down list.
 - Select **Create a new restore session**, and then enter a display name and your Outlook credentials for the restore session.

3. On the Destination page, select the target location of the restored item from the following options, and then click **Next**:

Table 140: Mail restore destinations

Option	Description
Restore to the recovery folder	Recovers the selected items (including the folder hierarchy) to a recovery folder in an online mailbox of your choice. Go to step 4 .
Restore to the original location	Directs the selected item (including the folder hierarchy) to the email box in the online data store in which it originally resided. Go to step 5 .
Restore to the PST file	Saves the selected items (including the folder hierarchy) by creating a Personal Storage Table (PST) file or writing to an existing PST file. Go to step 6 .
Restore to the PST file(s) (separate file for each mailbox)	Saves each mailbox as a Personal Storage Table (PST) file. Go to step 6 .

4. If you selected **Restore to the recovery folder**, on the Configuration page, select a **Profile** from the drop-down list, browse for and select the Outlook address book, and then go to [step 7](#).
Optionally, select **Show advanced options**, to further customize the restore with the following options:

Table 141: Advanced mail restore options

Option	Description
Error handling	Determines the way to address and manage any errors that may occur during the restore process. Select one of the following options: <ul style="list-style-type: none">• Log and continue. Collects error messages in a log and continues with the restore process.• Notify user. Pauses the restore and displays a message in the Monitor Active Task dialog when it encounters an error, and gives you the option to continue with or cancel the restore.• Abort restore. Ends the restore process when an error occurs.
Restore deleted objects	For an Exchange 2010, 2013, and 2016 database, restores items that were permanently deleted. For an Exchange 2007 database, restores strikethrough items from the current folder.
Restore email rules	Restores any rules the user had in place at the time that the data was backed up.


5. If you selected **Restore to the original location**, on the Configuration page, select the target Outlook **Profile**, select a **Restore type** from the following options, and then go to [step 7](#):
 - **Restore only differences.** Identifies whether the item being restored is already present in the destination folder and completes the restore only if there is no duplicate item. Also known as a differential restore.
 - **Create duplicate entries.** Restores every item selected without overwriting existing items, resulting in duplicates of the previously existing items.
 - **Overwrite if more recent.** Restores newer items that are not present in the online data store. It also restores items that are present in the online data store but are older than the items in the copy of the Exchange database.

Optionally, select **Show advanced options**, to further customize the restore with the following options:

Table 142: Advanced mail restore options

Option	Description
Error handling	<p>Determines the way to address and manage any errors that may occur during the restore process. Select one of the following options:</p> <ul style="list-style-type: none"> • Log and continue. Collects error messages in a log and continues with the restore process. • Notify user. Pauses the restore and displays a message in the Monitor Active Task dialog when it encounters an error, and gives you the option to continue with or cancel the restore. • Abort restore. Ends the restore process when an error occurs.
Restore deleted objects	<p>For an Exchange 2010, 2013, and 2016 database, restores items that were permanently deleted.</p> <p>For an Exchange 2007 database, restores strikethrough items from the current folder.</p>
Restore email rules	Restores any rules the user had in place at the time that the data was backed up.
Restore user permissions	<p>Restores the custom permissions set for a public folder.</p> <p>i NOTE: This option is only available when you restore a public folder to its original location. If you do not select to restore permissions, then the default folder permissions are restored with the content.</p>

6. If you selected **Restore to the PST file** or **Restore to the PST file(s) (separate file for each mailbox)**, on the Configuration page, complete the following selections, and then go to [step 7](#):
 - a. **Profile**. Select an Outlook profile from the drop-down list.
 - b. **Primary PST storage**. To locate and select the initial destination folder for the PST file, enter the path or select an existing file.
 - c. **Overflow PST storage (optional)**. If the primary destination has insufficient space for the PST file, select a secondary destination for the PST file.

 **NOTE:** Do not assign the overflow location to the same disk as the primary location.

Optionally, select **Show advanced options**, to further customize the restore with the following options:

Table 143: Advanced mail restore options

Option	Description
Error handling	<p>Determines the way to address and manage any errors that may occur during the restore process. Select one of the following options:</p> <ul style="list-style-type: none">• Log and continue. Collects error messages in a log and continues with the restore process.• Notify user. Pauses the restore and displays a message in the Monitor Active Task dialog when it encounters an error, and gives you the option to continue with or cancel the restore.• Abort restore. Ends the restore process when an error occurs.
Restore deleted objects	<p>For an Exchange 2010, 2013, and 2016 database, restores items that were permanently deleted.</p> <p>For an Exchange 2007 database, restores strikethrough items from the current folder.</p>

7. Click **Finish**.

The items restore to your selected destination. You can monitor the progress of the job in on the *Events* page.

Bare metal restore

When operating as expected, servers perform the tasks for which they are configured. If a server protected in your Rapid Recovery Core suffers a catastrophic failure that renders the server inoperable, administrators must take immediate action to restore the full functionality of that machine.

In such cases, especially when the data loss includes a system volume, you can use Rapid Recovery to perform a *bare metal restore* (BMR) for your protected machines. BMR is a process that restores the full software configuration for a specific system. It uses the term “bare metal” because the restore operation recovers not only the data from the server, but also reformats the hard drive and reinstalls the operating system and all software applications.

Rapid Recovery Core lets you perform bare metal restore for protected Windows or Linux machines. The protected system can be restored to similar or dissimilar hardware.

As of release 6.8, in addition to BMR of volume-level backups, you can also perform BMR for disk-level backups of Linux machines.

This section describes how to restore a recovery point from a protected machine to bare metal using similar or dissimilar hardware. Most of the tasks for a BMR are performed from the Restore Machine Wizard. When restoring a Linux machine to bare metal, you can also accomplish several tasks from the command line. These procedures are also included in this section.

Topics include:

- [About bare metal restore](#)
- [Differences in bare metal restore for Windows and Linux machines](#)
- [Managing a Windows boot image](#)
- [Managing a Linux boot image](#)
- [Performing a bare metal restore using the Restore Machine Wizard](#)
- [Using the Universal Recovery Console for a BMR](#)
- [Performing a bare metal restore for Linux machines](#)
- [Verifying a bare metal restore](#)

About bare metal restore

Bare metal restore is the process of restoring **all** content from a specific computer system — data, applications, user accounts, and the operating system — from a recovery point.

i | **NOTE:** Before performing bare metal restore, ensure you have a healthy hardware system with which to replace the failed system.

BMR is used not only in disaster recovery scenarios, but also to migrate data when upgrading or replacing servers. To perform a BMR, Rapid Recovery uses an ISO image as a boot disk, which lets you connect the BMR target machine with the Rapid Recovery Core using a restore interface called the Universal Recovery Console.

NOTE: An ISO image is a single archive that contains data for every sector of a disk, including the disk file system. ISO images are saved in ISO-9660 format, set by the International Organization for Standardization (ISO). The file format uses the `.iso` file extension.

For more flexibility, Rapid Recovery supports BMR both to similar and dissimilar hardware. Examples of restoring to similar hardware include replacing the hard drive only of the existing system, or swapping out a failed server with an identical machine. An example of restoring to dissimilar hardware is the case where you replace the failed system with a server produced by a different manufacturer or with a different configuration.

The process of performing a BMR includes several separate procedures. At the top level, it involves creating or downloading a bootable ISO image; transmitting that image to an accessible location (removable media, network location, or hypervisor); starting up the BMR target server from the boot image; connecting it to the recovery console instance; mapping volumes; initiating the recovery; and then monitoring the restore progress. Once the bare metal restore is complete, you can continue with the task of loading the operating system and the software applications on the restored server, followed by establishing unique settings required for your configuration.

NOTE: Bare metal restore is supported for virtual machines (VMs) as well as for physical machines. However, to be practical, if the machine you want to replace is a VM, it is generally quicker and easier to perform virtual export from a recovery point to achieve the same goal. For more information on performing a VM export, see [Exporting to virtual machines using Rapid Recovery](#).

Rapid Recovery supports BMR for both Windows and Linux machines using the Restore Machine wizard from the Rapid Recovery Core Console. Some steps differ. For a list of general BMR steps, differentiated for Windows and Linux restores, see [Differences in bare metal restore for Windows and Linux machines](#).

For Linux, you can also accomplish many tasks required for BMR from the command line. If that is your preference, both approaches are included in this User Guide when applicable.

This section includes conceptual topics throughout, including prerequisites and information about boot ISO images used for BMR for Windows and Linux machines. Before performing any BMR, consider requirements as described in [Prerequisites for performing a bare metal restore for Windows or Linux machines](#).

An ISO image is a single archive that contains data for every sector of a disk, including the disk file system. The topic [Performing a bare metal restore using the Restore Machine Wizard](#) describes the process to start a BMR from the wizard in the Core Console. Following that process, and referring to other topics for additional information, users can perform a BMR from the Restore Machine Wizard for both Windows and Linux machines.

Differences in bare metal restore for Windows and Linux machines

The main tasks for performing BMR are described in the following table. Differences are noted between the process for Windows and Linux bare metal restores.

Step	General BMR Steps	Windows	Linux
1	Prepare the destination machine. Repair or prepare hardware to replace the failed system. The fixed or replaced system is referred to as the BMR target machine.		
2	Specify a recovery point. When restoring to bare metal, you must select the recovery	Select the appropriate recovery point from the Rapid Recovery Core	You have two options:

Step	General BMR Steps	Windows	Linux
	<p>point from which to restore all data, the OS, and applications. Often this is the most recent snapshot. However, there are cases in which you want to select an earlier recovery point (for example, if the failure was due to a recent change in software configuration).</p> <p>For restore of any machine (Windows or Linux), including BMR, you can identify the recovery point from the Rapid Recovery Core Console in two ways:</p> <ol style="list-style-type: none"> 1. Navigate to a protected machine, view its recovery points, and launch a restore from the <i>Recovery Points</i> page of the Core Console. 2. Launch the Restore Machine Wizard, select the machine, and then the recovery point, and select Restore. <p>In either case, for BMR, choose to restore from a boot image such as a Windows boot CD or the Linux Live DVD.</p>	<p>Console.</p> <p>To perform a BMR from the <i>Recovery Points</i> page of a specific machine, see About restoring volumes from a recovery point.</p> <p>See steps 2 and 3 in the task Performing a bare metal restore using the Restore Machine Wizard.</p>	<ol style="list-style-type: none"> 1. Use the Restore Machine wizard (see steps 2 and 3 in the task Performing a bare metal restore using the Restore Machine Wizard). 2. Use the command line local_mount utility. If using the command line, do so following step 3 as described in this table. For more information, see Launching a bare metal restore for a Linux machine using the command line.
3	<p>Manage the boot image. This step involves the following sub-steps:</p>	<p>For descriptive information on managing a boot CD image for a Windows machine, see Managing a Windows boot image.</p>	<p>For more descriptive information about managing the boot image for Linux, see Managing a Linux boot image.</p>
3a	<ul style="list-style-type: none"> • Injecting missing drivers. If additional Ethernet controller, storage, network adapter, or other drivers are needed, inject drivers into the boot image. <p>i NOTE: The process of injecting drivers is particularly relevant when restoring to dissimilar Windows hardware.</p>	<p>Windows machines may require driver injection, which can be performed on the <i>Driver Injection</i> page of the Restore Machine Wizard. See step 11 in the task Performing a bare metal restore using the Restore Machine Wizard.</p>	<p>The Live DVD has a variety of necessary drivers. Driver injection is typically not required for Linux machines.</p>
3b	<ul style="list-style-type: none"> • Obtaining a bootable ISO image. Define and create the image for Windows, or download the image for Linux bare metal restores. 	<p>Define the requirements for the boot ISO image from</p>	<p>Download the Linux Live DVD boot ISO image for the appropriate version of Rapid</p>

Step	General BMR Steps	Windows	Linux
		the Restore Machine Wizard. Define an export path and generate the boot CD to a location you specify. See steps 6 to 10 in the task Performing a bare metal restore using the Restore Machine Wizard .	Recovery Core from the <i>Downloads</i> page on the QorePortal or the Rapid Recovery License Portal.
3c	<ul style="list-style-type: none"> • Transmitting the image to an accessible location. Place the boot disk image in a location from which the BMR target machine can boot from it. Options include: <ul style="list-style-type: none"> • Transferring the ISO image to physical storage media (for example, a CD, DVD, or a bootable USB flash drive) and moving it to the physical BMR target server. • Storing the ISO image on a network location visible to the BMR target machine. 		
3d	<ul style="list-style-type: none"> • Mounting the ISO image. Boot the BMR target machine from the ISO image. <ul style="list-style-type: none"> • If restoring a hypervisor guest, mount the boot media on the hypervisor. • If restoring a physical machine, access the boot menu options, point to the ISO boot image, and reboot. 		
4	Obtain URC login credentials. On the BMR target machine, launch the Universal Recovery Console (URC) user interface, and capture the session-unique IP address and authentication key credentials provided in the URC.		If the machine is not able to find the IP address, you may be required to manually configure it.
5	Connect the BMR target machine to the Core Console. From the Core Console,	See step 14 in the task Performing a bare	

Step	General BMR Steps	Windows	Linux
	connect to the BMR target machine using the credentials from the URC.	metal restore using the Restore Machine Wizard.	
6	Map volumes to be restored. From the Core Console, map volumes to be created during the restore process.	See step 16 in the task Performing a bare metal restore using the Restore Machine Wizard.	If using complex LVM or RAID volumes, use the automatic mapping feature in the Restore Machine Wizard if you want to automatically create the appropriate partitions. Otherwise, the partitions must be created manually before you begin the restore process. If mapping volumes manually, you must first ensure the number and size of volumes from the recovery point matches the number and size of volumes on the machine to which you are restoring data. To manually create partitions, see Creating partitions on the destination drive.
7	Begin restoring data. The next step is to begin the actual restore process.	From the Core Console, start the restore process.	Starting the restore process can be performed both from the Restore Machine Wizard in the Core console and from the command line. For more information on using the command line, see Launching a bare metal restore for a Linux machine using the command line.
8	Monitor the restore process. You can monitor the progress on the <i>Events</i> page. For more information, see Viewing events using tasks, alerts, and journal pages.		
9	Verify the restore process. During the restore process, and when it completes, you can verify the restore process.	You can verify from the Core Console. See Verifying a bare metal restore.	You can verify from the Core Console or from the command line. For this second option, see Verifying the bare metal restore from the command line.
10	Complete your custom configuration.		

Step	General BMR Steps	Windows	Linux
	Optionally, establish any unique settings required for your new configuration.		

Prerequisites for performing a bare metal restore for Windows or Linux machines

Before you can begin the process of performing a bare metal restore for a Windows machine, you must ensure that the following conditions and criteria exist:

- **Backups of the machine you want to restore.** You must have a functioning Rapid Recovery Core containing recovery points of the protected server you want to restore.
- **Hardware to restore (new or old, similar or dissimilar).** The BMR target machine must meet the installation requirements for a protected machine. For details, see the *Rapid Recovery Installation and Upgrade Guide*. At minimum, the BMR target must have a 64-bit central processing unit (CPU). The Windows boot CD created by the Core uses the Windows Preinstallation Environment (Win PE) 10 operating system. Rapid Recovery BMRs are not compatible with x86-based CPUs. You can only perform a BMR on a 64-bit CPU.

i **NOTE:** Prior to generating a boot CD, you must install Windows Preinstallation Environment (WinPE). If not installed, you are prompted to download the [Windows Assessment and Deployment Kit](#) and the [Windows Assessment and Deployment Kit Add-Ons](#) from the Microsoft website.
- **Compatible storage drivers and network adapter drivers.** If restoring to dissimilar hardware, then you must have compatible storage drivers and network adapter drivers for the target machine, including RAID, AHCI, and chipset drivers, as appropriate. These should be injected to the boot ISO image.
- **Bootable ISO image.** You must have an ISO image from which to boot the BMR target machine. The steps are different for Windows or Linux BMR targets:
 - **Windows: Generate a Boot CD.** You must generate a boot CD ISO image from the Rapid Recovery Core from which to boot the BMR target machine. The boot CD contains the Rapid Recovery Universal Recovery Console UI. From the URC, you can connect the BMR target to the Core to perform the restore. If restoring from a network location, you may be able to boot directly from the ISO image.
 - **Linux: Live DVD boot image.** Obtain the Linux Live DVD ISO image, which includes a bootable version of Ubuntu Linux. You can download it from the QorePortal at <https://qoreportal.quest.com/> or from the Rapid Recovery License Portal at <https://rapidrecovery.licenseportal.com>. If you have any issues downloading the Live DVD, contact Quest Data Protection Support. When booted from this ISO image, a simplified version of the Universal Recovery Console appears.
- **Image media and software.** If you cannot boot directly from the ISO image, move the image to physical storage media such as a bootable USB flash drive, CD, or DVD. This process requires blank storage media and disk burning software, or software to create an ISO image or make one bootable. For BMR on Windows, if managing machines remotely using virtual network computing software such as UltraVNC, then you must have VNC Viewer.

- **Storage space and partitions, as appropriate.** Ensure that there is enough space on the hard drive of the BMR target to create destination partitions to contain the volumes you want to restore from the recovery point source volumes. Any destination partition should be at least as large as the original source partition. To perform a BMR, you must restore at least the system volume.



i **NOTE:** If recovering LVM volumes and RAID volumes on Linux machines (including LVMs and RAID with partitions, and complex LVMs and RAID), you must either specify automatic volume mapping, or create the partitions prior to starting the BMR.

- **Windows: compatible partitions.** Windows 8.1, Windows 10, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019 operating systems that are booted from FAT32 EFI partitions are available for protection or recovery, as well as are Resilient File System (ReFS) volumes. UEFI partitions are treated as simple FAT32 volumes. Incremental transfers are fully supported and protected. Rapid Recovery provides support of UEFI systems for BMR including automatic partitioning GPT disks, LVM and software-based RAID volumes.
- **Linux: Restore path.** Identify the path for the restore, which is the path for the device file descriptor. To identify the path for the device file descriptor, use the `fdisk` command from a terminal window.

Managing a Windows boot image

A bare metal restore for Windows requires a boot CD, which you create from the Rapid Recovery Core Console. This is a bootable ISO image which contains the Universal Recovery Console (URC) interface, an environment that is used to restore the system drive or the entire server by connecting the target machine with the Rapid Recovery Core.

i **NOTE:** BMR for a Linux protected machine requires an ISO image called the Live DVD. For information about the Live DVD, see [Managing a Linux boot image](#).

You can create a Windows boot CD using one of two methods: by defining parameters in the Rapid Recovery Core using the Restore Machine Wizard, or from the  [Boot CDs](#) page of the Core Console accessible under the  (More) menu.

Tailor each boot CD ISO image to your specific needs. For example, each boot CD must contain the correct Ethernet network and mass storage drivers. If you anticipate that you will be restoring to different hardware from the machine on which the recovery point originated, then you must include storage controller and other drivers in the boot CD. For information about injecting those drivers in the boot CD, see [Understanding driver injection in a boot CD](#).

If you plan to connect to the BMR target remotely, you cannot use RDS, but you can use UltraVNC. If using UltraVNC, you will need to provide the UltraVNC password defined when creating the boot CD to access the target machine. For more information, see [Using UltraVNC for remote access](#).

After it is created, you will use the boot CD image to boot your Windows-based BMR target machine when restoring to bare metal. To boot from it, based on the specifics of your environment, you may need to transfer the image to physical media. Then virtually or physically load the boot image, and boot the Windows server from the boot image. You will then start the URC, note the single-session authentication information provided, and use that information on the Rapid Recovery Core to connect the target machine to the Core to perform the restore process.

This section includes the following relevant topics:

- [Understanding driver injection in a boot CD](#)
- [Using UltraVNC for remote access](#)
- [Creating a boot CD ISO image](#)

- [Transferring the boot CD ISO image to media](#)
- [Loading the boot CD and starting the target machine](#)

Understanding driver injection in a boot CD

The boot CD image requires storage drivers to recognize the drives of the server, and network adapter drivers in order to communicate with the Rapid Recovery Core over the network.

To suit this purpose, a generic set of Windows10 x64 storage controller and network adapter drivers are included automatically when you generate a boot CD for Windows. These generic drivers satisfy the requirements of many newer systems. Data restored from the recovery point also includes drivers from the hardware previously in place. If restoring to the same or similar hardware, the included drivers or restored drivers may be sufficient.

Creation of a successful boot CD can be a trial and error effort. When creating the boot CD, you can use driver injection to facilitate interoperability between the recovery console, network adapter, and storage on the target server. If performing a BMR to dissimilar hardware, or if restoring an older system, you may need to inject storage controller or network adapter drivers when creating the boot CD. If you discover the boot CD you created does not contain the drivers necessary to complete the restore, you can also load drivers on to the target machine using the URC. After successfully completing the restore process and booting the OS, you can download and install any additional drivers needed by the OS to interact with its new hardware. For more information, see [Loading drivers using the Universal Recovery Console](#).

Using UltraVNC for remote access

You cannot log onto the BMR target remotely using Remote Desktop Services while using the boot CD. To remotely access the BMR target and interact with the Universal Recovery Console, you can use UltraVNC (a third-party tool). UltraVNC is an open-source utility that lets users access a computer remotely as if they were in front of it. Use of this product with Rapid Recovery to perform BMR requires two components, VNC Server and VNC Viewer. If you need to remotely access the BMR target machine, then UltraVNC Server can be added to the boot CD when you create it in the Rapid Recovery Core Console.

Before you can select the option to embed UltraVNC into the boot CD, you must have a qualifying version of UltraVNC installed on the Core machine.

For this release, supported VNC versions include 1.1.8.9, 1.0.9.6, 1.0.9.6.1, 1.0.9.6.2, and 1.1.9.3. Install both the UltraVNC Server and UltraVNC Viewer components.

You can visit <https://www.uvnc.com/downloads/ultravnc/> to download UltraVNC for x64 architecture.

After installing, copy the executable file for the server component only (winvnc.exe) to the following path:

`C:\Program Files\AppRecovery\Core\BootCdKit\Tools\UltraVnc_x64.`

i **NOTE:** Do not move the executable from its original installation location. Place a duplicate copy in the specified directory.

Creating a boot CD ISO image

A boot CD is the term Rapid Recovery uses to refer to the ISO image reserved for performing a bare metal restore (BMR) for Windows machines. The image includes the Rapid Recovery Universal Recovery Console (URC).

To perform a BMR on a machine, you must start the machine from the boot CD, which launches the URC. The URC enables you to connect the BMR target to the recovery point you want to use to complete the restore.

You cannot log onto the BMR target machine using Remote Desktop Services while using the boot CD. If you need to connect remotely to the BMR target machine, see [Using UltraVNC for remote access](#).

NOTE: You can also create a boot CD image from the Restore Machine Wizard workflow. When you select **Recover to any target machine using a Boot CD**, the next steps in the wizard let you create a boot CD. If you create a boot CD from that wizard, this task is not required.

Use this procedure to create a boot CD for a Windows machine independent of the Restore Machine Wizard.

1. From the Rapid Recovery Core Console where the server you need to restore is protected, in the icon bar, click the **⋮** (More) menu, and then click **⚙️ Boot CDs**.
2. On the *Boot CDs* page, click **+ Create Boot CD**.
The *Create Boot CD* dialog box displays.
3. In the *Create Boot CD* dialog box, in the **Output path** text box, enter the path in which you want to store the boot CD ISO image. For example, to create a boot CD for a protected machine named *Protected Machine 1*, enter the path `C:\ProgramData\AppRecovery\Boot CDs\ProtectedMachine1-BootCD.iso`.

NOTE: If you want to store the image on a shared drive which has insufficient disk space, you can create a disk as needed in the path; for example, `F:\filename.iso`.

NOTE: The file extension must be `.iso`. When specifying the path, use only alphanumeric characters, the hyphen, the backslash (only as a path delimiter), and the period (only to separate host names and domains). The letters a to z are not case sensitive. Do not use spaces in the name of the boot CD image. No other symbols or punctuation characters are permitted.

4. Under *Connection Options*, do one of the following:
 - To obtain the IP address dynamically using Dynamic Host Configuration Protocol (DHCP), select **Obtain IP address automatically**.
 - To specify a static IP address for connecting to the URC on the BMR target, select **Use the following IP address**, and then enter the information described in the following table.

Table 144: Network connection options

Option	Description
IP address	The IP address for the restored machine.
Subnet mask	The subnet mask for the restored machine.
Default gateway	Specify the default gateway for the restored machine.
DNS server	Specify the domain name server for the restored machine.

NOTE: You must specify all four of these parameters.

5. UltraVNC is a third-party utility that lets you manage the URC remotely. If you require remote access to the recovery console, and you have UltraVNC installed, under **UltraVNC Options**, select **Add UltraVNC**, and then enter the information described in the following table.



Table 145: UltraVNC connection credentials

Option	Description
UltraVNC password	Define a password. You must enter this password to connect from the boot CD using UltraVNC.
UltraVNC port	The port you want to use to connect to the BMR target. The default port is 5900.

6. If you plan to restore to dissimilar hardware, inject the appropriate storage controller and other drivers for your target system by completing the following steps:
 - a. Download the drivers from the server manufacturer's website and unpack them.
 - b. Compress each driver into a `.zip` file using an appropriate compression utility (for example, WinZip).
 - c. In the *Create Boot CD* dialog box, in the Drivers pane, click **Add an Archive of Drivers**. The *Choose File to Upload* dialog box appears.
 - d. From the dialog box, navigate through the filing system to locate the compressed driver file. Select the file, and then click **Open**. The *Choose File to Upload* dialog box closes. The driver file you selected appears in the Drivers pane of the *Create Boot CD* dialog box.
 - e. Repeat step c and step d, as appropriate, until you add all necessary drivers.
 - f. In the Drivers pane, select the drivers that you want to inject.

i **NOTE:** Automatic driver injection is supported by Windows 8.1 and higher, and Windows Server 2012 R2 and higher. If creating a boot CD ISO image for earlier versions of Windows, manually save drivers to `C:\Program Files\AppRecovery\Core\BootCdKit\Drivers\`.

For more information about injecting drivers, see [Understanding driver injection in a boot CD](#).

7. Click **Create Boot CD**.
Rapid Recovery creates the boot CD and saves it with the file name you provided.
8. To monitor the progress of this task, go to the icon bar and click the  Events icon.
For more information about monitoring Rapid Recovery events, see [Viewing events using tasks, alerts, and journal pages](#).
When the ISO image creation is complete, a record of the image appears on the *Boot CDs* page, which you can access from the  (More) menu in the icon bar.

To access the ISO image, navigate to the output path you specified or click the link on the *Boot CDs* page. Then save the image to a location from which you can then load it onto the new system, such as a network drive.

Transferring the boot CD ISO image to media

When you create the boot CD file, it is stored as an ISO image in the path you specified. You must be able to mount this image as a drive on the server on which you are performing a bare metal restore.

You can burn the boot CD ISO image onto compact disc (CD) or digital video disk (DVD) media accessible at system startup.


When you start the machine from the boot CD, the Universal Recovery Console launches automatically.


If performing a BMR on a virtual machine, this step is not required. Simply load the ISO image in a drive and edit settings for that VM to start from that drive.


Loading the boot CD and starting the target machine

After you create the boot CD image, you need to boot the target server using that image.

To connect the BMR target machine to the Rapid Recovery Core Console or to use Chromium for downloading additional drivers, you must first load an Ethernet controller and network adapter. For more information, see [Loading drivers using the Universal Recovery Console](#).

1. On the BMR target machine, load the boot CD image from the appropriate location, and then start the server from the boot CD image to load Win PE 10 and the Universal Recovery Console (URC) environment. The target machine displays a blue Quest screen with three URC function buttons at the top of the screen.
2. To start the URC user interface, from the buttons at the top of the screen, click  (Start URC). The URC splash screen appears, and you are prompted to choose a display language.
3. From the language drop-down menu, select a display language and click **OK**. The Universal Recovery Console appears. The URC applies network settings, starts Rapid Recovery Agent, and searches the boot CD image for available drivers.
4. On the right side of the console, under **Authentication**, if the IP address does not appear, the target machine cannot find the appropriate network adapter. Stop this task, and perform one of the tasks described in the topic [Loading drivers using the Universal Recovery Console](#).
5. On the right side of the console, under **Authentication**, if the IP address of the target machine populates, you also see a single-use password.

 **NOTE:** If you specified an IP address in the *Create Boot CD* dialog box, the Universal Recovery Console displays that IP address in the Authentication area.

6. If you want to change the IP address, do the following:
 - a. To the right of the IP address, click  (Change). The *Change network adapter settings* dialog box appears.
 - b. Enter the appropriate IP address, completing values in IP address, gateway, subnet mask, and DNS server areas as appropriate.
 - c. When satisfied, click **OK**.

The *Change network adapter settings* dialog box closes, and the new IP address appears in the Authentication area.

7. If you want to keep this IP address, record the authentication information. You will use it to connect the URC to the Rapid Recovery Core Console.

The machine is ready for you to connect to the Core, select a recovery point, and continue the bare metal restore process.

Managing a Linux boot image

A bare metal restore for Linux requires a Live DVD boot image, which you download from an appropriate download location, such as the QorePortal or the Rapid Recovery License Portal. You will use this image to boot the destination Linux machine and connect to the Rapid Recovery Core to restore from a specified recovery point. Based on the specifics of your environment, you may need to transfer this image to physical media. You must then

virtually or physically load the boot image, and start the Linux server from the boot image.

i | **NOTE:** The Live DVD was previously known as the Live CD.

Managing a Linux boot image is a step in [Performing a bare metal restore for Linux machines](#).

This section includes the following relevant topics:

- [About the boot ISO image for Linux](#)
- [Downloading a boot ISO image for Linux](#)
- [Saving the Live DVD ISO image to media](#)
- [Loading the Live DVD and starting the target machine](#)
- [Connecting to the BMR target from the Rapid Recovery Core](#)

About the boot ISO image for Linux


The first step when performing a bare metal restore (BMR) for a Linux machine is to download the Linux Live DVD ISO image from the *Downloads* page of either the QorePortal or the Rapid Recovery License Portal. The Live DVD functions with all Linux file systems supported by Rapid Recovery, and includes a bootable version of Ubuntu Linux, the GNU Screen utility (a terminal multiplexer), and a limited version of the Rapid Recovery Universal Recovery Console (URC) interface. The Rapid Recovery Universal Recovery Console is an environment that is used to restore the system drive or the entire server directly from the Rapid Recovery Core.

Downloading a boot ISO image for Linux

To complete a bare-metal restore for a Linux machine, you need a Live DVD ISO image that matches your version of Rapid Recovery Core. The current version of Live DVD is available from the *Downloads* page from both the QorePortal and the Rapid Recovery License Portal. If you need a different version, contact Quest Data Protection Support.

This task is a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Managing a Linux boot image](#).

To download the Live DVD ISO image, complete the steps in this procedure.

1. In a web browser, open the QorePortal at <https://qoreportal.quest.com/>.
2. Click **Settings**, and then click  **Downloads**.
3. Under *Utilities*, locate **Linux Live DVD**, and then click **Download**.
The Linux boot ISO image, for example `rapidrecovery-livedvd-6.3.x.iso`, saves to the downloads destination folder.
4. If restoring from a physical machine, transfer the Live DVD ISO image onto physical media. For more information, see [Saving the Live DVD ISO image to media](#).
5. Optionally, if restoring a Linux virtual machine, you can save the ISO image to a network location, and then edit the VM settings to boot from it. Or you can transfer the Live DVD ISO image onto physical media, and change your VM settings to boot from a DVD or CD drive containing that physical media.

Saving the Live DVD ISO image to media

When you download the Linux Live DVD file, it is stored as an ISO image in the path you specified. You must be able to boot the target Linux machine from the Live DVD image.

This task is a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Managing a Linux boot image](#).

1. Copy the boot CD ISO image onto removable media supported by your Linux machine; for example, a bootable USB flash drive, digital video disk (DVD), or compact disc (CD).
2. If necessary, take the proper steps to ensure that the media with the Live DVD ISO image is bootable. For example, change the boot order of devices to ensure the volume or drive containing the Live DVD boot image.

If performing a BMR on a virtual machine, this task is not required. Simply edit the VM settings to boot from the volume or drive containing the Rapid Recovery Live DVD ISO image instead of from physical media.

You can also use virtual export to restore a Linux VM. For more information, see [VM export](#).

Loading the Live DVD and starting the target machine

After you obtain the Live DVD ISO image, you need to start the Linux machine from the Live DVD ISO image that matches your version of Rapid Recovery Core.

This task is a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Managing a Linux boot image](#).

1. Navigate to the BMR target machine and load the Live DVD image from the appropriate location.
2. If necessary, access settings for the BMR target machine, and modify the boot order so that the server will start from the device containing the Live DVD image.
3. Boot the machine.
A Quest splash screen displays. A terminal window opens, displaying the IP address and authentication password for the machine.

i | **NOTE:** A new temporary password is generated each time the machine is started with the Live DVD image.

4. Write down the IP address and the authentication password displayed on the introduction screen. You will need this information to connect the BMR target machine with the Rapid Recovery Core Console to complete the bare metal restore.

Connecting to the BMR target from the Rapid Recovery Core


After you boot the target Linux machine from the Live DVD ISO image, this machine is ready for you to connect to it from the Core and begin the bare metal restore process. You can perform this process using either of two methods:

- Launching a restore using the Restore Machine wizard accessed from the Rapid Recovery Core Console. For more information, see [Launching a bare metal restore for Linux](#).
- Launching a Restore from the command Line using the local_mount utility. For more information, see [Launching a bare metal restore for a Linux machine using the command line](#).

Performing a bare metal restore using the Restore Machine Wizard

You can use the Restore Wizard to perform a bare metal restore for Windows or Linux machines.

i NOTE: A bootable ISO image is required to complete a BMR. For Windows restores, you need to create a boot CD image from the Core Console. For Linux restores, you need to download the Live DVD boot ISO image for the appropriate version of Rapid Recovery Core from the QorePortal or Rapid Recovery License Portal.


This procedure includes a description of creating and managing a Windows boot CD image from the Restore Machine Wizard. Independent of the wizard, you can also create a Windows boot CD from the *Boot CDs* page, accessible from the  (More) menu of the Core console.

For information on managing a Windows boot CD image outside of the Restore Machine Wizard, see [Managing a Windows boot image](#). For information on downloading and managing the Live DVD image, see [Managing a Linux boot image](#).

Before performing a BMR, see [Prerequisites for performing a bare metal restore for Windows or Linux machines](#).

You cannot log onto the BMR target machine using Windows Remote Desktop Services while using the boot CD. If you need to connect remotely to the BMR target machine, see [Using UltraVNC for remote access](#).

Complete the steps in the following procedure to perform BMR using the Restore Machine Wizard.

1. To restore volumes (including a system volume) on a protected machine, navigate to the Rapid Recovery Core Console and click  **Restore** from the Rapid Recovery button bar. The Restore Machine Wizard appears.
2. On the *Machines* page, select the protected machine you want to restore, and then click **Next**. The *Recovery Points* page appears.
3. Select the recovery point you want to use to restore the machine. Optionally, if you want to limit the number of recovery points displayed, you can filter by volumes (if defined) or by creation date of the recovery point. You can also conduct a search for a specific recovery point.
4. Click **Next**.
5. On the *Destination* page, select **Recover to any target machine using a Boot CD**.

6. If restoring a Windows machine, on the *Destination* page, do one of the following:
 - If you have already created a boot CD ISO image for restoring a Windows machine, load the boot CD into the BMR target machine now, reboot the target machine, and then on the Core machine, continue to [step 8](#).
For conceptual information about Windows boot CDs, see [Managing a Windows boot image](#). For information about transferring the ISO image to removable media, see [Transferring the boot CD ISO image to media](#). For information about how to load the boot CD, see [Loading the boot CD and starting the target machine](#).
 - If you have not yet created a boot CD for restoring a Windows machine, click **Next**, and then continue to [step 9](#).
7. If restoring a Linux machine, do the following:
 - a. Download the Live DVD boot ISO image, as described in [Saving the Live DVD ISO image to media](#).
 - b. Transfer the image to removable media and boot the BMR target server, as described in [Loading the Live DVD and starting the target machine](#).
 - c. On the Core machine, return to this procedure and continue with [step 15](#).
8. On the *Destination* page, under the selected **Recover to any target machine using a Boot CD** option, select **I already have a boot CD running on the target machine**, click **Next**, and then go to [step 15](#).
9. On the *Boot CD* page, in the **Output path** text box, enter the path in which you want to store the boot CD ISO image. For example, to create a boot CD for a protected machine named *Protected Machine 1*, enter the path `C:\ProgramData\AppRecovery\Boot CDs\ProtectedMachine1-BootCD.iso`.
 - i** **NOTE:** If you want to store the image on a shared drive which has insufficient disk space, you can create a disk as needed in the path; for example, `F:\filename.iso`.
 - i** **NOTE:** The file extension must be `.iso`. When specifying the path, use only alphanumeric characters, the hyphen, the backslash (only as a path delimiter), and the period (only to separate host names and domains). The letters a to z are not case sensitive. Do not use spaces in the name of the boot CD image. No other symbols or punctuation characters are permitted.

10. Optionally, to set up network parameters for the target machine, or to add UltraVNC capabilities, select **Show advanced options**, and then complete the following steps:

- To specify a static IP address for connecting to the URC on the BMR target, select **Use the following IP address**, and then enter the information described in the following table.

Table 146: Network connection options

Option	Description
IP address	The IP address for the restored machine.
Subnet mask	The subnet mask for the restored machine.
Default gateway	Specify the default gateway for the restored machine.
DNS server	Specify the domain name server for the restored machine.

- If you have UltraVNC installed and would like to use it to complete the BMR remotely, select **Add UltraVNC**, and then enter the information described in the following table.

Table 147: UltraVNC connection credentials

Option	Description
Password	Define a password. You must enter this password to connect from the boot CD using UltraVNC.
Port	The port you want to use to connect to the BMR target. The default port is 5900.

i **NOTE:** UltraVNC Options are only available if you already have UltraVNC installed on your Core. For more information, see [Using UltraVNC for remote access](#).

11. When you are satisfied with your selections on the *Boot CD* page, click **Next**.

12. Optionally, on the *Driver Injection* page, especially if you plan to restore to dissimilar hardware, inject the appropriate storage controller and other drivers for your target system by completing the following steps:
 - a. Download the drivers from the server manufacturer's website and unpack them.
 - b. Compress each driver into a .zip file using an appropriate compression utility (for example, WinZip).
 - c. On the *Driver Injection* page of the Restore Machine Wizard, click **+ Add an Archive of Drivers**.
 - d. Navigate through the filing system to locate the compressed driver file, select the file, and then click **Open**.
 - e. Repeat [step c](#) and [step d](#), as appropriate, until you inject all necessary drivers.
For more information about injecting drivers, see [Understanding driver injection in a boot CD](#).

i **NOTE:** Automatic driver injection is supported by Windows 8.1 and higher, and Windows Server 2012 R2 and higher. If creating a boot CD ISO image for earlier versions of Windows, manually save drivers to C:\Program Files\AppRecovery\Core\BootCdKit\Drivers\.

- f. On the *Driver Injection* page, click **Next**.

Rapid Recovery creates the boot CD ISO image.

13. On the *ISO Image* page, click **Next**.
14. If you can start the BMR target machine from the boot CD ISO image, do so now. If you cannot boot the BMR target machine directly from the image, do the following:
 - Copy the ISO image to a physical medium (such as a bootable USB flash drive, DVD, or CD).
 - Load the storage medium in the target machine.
 - Configure the machine settings to load from the selected storage medium, and restart the BMR target machine from the boot image.

i **NOTE:** You may need to change the BIOS settings of the target machine to ensure that the volume that loads first is the boot image.

The target machine, when started from the boot image, displays the Universal Recovery Console (URC) interface. This environment is used to restore the system drive or selected volumes directly from the Rapid Recovery Core. Note the IP address and authentication key credentials in the URC, which refresh each time you start from the boot image.

15. From the Core Console, on the *Destination* or *Connection* page of the Restore Machine Wizard, enter authentication information from the URC instance of the machine you want to restore as follows:


Table 148: Authentication options

Option	Description
IP Address	The IP address provided in the URC on the target machine.
Authentication Key (or Password)	The authentication key (or password, for Linux machines) provided in the URC on the target machine.

16. Click **Next**.
A connection is made between the Core and the BMR target machine. Then, the *Disk Mapping* page of the wizard appears.

17. On the *Disk Mapping* page, if you want to map volumes manually, proceed to [step 18](#). If you want to map volumes automatically, complete the following steps:


- a. From the **Volume mapping** drop-down menu, select **Automatic**.
- b. From the list of volumes in the grid on the left, ensure that the volumes you want to restore are selected. All volumes are selected by default.
If you do not want to restore a listed volume, clear the option.

 **NOTE:** At least one volume must be selected to perform the restore.

- c. From the list of disks in the grid on the right, select the destination disk (or disks) for the restore.
- d. Click **Next**.
- e. In the *Disk Mapping Preview* page, review the parameters of the restore actions you selected. Do one of the following:
 - If you see a warning that the automatically mapped volumes are not properly mapped, click **Back**, and proceed to [step 18](#).
 - If the mapping is sufficient, select the option to overwrite selected volumes, and then click **Finish**.

18. To map volumes manually, on the *Disk Mapping* page, complete the following steps:


- a. From the Volume mapping drop-down menu, select **Manual**.
- b. In the Destination column for each volume listed, select a destination volume you want to restore. Optionally, for each volume listed, if you do not wish to restore the volume, select **Do not restore** from the Destination drop-down menu.

 **NOTE:** At least one volume must be selected to perform the restore.

19. Click **Next**.
20. If the volumes you want to restore contain Oracle, SQL Server, or Microsoft Exchange databases, and if you are performing a Live Recovery, then on the *Dismount Databases* page, you are prompted to dismount them. Optionally, if you want to remount these databases after the restore is complete, select **Automatically remount all databases after the recovery point is restored**.

21. On the *Summary* page, do the following:

- a. Verify the bare metal restore details.
- b. When satisfied, select **IMPORTANT! I understand that this operation will overwrite selected volumes with the data from the selected recovery point**.

 **CAUTION:** All existing partitions and data on the target drive will be permanently removed and replaced with the contents of the selected recovery point, including the operating system and all data.

- c. Click **Finish**.

The Restore Machine Wizard closes. The restore begins. If Toast alerts are enabled for this type of event, you see a message that the restore process has begun. You can monitor the progress on the *Events* page. For more information, see [Viewing events using tasks, alerts, and journal pages](#).

Using the Universal Recovery Console for a BMR




The Universal Recovery Console (URC) is a recovery environment embedded into a bootable ISO image, and is used to perform bare metal restore. When you boot the BMR target from the boot CD ISO image, the URC environment appears. The user interface appears slightly differently for Windows and Linux targets.

- The URC for Windows targets uses a graphical user interface based on Windows Preinstallation Environment (Win PE) 10 OS.
- The URC for Linux targets uses a Linux Ubuntu 16 command line interface.


On a Linux recovery target machine, the sole purpose of the URC is to provide single-use credentials to connect the BMR target machine with a running Rapid Recovery Core instance to perform the restore process.

In addition to providing single-use credentials for BMR, the URC for Windows contains a full-featured recovery environment. It includes function buttons and (when launched) a console.

The buttons in the Windows-based URC perform the following functions:



Button	Button Label	Description
	Start Universal Recovery Console	Launches the console from which you can manage drivers on the boot CD, manage additional drivers on the BMR target, perform bare metal restore from a Rapid Recovery archive, and monitor the restore progress.
	Useful Tools	Menu to access tools that may be required to help with your bare metal restore. For example, to launch a browser that runs on Windows PE, select Chromium from the Tools menu. For specific information, see About Windows Universal Recovery Console tools .
	Power menu	This menu includes options to reboot or shut down the BMR target machine. Each time you reboot, the authentication key is refreshed.

The Windows-based URC includes the following tabs:

Tab Name	Description
Boot CD Driver Manager	Lets you manage the drivers available on the boot CD. Click the arrow next to each item to expand to show its child objects. After you make changes, click Force Load to apply the changes and test the drivers.  NOTE: Items listed under Other devices do not yet have the correct drivers associated with them.
Existing Windows Driver Manager	Lets you load and manage drivers not included on the boot CD.

Tab Name	Description
Restore from Archive	Lets you perform a BMR from a Rapid Recovery archive.
Restore Progress	Lets you monitor the process of the bare metal restore. This tab only appears when a restore takes place.


For both Windows and Linux BMR target machines, the *Authentication* area shows the following information:

Button	Purpose	Description
	IP address	When an appropriate network adapter is loaded, the IP address of the BMR target machine is displayed.
	Authentication key	A new single-use password generates each time the BMR target machine is started using the boot ISO image.

Write down the authentication information. You will need this information to connect the BMR target machine with the Rapid Recovery Core Console to complete the restore process.

About Windows Universal Recovery Console tools

The Windows-based Universal Recovery Console (URC) includes access to tools that may assist you in completing a bare metal restore (BMR).

You can find the following tools by clicking  (Useful tools) from the top buttons displayed on the BMR target machine when booted into the URC. These tools include the following:

- **Far Manager.** This tool is similar to Windows Explorer. It provides a way to browse for files on the server until you complete the BMR and install an operating system with its own browsing function, such as Windows Explorer.
- **Chromium.** This open-source browser lets you access the Internet on a server that has a network controller loaded through the URC.
- **PuTTY.** This tool is an open-source terminal emulator. In the context of performing a BMR in Rapid Recovery, it lets you connect to a NAS storage device that does not include a user interface. This capability may be necessary if you want to restore from an archive stored on a NAS.
- **Notepad.** As in a Windows operating system, this text tool lets you type unformatted notes and view log files.
- **Task Manager.** As in a Windows operating system, this tool lets you manage processes and monitor the performance of the server while the restore is in progress.
- **Registry Editor.** As in a Windows operating system, this tool lets you change the system registry of the BMR target.
- **Command Prompt.** This tool lets you perform commands on the BMR target outside of the URC until you install a user interface.

Loading drivers using the Universal Recovery Console

The following tasks are prerequisites for this procedure.

- [Creating a boot CD ISO image](#)
- [Transferring the boot CD ISO image to media](#)
- [Loading the boot CD and starting the target machine](#)

The Universal Recovery Console lets you add any drivers that were not included in the ISO image but are required for a successful bare metal restore.

This task is part of the process for [Using the Universal Recovery Console for a BMR](#).

When creating a boot CD, you can add necessary drivers to the ISO image. After you boot into the target machine, you also can load storage or network drivers from within the Universal Recovery Console (URC).


If you are restoring to dissimilar hardware, you must inject storage controller, RAID, AHCI, chipset, and other drivers if they are not already on the boot CD. These drivers make it possible for the operating system to operate all devices on your target server successfully after you restart the system following the restore process.

Complete the steps in one of the following procedures to load drivers using the URC:

- [Loading drivers in the Universal Recovery Console using portable media](#)
- [Loading a driver in the URC using Chromium](#)

Loading drivers in the Universal Recovery Console using portable media



Complete the following procedure to use a portable media device to load drivers in the Universal Recovery Console (URC).

1. On an internet-connected machine, download the drivers from the manufacturer's website for the server and unpack them.
2. Compress each driver into a `.zip` file using an appropriate compression utility (for example, WinZip).
3. Copy and save the `.zip` file of drivers onto a portable media device, such as a USB drive.
4. Remove the media from the connected machine and insert it into the BMR target machine.
5. On the target server, load the boot CD ISO image from removable media and start the machine. The Quest splash screen appears.
6. To start the URC, click the  (Start URC) button. The URC opens to the Boot CD driver manager tab.
7. Expand the **Other devices** list. This list shows the drivers that are necessary for the hardware but are not included in the boot CD.
8. Right-click a device from the list, and then click **Load Driver**.
9. In the *Select driver load mode* window, select one of the following options:
 - Load single driver package (driver will be loaded without verification for device support)
 - Scan folder for driver packets (drivers for selected device will be searched in selected folder)

10. Expand the drive for the portable media device, select the driver (with file extension `.inf`), and then click **OK**.
The driver loads to the current operating system.
11. In the *Info* window, click **OK** to acknowledge that the driver successfully loaded.
12. Repeat this procedure as necessary for each driver you want to load.

Loading a driver in the URC using Chromium

Complete the following procedure to use the Chromium browser that comes installed on the boot CD to load drivers while in the URC.

1. On the target server, load the boot CD and start the machine.
The Quest splash screen appears.
2. To start the URC, from the buttons at the top of the screen, click  (Start URC).
The URC opens to the Boot CD driver manager tab.
3. From the buttons at the top of the screen, click  (Open useful tools), and then click **Chromium**.
The Chromium browser opens.
4. Using Chromium, navigate to a website where you can download the necessary driver.
5. Download the driver or drivers to your location choice, such as a local folder or a network file share.
If you do not specify a download location, by default Chromium downloads files to the path `Boot (X:)\Program Files\Chromium`.
6. Expand the **Other devices** list.
This list shows the drivers that are necessary for the hardware but are not included in the boot CD.
7. Right-click a device from the list, and then click **Load Driver**.
8. In the Select driver load mode window, select one of the following options:
 - Load single driver package (driver is loaded without verification for device support)
 - Scan folder for driver packets (drivers for selected device are searched in selected folder)
9. Navigate to the location where you saved the driver, select the driver, and then click **OK**.
The driver loads to the current operating system.
10. In the Info window, click **OK** to acknowledge that the driver successfully loaded.
11. Repeat this procedure as necessary for each driver you want to load.

Selecting a recovery point and initiating a BMR

After the Universal Recovery Console (URC) is accessible on the bare metal restore (BMR) target machine, you must select the recovery point that you want to restore.

Navigate to the Core Console to select which recovery point you want to load, and then designate the recovery console as the destination for the restored data.

i **NOTE:** This step is required to perform BMR on all Windows machines and optional to perform BMR on Linux machines.

This task is part of the process for [Using the Universal Recovery Console for a BMR](#).

If performing a BMR for a Linux machine from the Core Console, then this task is also a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Launching a bare metal restore for a Linux machine using the command line](#).

1. In the Rapid Recovery Core Console, from the list of protected machines (or attached archives, if relevant), click the name of the protected machine you want to restore.
The *Summary* page for the selected machine appears.
2. Click **Recovery Points**.
3. Next to the recovery point you want to use for the BMR, click the drop-down menu, and then click **Restore**.
The Restore Machine Wizard appears.
4. Select **Recover to any target machine using a Boot CD**.
5. Select **I already have a Boot CD running on the target machine**.
The authentication text boxes appear.
6. Enter the information about the machine you want to restore as described in the following table.

Table 149: Target machine information

Text Box	Description
IP Address	The IP address of the machine to which you want to restore. This is identical to the IP address displayed in the URC.
Authentication Key	The specific password to connect to the selected server. This is identical to the Authentication Key displayed in the URC.

7. Click **Next**.
If the connection information you entered matches the URC, and if the Core and the target server can identify each other properly on the network, then the volumes for the selected recovery point are loaded, and the *Disk Mapping* page appears. In this case, your next step is to map volumes.
8. Proceed to [About disk mapping for a bare metal restore](#) to learn about your disk-mapping options.

About disk mapping for a bare metal restore

After you connect to the Universal Recovery Console, you need to map volumes between those listed in the recovery point and the volumes existing on the target hardware.

Rapid Recovery attempts to automatically map volumes. If you accept the default mapping, then the disk on the destination machine is cleaned and re-partitioned and any previously existing data is deleted. The alignment is performed in the order the volumes are listed in the recovery point, and the volumes are allocated to the disks appropriately according to size, and so on. Assuming there is enough space on the target drive, no partitioning is required when using automatic disk alignment. A disk can be used by multiple volumes. If you manually map the drives, note that you cannot use the same disk twice.

For manual mapping, you must have the new machine correctly formatted already before restoring it.

i **NOTE:** When restoring a Linux machine and using manual mapping, additionally, the partitions must be mounted prior to completing the restore.

The destination machine must have a separate partition for each volume in the recovery point, including the system reserved volume. For more information, see [Using the Universal Recovery Console for a BMR](#).

As of Rapid Recovery Core release 6.8, you can use automatic mapping when performing BMR for LVM and software-based RAID volumes on Linux machines without first creating the partitions.

Complete the procedure for one of the following disk-mapping options:

- [Automatically mapping disks for a BMR](#)
- [Manually mapping disks for a BMR](#)

NOTE: ReFS is supported for Rapid Recovery on machines with Windows Server 2012 R2, Windows Server 2016, and the most recent release of Windows Server 2019. If protecting machines with the ReFS file system, the version of Windows that hosts the Rapid Recovery Core must be newer than the Windows version on the protected machine.

CAUTION: Bare metal restore of Storage Spaces disks configuration (a feature of Windows 8.1) is not supported. For details, see the *Rapid Recovery Installation and Upgrade Guide*.

This task is part of the process for [Using the Universal Recovery Console for a BMR](#).

If performing a BMR for a Linux machine from the Core Console, then this task is also a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Launching a bare metal restore for Linux](#).

Automatically mapping disks for a BMR

This procedure lets you automatically map disks during a bare metal restore (BMR) using the Restore Machine Wizard.

Complete the steps in the following procedure to automatically select the volumes you want to recover and where to restore them.

1. On the *Disk Mapping* page of the Restore Machine Wizard, next to Volume mapping, select **Automatic** from the drop-down menu.
2. In the left table, verify that the appropriate volumes are listed and are selected.

NOTE: Typically for a BMR, you should restore, at minimum, the system reserved volume and the system volume (usually, but not always, the `C:\` volume). You must select at least one volume to perform a BMR.

3. In the right table, select the disk or disks to which you want to map volumes on the target machine.
4. Click **Next**.
5. On the *Disk Mapping Preview* page, review the mapping of the recovery point volumes and the destination volume for the restore.
6. To begin the restore, click **Finish**.

CAUTION: If you select *Begin Restore*, all existing partitions and data on the target drive are permanently removed and replaced with the contents of the selected recovery point, including the operating system and all data.

Manually mapping disks for a BMR

This procedure describes how to designate the drive or volume locations on the BMR target volume when performing BMR from a recovery point.

To manually map disks on Linux BMR target machines, you must first use DiskPart from the Command Line on the BMR target machine to create and format target volumes. For more information, see [DiskPart Command-Line Options \(Standard 7 SP1\)](#) on the Microsoft Developer Network.

Complete the steps in the following procedure to manually select the volumes you want to recover and where to restore them.

1. On the *Disk Mapping* page of the Restore Machine Wizard, next to Volume mapping, select **Manual** from the drop-down menu.

i | **NOTE:** If no volumes exist on the drive of the machine on which you are performing a bare metal restore (BMR), you cannot see this option or manually map volumes.

2. In the *Volume Mapping* area, under Source Volume, verify that the source volume is selected, and that the appropriate volumes are listed beneath it and are selected.
3. For each source volume, from the *Destination* drop-down menu, select the appropriate destination to restore on the BMR target volume. If you do not want to restore a specific volume, select **Do not restore**.
4. When mappings for all listed volumes are selected, then click **Next**.
5. In the confirmation dialog box, review the mapping of the source of the recovery point and the destination volume for the restore.
6. To begin the restore, click **Finish**.

! | **CAUTION:** When you proceed with the restore, all existing partitions and data on the target drive is removed permanently, and replaced with the contents of the selected recovery point, including the operating system and all data.

Performing a BMR from an archive

Rapid Recovery lets you restore a machine from bare metal using an archived recovery point.

i | **NOTE:** This topic applies to Windows machines. If you want to restore a Linux machine from an archive, mount or attach the archive, and then restore the data from the desired recovery point.

The following tasks are prerequisites for this procedure.

- [Creating an archive](#)
- [Creating a boot CD ISO image](#)
- [Loading the boot CD and starting the target machine](#)

From the Universal Recovery Console (URC) on a Windows-based BMR target machine, you can perform a bare metal restore from a recovery point stored in a Rapid Recovery archive. The URC lets you reach this archive whether it is on a local drive, a network share, or a cloud account.

i | **NOTE:** This procedure assumes you have an archive from which to restore, and a boot CD ISO image; and that you have started the BMR target machine from the boot ISO image.

1. From the Universal Recovery Console of the BMR target machine, click **Restore from Archive**. The **Restore from Archive** tab displays the *Location* page, and shows the remaining steps required to restore from an archive.
2. From the **Location type** drop-down menu, select the location of your archive. You can choose from the following options.
 - Local
 - Network
 - Cloud

3. Based on the location type selected in the previous step, enter the credentials described in the following table.

Table 150: Location type credentials options

Location Type	Option	Description
Local	Local path	The current location of the archive. For example, D:\work\archive.
Network	Network path	The current location of the archive. For example: \\machine\shared_folder\archive.
	User	The user name for network share access.
	Password	The password for network share access.
Cloud	Cloud type	The provider of your cloud storage location. Select from the following options:
		<ul style="list-style-type: none">• Microsoft Azure
		<ul style="list-style-type: none">• Amazon S3
		<ul style="list-style-type: none">• Powered by OpenStack
		<ul style="list-style-type: none">• Rackspace Cloud Files
		<ul style="list-style-type: none">• Google Cloud
		<ul style="list-style-type: none">• Microsoft Azure Resource Management

4. If you selected a cloud type, complete the credentials that pertain to your cloud provider.
- For Microsoft Azure, complete the following steps:
 - a. Enter the following credentials:
 - Storage Account Name
 - Access Key
 - b. For the Container name, from the drop-down list, select a container.
 - c. For the Cloud path, from the drop-down list, select the path to the archive.
 - For Amazon S3, complete the following steps:
 - a. Enter the following credentials:
 - Access key
 - Secret key
 - b. For the Container name, from the drop-down list, select a container.
 - c. For the Cloud path, from the drop-down list, select the path to the archive.
 - For Powered by OpenStack or Rackspace Cloud Files accounts, complete the following steps:
 - a. Enter the following information:
 - Region
 - User
 - b. Select one of the following options:
 - Password
 - API Key
 - c. In the text box, enter the information based on your selection in Step c.
 - d. Enter the following information:
 - Tenant ID
 - Authentication URL
 - e. For the Container name, from the drop-down list, select a container.
 - f. For the Cloud path, from the drop-down list, select the path to the archive.
 - For Google Cloud accounts, complete the following steps:
 - a. Upload the certificate file
 - b. Enter the following credentials:
 - Private key
 - c. Enter the following information:
 - Project ID
 - Service account email
 - d. For the Container name, from the drop-down list, select a container.
 - e. For the Cloud path, from the drop-down list, select the path to the archive.
 - f. If prompted for the Bucket name, from the drop-down list, select the appropriate bucket.

5. Click **Next**.
6. On the Machines page, select the machine you want to restore, and then click **Next**.
7. On the Recovery Points page, select the recovery point you want to use to restore the machine, and then click **Next**.
8. On the Mapping page, select one of the following options, and then complete the corresponding steps:
 - From the **Volume Mapping** drop-down list, select **Automatic**.
 - a. In the left table, verify that the appropriate volumes are listed and are selected.

i **NOTE:** Typically for a BMR, you should restore, at minimum, the system reserved volume and the system volume (usually, but not always, the C: \ volume). You must select at least one volume to perform a BMR.
 - b. In the right table, select the disk or disks to which you want to map volumes on the target machine.
 - From the **Volume Mapping** drop-down list, select **Manual**.

i **NOTE:** To manually map disks, you must first use DiskPart on the Command Line to create and format target volumes. For more information, see [DiskPart Command-Line Options \(Standard 7 SP1\)](#) on the Microsoft Developer Network.

i **NOTE:** If no volumes exist on the drive of the machine on which you are performing a bare metal restore (BMR), you cannot see this option or manually map volumes.
 - Under **Destination Volumes**, from the drop-down menu, select the appropriate target volume for each volume in the recovery point.
9. In the **mount maps path** text box, enter a destination for the temporary storage of mapping files. The default location is X:\ProgramData\AppRecovery\IndexEntriesMaps.


i **NOTE:** To ensure that your destination has sufficient free space, divide the total mount volume capacity by 1,024. For example, using the formula (Mount volume total capacity) / 1024 = Free space, then 1 TB / 1024 = 1 GB.
10. Click **Restore**.
The URC maps the volumes to the new disk or disks.
11. Click **Restore**.
The URC restores the data to the target machine. You can view the progress on the Restore progress tab.
12. After the restore is complete, remove the boot CD.
13. To boot the BMR target machine into Windows, restart the machine.

Loading drivers to the operating system

This procedure describes how to load drivers to the operating system on a bare metal restore (BMR) target. You can use this procedure to troubleshoot driver conflicts for the restored machine.

After performing a BMR, you can load or inject additional drivers to the OS of the restored machine from the URC. You must have the drivers accessible in a compressed format.

1. From the URC of the BMR target machine, click **Existing Windows driver management**.
The **Existing Windows driver management** tab of the console appears.

2. From the drop-down list, select an operating system.
The URC searches for available drivers.
3. To load additional drivers, click  **Force Load**.
4. Navigate through the filing system to locate the compressed driver file, and then select the file.
5. Click **OK**.
The URC loads the driver into the operating system you selected.
6. Repeat [step 3](#) through [step 5](#) for each additional driver you need to load.
7. Restart the BMR target machine.
The BMR is complete. If you experience an issue when you restart, see [Repairing boot problems](#).

Performing a bare metal restore for Linux machines

In Rapid Recovery, you can perform a Bare Metal Restore (BMR) for a protected Linux machine, including a restore of the system volume. BMR functionality is supported for Linux using the Restore Machine Wizard from the Core Console, and also using the command line `local_mount` utility.

CAUTION: Rapid Recovery supports ext2 partition types only if the kernel is version 3.10 and above. If using an earlier kernel, convert any ext2 partitions to ext3, ext4, or XFS before you begin protecting and backing up the machine.

CAUTION: When you boot a restored Linux machine for the first time after a BMR, Rapid Recovery Core first attempts to capture an incremental snapshot of the restored machine. If incremental capture is not possible due to the amount of data and the state of the machine, then Rapid Recovery Core captures a base image of the restored machine. This process takes more time than taking an incremental snapshot. For more information about base images and incremental snapshots, see [Understanding protection schedules](#).

To perform a bare metal restore for Linux machines, perform the following tasks.

- **Manage a Linux boot image.** The Linux Live DVD boot ISO image is used to start up the BMR target machine, from which you can access the Universal Recovery Console to communicate with backups on the Core. See [Managing a Linux boot image](#).
 - To obtain the boot image for BMR, make sure the Live DVD version matches your Core version. For more information, see [About the boot ISO image for Linux](#) followed by [Downloading a boot ISO image for Linux](#).
 - If you require physical media to start up the destination Linux machine, you will need to transfer the ISO image to media. See [Saving the Live DVD ISO image to media](#).
 - In all cases, you will need to load the boot ISO image into the BMR target machine and boot the server from that image. See [Loading the Live DVD and starting the target machine](#).
 - After you load the media and boot, you must connect the Linux BMR target machine to the Rapid Recovery Core. See [Connecting to the BMR target from the Rapid Recovery Core](#).

- **Manage Partitions.** You may need to create or mount partitions before performing a BMR on a Linux machine. See [Managing Linux partitions](#).
 - The Linux system on which you are performing a BMR must have the same partitions as the source volumes in the recovery point. You may need to create additional partitions on the target system, if required. See [Creating partitions on the destination drive](#).
 - If you are performing a manual BMR, you must first mount partitions. See [Mounting partitions from the command line](#). Steps to mount partitions are included in the process to perform a BMR from the command line. See [Launching a bare metal restore for a Linux machine using the command line](#). If you are using auto-partitioning for BMR within the Core Console, you do not need to mount partitions. Rapid Recovery will restore the same partitions as those included in the recovery point being restored.
- **Launch a Bare Metal Restore for Linux.** Once the destination machine is started from the Live DVD boot image, you can launch the BMR. The tasks required depend on whether you will perform this from the Rapid Recovery Core Console user interface or from the command line using the `local_mount` utility. See [Launching a bare metal restore for Linux](#).
 - If using the Core Console, you will need to initiate a restore from a recovery point on the Core. See [Selecting a recovery point and initiating a BMR](#).
 - If using the Core Console, you will need to map the volumes from the UI. See [About disk mapping for a bare metal restore](#).
 - Optionally, if restoring from the command line, you can use the GNU Screen utility to enhance your ability to scroll and see commands in the terminal console. This utility opens by default. If you close it, you can start it again. For more information, see [Starting the Screen utility](#).
 - If using `local_mount`, all tasks will be performed at the command line. For more information, see [Launching a bare metal restore for a Linux machine using the command line](#).
- **Verify a Bare Metal Restore.** After starting the bare metal restore, you can verify and monitor your progress. See [Verifying the bare metal restore from the command line](#).
 - You can monitor the progress of your restore. See [Viewing the recovery progress](#).
 - Once completed, you can start the restored server. See [Starting a restored target server](#).
 - Troubleshoot the BMR process. See [Troubleshooting connections to the Universal Recovery Console](#) and [Repairing boot problems](#).

Managing Linux partitions

When performing a BMR, the destination drive onto which you will be restoring data must have the same partitions as in the recovery point you are restoring. You may need to create partitions to meet this requirement.

You can launch the restore from the command line using the `local_mount` utility, or you can launch the restore from the Rapid Recovery Core Console. If restoring using the user interface, you must first mount the partitions.

Managing Linux partitions is a step in [Performing a bare metal restore for Linux machines](#).

You can perform the following tasks:

- [Creating partitions on the destination drive](#)
- [Formatting partitions on the destination drive](#)
- [Mounting partitions from the command line](#)

Creating partitions on the destination drive

Often, when performing a BMR, the destination drive is a new volume that may consist of a single partition. The drive on the destination machine must have the same partition table as in the recovery point, including the size of the volumes. If the destination drive does not contain the same partitions, you must create them before performing the bare metal restore. Use the fdisk utility to create partitions on the destination drive equal to the partitions on the source drive.

Formatting partitions on the destination drive

After creating partitions on a new volume on the destination drive to perform bare metal restore, if you are not using auto partition, you must format the partitions before they can be mounted. If this situation applies to you, format partitions in ext3, ext4, or XFS formats.

For all other scenarios, you do not need to format partitions.

Mounting partitions from the command line

If performing manual partitioning for BMR of a Linux machine using the Restore Machine Wizard, you must first mount the appropriate partitions on the destination machine. Perform this action from the command line in the Universal Recovery Console.

This process is a step in performing a BMR for Linux machines from the command line.

Launching a bare metal restore for Linux

Before launching a bare metal restore (BMR) for a Linux machine, the following conditions are required:

- To restore a recovery point saved on the Core, you must have the appropriate hardware in place. For more information, see [Prerequisites for performing a bare metal restore for Windows or Linux machines](#).
- The BMR target machine must be started using the Live DVD boot ISO image. For more information, see [Managing a Linux boot image](#).
- The number of volumes on the Linux machine to be restored must match the number of volumes in the recovery point. You must also decide whether to restore from the Restore Machines wizard in the Rapid Recovery Core Console, or from the command line using local_mount. For more information, see [Managing Linux partitions](#).
- If restoring from the Core Console UI, the first step in launching a BMR is to select the appropriate recovery point, then initiate the restore to the hardware by specifying the IP address and temporary password you obtained from the Universal Recovery Console. You must then map the drives and start the restore.

This process is a step in [Performing a bare metal restore for Linux machines](#).

To launch a BMR from the Rapid Recovery Core Console, perform the following tasks.

- [Selecting a recovery point and initiating a BMR](#)
- [About disk mapping for a bare metal restore](#)

If restoring from the command line using the local_mount utility, then you must first set appropriate privileges, mount volumes, execute local_mount, obtain information about the Core from the list of machines, connect to the Core, obtain a list of recovery points, select the recovery point you want to roll back onto bare metal, and launch the restore.

Optionally, you may want to start the Screen utility.

To launch a BMR from the command line, perform the following tasks.

- [Starting the Screen utility](#)
- [Launching a bare metal restore for a Linux machine using the command line](#)

Starting the Screen utility

Included on the Live DVD is Screen, a GNU utility which is available when you boot from the Live DVD into the Universal Recovery Console. Screen allows users to manage multiple shells simultaneously over a single Secure Shell (SSH) session or console window. This allows you to perform one task in a terminal window (such as verify mounted volumes) and, while that is running, open or switch to another shell instance to perform another task (such as to run the `local_mount` utility).

The Screen utility also has its own scroll-back buffer, which enables you to scroll the screen to view larger amounts of data, such as a list of recovery points.

i | **NOTE:** This utility is provided for convenience; use of the Screen utility is optional.

The Screen utility starts on the machine booted with the Live DVD by default. However, if you have closed this application, you must start the Screen utility from the Live DVD using the procedure below.

1. If the machine was booted from the Live DVD, then in the terminal window, type `screen` and press **Enter**. The Screen utility starts.

Launching a bare metal restore for a Linux machine using the command line

Once the Live DVD ISO image is accessible on the machine on which you want to perform a BMR, and the number and size of volumes matches between the target machine and the recovery point you want to restore to bare metal, then you can launch a restore from the command line using the `local_mount` utility.

If you want to perform a BMR using the Rapid Recovery Core Console user interface, see [Selecting a recovery point and initiating a BMR](#).

i | **NOTE:** When performing this procedure, do not attempt to mount recovery points to the `/tmp` folder, which contains the `rapidrecovery-vdisk` files.

This task is a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Launching a bare metal restore for a Linux machine using the command line](#).

Complete the steps in this procedure to select a recovery point on the Core to restore to the physical or virtual BMR target machine.

1. To run the Rapid Recovery `local_mount` utility as root, type the following command and then press **Enter**:

```
sudo local_mount
```

2. To list the protected machines, type the following command and then press **Enter**:

```
lm
```

3. When prompted, enter the connection information for the Rapid Recovery Core as described in the following table, pressing **Enter** after each required command:

Table 151: Rapid Recovery Core connection information

Text Box	Description	Required
Rapid Recovery Core IP address or hostname	The IP address or hostname of the Rapid Recovery Core.	Yes
Domain	The domain of the Rapid Recovery Core. This is optional.	No
User	The user name for an administrative user on the Core	Yes
Password	The password used to connect the administrative user to the Core.	Yes

A list displays showing the machines protected by the Rapid Recovery Core. It lists the machines found by line item number, the host display name or IP address, and an ID number for the machine.

4. To list the recovery points for the machine that you want to restore, type the list recovery points command using the following syntax and then press **Enter**:

```
lr <machine_line_item_number>
```

i | **NOTE:** You can also enter the machine ID number in this command instead of the line item number.

A list displays the base and incremental recovery points for that machine. This list includes:

- A line item number
- Date and time stamp
- A numbered list of volumes within the recovery point
- Location of the volume
- Size of the recovery point
- An ID number for the volume that includes a sequence number at the end, which identifies the recovery point

5. To select the recovery point for a restore, enter the following command and then press **Enter**:

```
r <recovery_point_ID_number> <path>
```

CAUTION: You must ensure that the system volume is not mounted.

NOTE: If you started the machine from the Live DVD, then the system volume is not mounted.

This command restores the volume image specified by the ID from the Core to the specified path. The path for the restore is the path for the device file descriptor and is not the directory to which it is mounted.

NOTE: You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. In that case, use the agent/machine line number (from the `lm` output), followed by the recovery point line number and volume number (from the numbered list of volumes within the recovery point), followed by the path. For example: `r <machine_line_item_number> <base_image_recovery_point_line_number> <volume_number> <path>`

In this command, `<path>` is the file descriptor for the actual volume.

For this syntax, when variables shown in brackets are replaced with values, the command looks like this example: `r 1 24 1 /dev/sda1`

6. When prompted to proceed, enter **y** for Yes and then press **Enter**.

After the restore begins, a series of messages display that notify you of the restore completion status.

NOTE: If you receive an exception message, the details regarding that exception can be found in the `local_mount.log` file. The `local_mount.log` file is located in `/var/log/apprecovery`.

7. Upon a successful restore, exit `local_mount` by typing `exit` and then press **Enter**.
8. Your next step is to verify the restore. For more information, see [Verifying the bare metal restore from the command line](#).


Restoring volumes for a Linux machine using the command line

Before restoring from a recovery point using the command line, you must dismount the disk on which you will be restoring data.

This procedure describes how to restore the volumes on a protected Linux machine from the recovery points stored in the Rapid Recovery Core using the command line `local_mount` utility.

NOTE: In previous releases, this process was referred to as performing a rollback.

For information about performing restore for Linux volumes from the Rapid Recovery Core Console, see [Performing a bare metal restore using the Restore Machine Wizard](#).

You can begin a restore from any location on the Rapid Recovery Core Console by clicking the  Restore icon in the Rapid Recovery button bar. When you start a restore in this manner, you must specify which of the machines protected on the Core you want to restore, and then drill down to the volume you want to restore.

In Rapid Recovery, you can restore volumes on your protected Linux machines using the command line `local_mount` utility.

CAUTION: To restore the system or root (`/`) partition or entire operating system, see [Performing a bare metal restore for Linux machines](#).

1. Run the Rapid Recovery `local_mount` utility as root, for example:

```
sudo local_mount
```

2. At the Rapid Recovery mount prompt, enter the following command to list the protected machines.

```
lm
```

3. When prompted, enter the IP address or hostname of your Rapid Recovery Core server.

4. Enter the logon credentials, that is, the user name and password, for this server.

A list displays showing the machines protected by this Rapid Recovery Core server. It lists the protected machines found by line item number, host or IP address, and an ID number for the machine (for example: 7d658e5f-fa08-4600-95f0-5f486bc1b6a4#de0896fd-571a-4cc5-aeed-264d2c3c72f4#f377e145-dd4d-3ac3-5b15-37ce8f4913ba:2).

5. Enter the following command to list the currently mounted recovery points for the specified machine:

```
lr <machine_line_item_number>
```

i **NOTE:** You can also enter the machine ID number in this command instead of the line item number.

A list displays that shows the base and incremental recovery points for that machine. This list includes a line item number, date and time stamp, location of volume, size of recovery point, and an ID number for the volume that includes a sequence number at the end (for example, "293cc667-44b4-48ab-91d8-44bc74252a4f:2"), which identifies the recovery point.

6. Enter the following command to select a recovery point to restore:

```
r <volume_recovery_point_ID_number> <device_path>
```

This command restores the volume image specified by the ID from the Core to the specified path. The path for the restore is the path for the device file descriptor, not the directory to which it is mounted.

You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. In that case, you would use the protected machine line number (from the `lm` output), followed by the recovery point line number and volume letter, followed by the path, such as, `r <machine_line_item_number> <recovery_point_line_number> <volume_letter> <path>`. In this command, `<path>` is the file descriptor for the actual volume.

For example, if the `lm` output lists three protected machines, and you enter the `lr` command for protected machine number 2, and you want to restore the 23 recovery point volume `b` to the volume that was mounted to the directory `/dev/sda5`, the command would be:

```
r 2 23 b /dev/sda5
```

i **NOTE:** It is possible to restore to `/` if needed. If performing a Bare Metal Restore using a Live DVD, it is assumed you want to restore to a different machine. For more information, see [Launching a bare metal restore for Linux](#).

7. When prompted to proceed, enter `y` for Yes.

Once the restore proceeds, a series of messages will display to notify you of the status.

8. Upon a successful restore, the `local_mount` utility will automatically mount and re-attach the kernel module to the restored volume if the target was previously protected and mounted. If not, you will need to mount the restored volume to the local disk and then should verify that the files are restored (for example, you can use the `sudo mount` command and then the `ls` command.)

Verifying a bare metal restore


After you perform a bare metal restore (BMR), you can verify the progress of the restore. When the action is completed successfully, you can start the restored server. Some troubleshooting steps are included if you encounter difficulties connecting to the Universal Recovery Console to complete the restore, and if you need to repair startup problems with the restored machine.

You can perform the following tasks:

- [Viewing the recovery progress](#)
- [Starting a restored target server](#)
- [Troubleshooting connections to the Universal Recovery Console](#)
- [Repairing boot problems](#)

Viewing the recovery progress

Complete the steps in this procedure to view the progress of restoring data from a recovery point (including bare metal restore) initiated from the Rapid Recovery Core Console.


1. After you initiate the process restoring data from a recovery point, while the task is in process, you can view its progress from the  Running Tasks drop-down menu on the Core Console.
2. Optionally, you can view detailed information in the Events page. For more information about monitoring Rapid Recovery events, see [Viewing events using tasks, alerts, and journal pages](#).
3. Additionally, only in the case of performing a restore or BMR of a protected Linux machine from the command line, you can view the recovery progress within the same command shell (window) from which you initiated the restore.

Starting a restored target server

Complete the steps in this procedure to start the restored target server.

i **NOTE:** Before starting the restored target server, you should verify that the recovery was successful. For more information, see [Viewing the recovery progress](#).

This task is part of the process for [Verifying a bare metal restore](#).

1. On the target server, verify that the Rapid Recovery Universal Recovery Console is active.
2. Eject the boot CD (or disconnect physical media with the boot CD image) from the restored server.
3. In the Universal Recovery Console, from the top function buttons, click  (Power menu), and then select **Reboot**.
4. Specify that you want to start the operating system normally.
5. Log on to the machine. The system should be restored to the state captured in the recovery point.

Troubleshooting connections to the Universal Recovery Console

The following are troubleshooting steps for connecting to the boot CD image as part of the process for [Selecting a recovery point and initiating a BMR](#).

If an error displays indicating that the Core could not connect to the remote server, then any of several possible causes are likely.

- Verify that the IP address and Current Password displayed in the URC are identical to the information you entered in the Recovery Console Instance dialog box.
- To reach the server on which to restore data, the Core must be able to identify the server on the network. To determine if server identification is possible, you can open a command prompt on the Core and ping the IP address of the target BMR server. You can also open a command prompt on the target server and ping the IP address of the Rapid Recovery Core.
- Verify that the network adapter settings are compatible between Core and target BMR server.

Repairing boot problems

The following tasks are prerequisites for this procedure.

- [Creating a boot CD ISO image](#)
- [Loading the boot CD and starting the target machine](#)
- [Loading drivers using the Universal Recovery Console](#)

Complete the steps in this procedure to repair startup problems. Keep in mind that if you restored to dissimilar hardware, you must have injected storage controller, RAID, AHCI, chipset and other drivers if they are not already on the boot CD. These drivers make it possible for the operating system to operate all devices on your target server successfully. For more information, see [Loading drivers using the Universal Recovery Console](#). Complete the following procedure to repair startup problems on your target server.

1. From the Universal Recovery Console, click the Existing Windows driver manager tab.
2. Click Repair Boot Problems.
The startup parameters in the target server boot record are automatically repaired.

Verifying the bare metal restore from the command line

Quest recommends performing the following steps to verify a bare metal restore completed from the command line. This task is a step in [Performing a bare metal restore for Linux machines](#).

- [Performing a file system check on the restored volume](#)
- [Using the command line to make a restored Linux machine bootable](#)

Performing a file system check on the restored volume

Once you execute a bare metal restore from the command line, you should perform a file system check on the restored volume to ensure the data restored from the recovery point was not corrupted.

This task is a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Verifying the bare metal restore from the command line](#).

Perform the task below to perform a file system check on the restored volume.

1. From the command line in the Universal Recovery Console of the Linux machine you have restored, to verify whether the appropriate partitions are mounted, type the following command and then press **Enter**:

```
df
```

2. If the restored volume is not mounted, then skip to [Step 3](#). If the restored volume is mounted, unmount it by typing the following command and then pressing **Enter**:

```
umount <mount point>
```

3. Run a file system check on the restored volumes by typing the following command and then press **Enter**:

```
fsck -f <volume>
```

If the fsck returns clean, the file system is verified.

4. Mount the appropriate volumes once again by typing the following command in format `mount <volume> <folder>`, and then press **Enter**.

For example, if the volume path is `prod/sda1` and the folder you want to mount to is `mnt`, then type the following and then press **Enter**:

```
mount /dev/sda1 /mnt
```

Using the command line to make a restored Linux machine bootable

Once you complete a clean file system check on the restored volume, you must create bootable partitions.

GNU Grand Unified Bootloader (GRUB) is a boot loader that allows administrators to configure which operating system or specific kernel configuration is used to start the system. After a BMR, the configuration file for GRUB must be modified so that the machine uses the appropriate universally unique identifier (UUID) for the root volume. Before this step you must mount the root and boot volumes, and check the UUIDs for each. This ensures that you can boot from the partition.

NOTE: This procedure applies to Linux machines that use GRUB1 or GRUB2. When using this procedure, ensure that the boot partition is healthy and protected.

GRUB or GRUB2 is typically installed with Linux operating systems. You can perform this procedure using the version that comes with your Linux distribution. If a version of GRUB is not installed, you will have to re-install the default version appropriate for your Linux distribution.

CAUTION: When you boot a restored Linux machine for the first time after a BMR, Rapid Recovery takes a base image of the restored machine. This process typically takes longer than taking an incremental snapshot. For more information about base images and incremental snapshots, see [Understanding protection schedules](#).

This task is a step in [Performing a bare metal restore for Linux machines](#). It is part of the process for [Verifying the bare metal restore from the command line](#).

Perform the task below to create bootable partitions using the command line.

1. You must mount the root volume first and then the boot volume. Mount each restored volume by using the following commands:

- a. To mount the root volume, type the following command and then press **Enter**:

```
mount /<restored volume[root]> /mnt
```

For example, if /dev/sda2 is the root volume, then type `mount /dev/sda2 /mnt` and then press **Enter**.

- b. To mount the boot volume, type the following command and then press **Enter**:

```
mount /<restored volume[boot]> /mnt/boot
```

For example, if /dev/sda1 is the boot volume, then type `mount /dev/sda1 /mnt/boot` and then press **Enter**.

i **NOTE:** Some system configurations may include the boot directory as part of the root volume.

2. If the volume size is increasing — that is, if the destination volume on the new Linux machine is larger than the volume was in the recovery point — then you must delete any existing bitmap data files.
3. Obtain the Universally Unique Identifier (UUID) of the new volumes by using the `blkid` command. Type the following and then press **Enter**:

```
blkid [volume]
```

i **NOTE:** You can also use the `ls -l /dev/disk/by-uuid` command.

4. If performing a BMR on a brand new disk on the destination machine, comment out the swap partition in `fstab` in your root volume.
5. Modifying `fstab` and `mtab` paths should occur on the restored volume, not the Live DVD. There is no need to modify paths on the Live DVD. Prepare for the installation of Grand Unified Bootloader (GRUB) by typing the following commands. Following each command, press **Enter**:

```
mount --bind /dev /mnt/dev
```

```
mount --bind /proc /mnt/proc
```

```
mount --bind /sys /mnt/sys
```

6. Change root directory by typing the following command and then press **Enter**:

```
chroot /mnt /bin/bash
```

7. Obtain the old UUID of the partition or partitions from the mounted recovery points `/etc/fstab` file and compare it to the UUIDs for the root (for Ubuntu and CentOS), boot (for CentOS and RHEL), or data partitions by typing the following command and then press **Enter**:

```
less /mnt/etc/fstab
```

8. Obtain the old UUID of the partition or partitions from the mounted recovery points `/etc/mtab` file and compare it to the UUIDs for the root (for Ubuntu and CentOS), boot (for CentOS and RHEL), and data partitions by typing the following command and then press **Enter**:

```
less /mnt/etc/mtab
```

9. If using SLES 11, install GRUB by typing the following commands, pressing **Enter** after each:

```
grub-install --recheck /dev/sda  
grub-install /dev/sda
```

10. If using Ubuntu, CentOS 6.x, RHEL 6.x, or Oracle Linux 6.x, install GRUB by typing the following command, and then press **Enter**:

```
grub-install /dev/sda
```

11. If using SLES 12, CentOS 7, RHEL 7, or Oracle 7, install GRUB2 by typing the following command, and then press **Enter**:

```
grub2-install /dev/sda
```

12. After you complete installation, run one of the following updates:

For SLES:

```
grub-install.unsupported --recheck /dev/sda  
grub-install.unsupported /dev/sda  
update-grub
```

i | **NOTE:** If the `update-grub` command does not exist on your Linux distribution, omit this option.

For other distributions:

```
grub-install /dev/sda  
update-grub
```

i | **NOTE:** If the `update-grub` command does not exist on your Linux distribution, omit this option.

13. Remove the Live DVD disk from the CD-ROM or DVD drive and restart the Linux machine.

Managing aging data

This section describes how to manage aging snapshot data saved to your repository. It includes information about retaining recovery points in your repository, retention policies, and the resulting process of rolling up recovery points to conserve space.

This section also describes how to manage retention policies that control rollup, and how to force rollup on demand.

Topics include:

[Data retention and archiving](#)

[Managing retention policies](#)

Data retention and archiving

Each time your Core captures a snapshot, the data is saved as a recovery point to your repository. Recovery points naturally accumulate over time. The Core uses a retention policy to determine how long snapshot data is retained in the repository. When nightly jobs run (specifically, during the rollup process), the Core enforces the retention policy to reduce the amount of storage space consumed. The date of each recovery point is compared to the date of the most recent recovery point. The Core then rolls up (combines) older recovery points. Over time, older recovery points in the repository are continually replaced with newer ones as the oldest recovery points eventually reach the oldest age defined in the retention period.

To keep recovery points that would otherwise be combined and eventually deleted, you can create an archive from the Core Console. An archive is a file containing a copy of the full set of recovery points for machines protected on your Core at the point in time in which it was created. You can later access archived information from the Core Console. In contrast with recovery points in the repository, recovery points in an archive do not get rolled up.

Archives are useful for maintaining compliance data; backing up your Core; seeding replication data to a remote replica Core; and for saving space in your Core for retaining recent business-critical transaction while maintaining backups for a longer period of time.

For more information about archives, see [Archiving](#).

Managing retention policies

A retention policy is a set of rules that dictates the length of time for the Core to retain recovery points before starting to roll them up. Retention policies can be set to roll up based on hours, days, weeks, months and years. You can set up to six rules (the default policy sets five rules).

Since you can back up as frequently as every 5 minutes, the first rule in the retention policy typically sets how long to retain all recovery points. For example, if you back up a machine every quarter hour, 96 recovery points are saved to the repository for that machine per day, until rollup begins. Without managing your retention policy, that amount of data can quickly fill a repository.

i **NOTE:** Administrators should note that frequent backups can have an impact on network traffic. Other factors affecting network traffic include other transfers (such as replication), the change rate of your data, and your network hardware, cables and switches.

The Core comes preset with a default retention policy. The default policy retains:

- All recovery points for three days
- One recovery point per hour for two days
- One recovery point per day for four days
- One recovery point per week for three weeks
- One recovery point per month for two months
- One recovery point per year for X years (disabled in default policy).

Following this default policy, the oldest recovery point is typically 92 days old. Data past that origination date for a default policy is deleted.

Setting the retention policy at the Core level applies automatically to all machines that the Core protects. You can change the default policy to suit your needs.

For any machine, you can also create a custom retention policy. Setting the policy at the machine level lets you specify a different retention policy than the default Core policy. For more information about configuring retention policies, see [Configuring Core default retention policy settings](#) and [Customizing retention policy settings for a protected machine](#).


Configuring Core default retention policy settings

The retention policy for the Core specifies how long the recovery points for a protected machine are stored in the repository.


The Core retention policy is enforced by a rollup process which is performed as one component of running nightly jobs. Then, recovery points beyond the age specified in the retention policy are “rolled up” (combined) into fewer recovery points that cover a less granular period of time. Applying the retention policy on a nightly basis results in the ongoing rollup of aging backups. This eventually results in the deletion of the oldest recovery points, based on the requirements specified in that retention policy.

Different retention settings can be configured for source and target Cores.

i **NOTE:** This topic is specific to customizing retention policy settings on the Rapid Recovery Core. When you save customized retention policy settings on the Core, you establish the default retention policy settings which can be applied to all machines protected by this Core. For more information on customizing retention policy settings for individual protected machines, see [Customizing retention policy settings for a protected machine](#).

1. Navigate to the Rapid Recovery Core Console.
2. On the icon bar, click  (Settings), and then do one of the following:
 - From the list of Core settings on the left side of the Settings page, click **Nightly Jobs**.
 - Scroll down on the right side of the Settings page until you can see the **Nightly Jobs** heading.

The Nightly Jobs settings for the Core appear.

3. Under **Nightly Jobs**, click  **Change**.
The *Nightly Jobs* dialog box displays.
4. To specify the time intervals for retaining the backup data as needed, in the Nightly Jobs pane, select **Rollup**, and then click **Settings**.
The *Retention Policy* dialog box for the Core default retention policy displays.
5. To restore Core retention policy settings to the default values at any time, at the bottom of the dialog box, click **Restore Defaults** and then click **Yes** to confirm.
All settings are restored to the default values described in the table in [step 6](#).

6. Enter the default schedule for retaining recovery points as described in the following table. The first setting specifies the primary retention period for all recovery points saved to the repository. Each additional setting, when enabled, provides a more granular level of retention by specifying the intervals between which recovery points should be rolled up. These settings define the duration for which recovery points are maintained.

The retention policy options are described in the following table.

Table 152: Schedule options for default retention policy

Text Box	Description
Keep all recovery points for n [retention time period]...	Specifies the primary retention period for all recovery points saved to the repository. Enter a number to represent the retention period and then select the time period. The default is 3 days. You can choose from days, weeks, months, and years.
...and then keep one recovery point per hour for n [retention time period]	If selected, this setting keeps one recovery point per hour for a time period that you specify. The default is 2 days. You can choose from days, weeks, months, and years. If you do not want to save at least one recovery point hourly, clear this option.
...and then keep one recovery point per day for n [retention time period]	If selected, this setting keeps one recovery point per day for a time period that you specify. The default is 4 days. You can choose from days, weeks, months, and years. If you do not want to save at least one recovery point daily, clear this option.
...and then keep one recovery point per week for n [retention time period]	If selected, this setting keeps one recovery point per week for a time period that you specify. The default is 3 weeks. You can choose from weeks, months, and years. If you do not want to save at least one recovery point weekly, clear this option.
...and then keep one recovery point per month for n [retention time period]	If selected, this setting keeps one recovery point per month for a time period that you specify. The default is 2 months. You can choose from months and years. If you do not want to save at least one recovery point monthly, clear this option.
...and then keep one recovery point per year for n [retention time period]	If selected, this setting keeps one recovery point per year for a number of years that you specify. The default (1 year) is disabled by default. If you do want to save at least one recovery point yearly, select this option and specify a number of years. If you do not want to save at least one recovery point yearly, clear this option.

The oldest recovery point is determined by the retention policy settings.

The following is an example of how the retention period is calculated.

Keep all recovery points for three days.

...and then keep one recovery point per hour for three days

...and then keep one recovery point per day for four days

...and then keep one recovery point per week for three weeks

...and then keep one recovery point per month for two months

...and then keep one recovery point per month for one year

In this example, the oldest recovery point would be one year and three months old.

7. When satisfied with your retention policy settings, click **Save**.
The *Retention Policy* dialog box closes.
8. In the *Nightly Jobs* dialog box, click **OK**.
The *Nightly Jobs* dialog box closes. The retention policy you defined is applied during the nightly rollup.
You can also to apply these settings when specifying the retention policy for any individual protected machine. For more information about setting retention policies for a protected machine, see [Customizing retention policy settings for a protected machine](#).


Customizing retention policy settings for a protected machine

The retention policy for a protected machine specifies how long recovery points are stored in the repository. By default, each protected machine uses the retention policy established for the Core unless you specify a custom retention policy, as described in this procedure.

Starting with AppAssure version 5.4.1, Rapid Recovery includes the ability to set disparate retention policies between a protected machine on the source Core and the corresponding replicated machine on the target Core.

Use this procedure to define a custom retention policy for a protected machine, including a replicated machine.

i NOTE: The following applies to environments upgrading from AppAssure release 5.3.x to release 5.4.1 or any version of Rapid Recovery Core. If you want to customize a retention policy for any replicated machine, first upgrade the source and target Cores to AppAssure Core release 5.4.1, and then perform the Checking Repository job on each repository in that target Core. Completing this job is likely to take a substantial amount of time, based on the size of your repository and the underlying storage system.

1. From the Protected Machines menu of the Rapid Recovery Core Console, click the name of the machine that you want to modify.
The *Summary* page for the selected machine displays.
2. Click the **Settings** menu.
The *Settings* page displays, showing configuration settings for the selected machine.
3. Optionally, click the **Nightly Jobs** link to scroll down in the Settings page to view nightly jobs settings.
4. Under the Nightly Jobs heading, click  **Change**.
The *Nightly Jobs* dialog box displays.
5. To specify the time intervals for retaining the backup data as needed, select **Rollup** and then click **Settings**.
The *Retention Policy* dialog box displays.

6. If customizing retention policies settings for a replicated machine, and if you see a caution notifying you to perform an Integrity Check on your repository, proceed with this step. Otherwise, skip to [step 7](#).
 - a. If you are prepared to perform the job, click **Check Integrity**
 - b. Click **Yes** to confirm the Integrity Check job.

i **NOTE:** Running this job could take a substantial amount of time, based on the size of your repository. During this time, you can perform no other actions (snapshots, replication, virtual export, and so on) in the repository.

Once the Checking Repository job completes all child jobs successfully, return to this procedure and continue with the next step.

7. In the *Retention Policy* dialog box, do one of the following:
 - To use the default retention policy for this protected machine, select **Use Core default retention policy**, and then click **Save**. The default policy is applied to this protected machine.
 - To define a custom retention policy for this agent, select **Use custom retention policy**, and then continue with the next step.

The *Retention Policy* dialog box expands to show custom retention policy information.

8. Enter the custom schedule for retaining recovery points as described in the following table. The first setting specifies the primary retention period for all recovery points saved to the repository. Each additional setting, when enabled, provides a more granular level of retention by specifying the intervals between which recovery points should be rolled up. These settings define the duration for which recovery points are maintained.

Table 153: Schedule options for custom retention policy

Text Box	Description
Keep all recovery points for n [retention time period]...	Specifies the primary retention period for all recovery points saved to the repository. Enter a number to represent the retention period and then select the time period. The default is 3 days. You can choose from days, weeks, months, and years.
...and then keep one recovery point per hour for n [retention time period]	If selected, this setting keeps one recovery point per hour for a time period that you specify. The default is 2 days. You can choose from days, weeks, months, and years. If you do not want to save at least one recovery point hourly, clear this option.
...and then keep one recovery point per day for n [retention time period]	If selected, this setting keeps one recovery point per day for a time period that you specify. The default is 4 days. You can choose from days, weeks, months, and years. If you do not want to save at least one recovery point daily, clear this option.
...and then keep one recovery point per week for n [retention time period]	If selected, this setting keeps one recovery point per week for a time period that you specify. The default is 3 weeks. You can choose from weeks, months, and years. If you do not want to save at least one recovery point weekly, clear this option.
...and then keep one recovery point per month for n [retention time period]	If selected, this setting keeps one recovery point per month for a time period that you specify. The default is 2 months. You can choose from months and years. If you do not want to save at least one recovery point monthly, clear this option.
...and then keep one recovery point per year for n [retention time period]	If selected, this setting keeps one recovery point per year for a number of years that you specify. The default (1 year) is disabled by default. If you do want to save at least one recovery point yearly, select this option and specify a number of years. If you do not want to save at least one recovery point yearly, clear this option.

The following is an example of how the retention period is calculated.

Keep all recovery points for three days.

...and then keep one recovery point per hour for three days


...and then keep one recovery point per day for four days

...and then keep one recovery point per week for three weeks
...and then keep one recovery point per month for two months
...and then keep one recovery point per month for one year
In this example, the oldest recovery point would be one year, 3 months old.

9. Click **Save**.
The *Retention Policy* dialog box closes.
10. In the *Nightly Jobs* dialog box, click **OK**.
The *Nightly Jobs* dialog box closes. The retention policy you defined for this machine is applied during the nightly rollup.

Forcing rollup for a protected machine

You can bypass your scheduled retention policy by forcing recovery points to roll up at the protected machine level.

1. From the Protected Machines menu of the Rapid Recovery Core Console, click the name of a specific protected machine.
The *Summary* page for the selection machine appears.
2. Click the **More** drop-down menu at the top of the protected machine view, and then select  **Retention Policy**.
The *Retention Policy* page for the specified machine appears.
3. Click **Force Rollup**.
4. In the dialog box, click **Yes** to confirm.
Rapid Recovery Core initiates rollup for this machine, regardless of the retention policy schedule.

Archiving

This section provides conceptual information about Rapid Recovery archives, including business cases for creating them, storage options, and uses. It also describes how to create a one-time archive, or how to create an archive that is continually updated on a schedule. Topics describe how to pause or edit scheduled archives, how to force or check an archive, and how to attach or import an archive.

Topics include:

- [Understanding archives](#)
- [Creating an archive](#)
- [Editing a scheduled archive](#)
- [Pausing or resuming a scheduled archive](#)
- [Forcing an archive job](#)
- [Checking an archive](#)
- [Attaching an archive](#)
- [Detaching an archive](#)
- [Importing an archive](#)

Understanding archives

The Rapid Recovery Core saves snapshot data to the repository. While a repository can reside on different storage technologies (such as SAN, DAS, or NAS), the most critical performance factor is speed. Repositories use short-term (fast and more expensive) media. To prevent a repository from filling up quickly, the Core enforces a retention policy, which over time rolls up recovery points and eventually replaces them with newer backup data.

If you need to retain recovery points, whether for historical significance, legal compliance, to fulfill offsite data storage policies, or other reasons, you can create an archive. An archive is a copy of recovery points from your repository for the specified machines over a date range that you designate. Archiving a set of recovery points does not delete the original recovery points in your repository. Instead, the archive freezes the collection of recovery points at the point in time in which the archive was created, as a separate copy in a storage location that you specify. Unlike recovery points in your repository, the data in an archive are not subject to rollup.

You can create, import, and attach archives from the  **Archive** option on the button bar, or from the *Archives* page accessible from the  (More) icon on the Core Console.

Related topics:

- [Archive creation and storage options](#)
- [Amazon storage options and archiving](#)
- [Recovery point chain options for archives](#)
- [Methods to access an archive](#)

- [Uses for archives](#)
- [Creating an archive](#)

Archive creation and storage options

You can create a one-time archive on demand at any time.

You can also define requirements for continual scheduled archive. This action creates an archive of recovery points for the machines you select, in the location you designate. Additional recovery points for those machines are then continually appended to the archive on a schedule you define (on a daily, weekly, or monthly basis).

When you create an archive, you specify where you want to save it. You can store an archive in a file system (locally or on a network), or in a storage account in the cloud.

i NOTE: Before archiving to a cloud account, you must first add the credentials to the storage account on your Rapid Recovery Core. For more information about defining a cloud account in the Core, see [Adding a cloud account](#).

If storing your archive in an Amazon cloud storage account, you must define the storage class when creating the account. To archive directly to Amazon Glacier, you can specify Glacier storage when defining the location in the Archive Wizard. For more information about Amazon storage classes, see [Amazon storage options and archiving](#).

- One-time archives are read-only. When creating a one-time archive, the destination location you specify must be empty.
- When using scheduled archive, the Core appends additional recovery points to the existing archive.
- If the storage medium you selected runs out of space, Rapid Recovery pauses the archive job, letting you specify another location. Your archive is then split into segments, which can reside in different locations, as space allows.

Amazon storage options and archiving

When archiving data in an Amazon Simple Storage Service (S3) account, you can choose from various storage classes. Each has different associated costs, benefits, and access restrictions. Prices may differ by region. Rapid Recovery release 6.2 extends support of Amazon storage accounts to all storage classes. This is useful when planning to store Rapid Recovery archives in the Amazon cloud.

Understanding storage classes available and the difference in lead times to access those classes is critical to control storage costs effectively while being able to access your archives within acceptable time frames.

i NOTE: Quest provides this information as a courtesy to its Rapid Recovery customers to help raise awareness of storage pricing factors. These concepts can help you plan and budget accordingly. You are responsible for all storage costs on Amazon or for any other cloud service provider. Since Amazon can change prerequisites, requirements, costs, storage tier definitions, and so on, use the Amazon website as the primary and authoritative source for that information.

In general, Amazon currently offers the following storage classes.

Standard. This storage class is for data you plan to access frequently or access quickly. This class is the default storage option from all Rapid Recovery wizards, windows, and dialog boxes. There are no separate fee to retrieve information in the standard storage class.

Standard Infrequent Access (IA). This storage class is more economical than standard, intended for data that you do not intend to access frequently. There is a retrieval fee associated with accessing data in the storage class. However, availability is immediate. Amazon charges fees for objects deleted from Infrequent Access storage before 30 days.

Glacier. This storage class is for long-term storage of data that does not require real-time access. It is most economical for long-term storage of data that is rarely retrieved. There is a retrieval fee associated with accessing data in this storage class. Amazon charges fees for objects deleted from Glacier before 90 days. Retrieval of data stored in Glacier is not immediate. Standard retrieval times require 3 to 5 hours; bulk retrieval for large amounts of data can take up to 12 hours. Expedited retrievals can take up to 5 minutes. Fees apply to each retrieval option.

NOTE: Unlike the other Amazon storage classes, Rapid Recovery does not support the creation of a cloud account that is specific to Glacier. To archive data in Glacier, choose an Amazon cloud account, and select the **Use Glacier storage** option in the *Location* page of the Archive Wizard.

Reduced Redundancy Storage (RRS). This category is a lower-cost storage class is designed for noncritical reproducible data with less redundancy than the standard storage class. There is a minimal retrieval fee associated with accessing data in this storage class. A fractional percentage (Amazon cites as much as .01%) of objects stored in this class are expected to be unrecoverable.

For the most recent information, always review materials on Amazon's website.

In general, the following guidelines apply:

- If you expect to restore data on any regular basis from a Rapid Recovery archive, Standard is likely to be the best option.
- From a cost perspective, if you plan to restore data occasionally, consider Standard Infrequent Access.
- Glacier is intended for cold storage of archived recovery points from which you rarely expect to restore. A good example of when to use Glacier storage is when saving data for regulatory compliance. Glacier is available as an archive option from the Archive Wizard.
- For storage of noncritical, reproducible data, consider RRS.

Recovery point chain options for archives

Before creating your archive, you must decide on the proper approach for recovery point chains. Use the following information to determine which option you select in the Options page of the Archive Wizard.

- **Build complete recovery point chains, including referenced base images outside the date range.** If you select the option to build complete recovery point chains, then you can perform the full range of restore actions for any recovery point in the archive. This range includes file-level restore, volume-level restore, and bare metal restore. When you select this option, full recovery point chains are saved with your archive. You can restore data even if the base image corresponding to the selected recovery point is earlier than the date range of your archive. However, the file size of this archive is larger, to ensure that you have access to data in the full recovery point chain.
- **Include only the recovery points in the date range. This saves space, but you are responsible for archiving any needed base images.** If you include only the recovery points in the specified date range in your archive, the file size of the archive is smaller. For data in which the base image is included in the date range you specified, you have access to the full range of restore options. However, if you want to recover data captured in a base image from a date earlier than the date range you specified, you may be restricted to file-level recovery only. Data outside the range of the archive is orphaned.

For more information on recovery point chains, see the topic [Recovery point chains and orphans](#).

Methods to access an archive

When you need to access the data in an archived recovery point, you have two options.

- For archives created with Rapid Recovery Core version 6.0.1 and later, you can *attach* the archive. The attached archive is displayed in the left navigation menu of the Core Console. You can browse the recovery points in the archive, and take the same actions on that data as with any other recovery points currently in your repository, without importing that data into your repository.
- You can *import* an archive, restoring those recovery points to your repository. You can then take the same actions on that data as with any other recovery points currently in your Core. Rapid Recovery Core is backward compatible, supporting import of archives from all AppAssure and Rapid Recovery versions.

CAUTION: Since the Core recognizes the original dates of recovery points in an archive, recovery points imported from an archive may be rolled up or deleted during the next nightly job period, if their age exceeds the retention period. If you want to retain older recovery points imported from an archive, you can disable rollup or extend the retention period for the relevant protected machines.

When you need to access the data in an archived recovery point, you can attach (for Rapid Recovery 6.x and later) or import the archive, restoring those recovery points to your repository.

Uses for archives

Once an archive is created, it can be used in the following ways:

- An archive can be used to move data between repositories.
- An archive can be attached to the Core Console, and mounted as a file system for simple file or folder recovery. Some restrictions apply.
- An archive can be imported into a repository. If any recovery points in the archive have already been rolled up or deleted from the repository, this action restores them.
- An archive can be used as the source for a bare metal restore, or for export to a virtual machine.

Creating an archive


You can use this procedure to create a one-time or scheduled archive.

If you plan on creating an archive to a cloud location, first add your cloud account to the Rapid Recovery Core Console. For more information, see [Adding a cloud account](#).

A one-time archive is an archive created on demand for a specified machine. A scheduled archive is an archive that automatically recurs on the date and time you specify in the wizard. When you schedule a recurring archive, you meet the need for archiving data from protected machines frequently without the inconvenience of repeatedly defining a one-time archive.


Before creating a scheduled archive, ensure that you are satisfied with the name of each folder and subfolder in the archive structure. After performing an initial archive, Quest recommends making no changes to the resulting folder structure. Changing any folder name in the archive path may result in capturing a base image instead of an incremental image. It could also cause difficulty restoring or importing data from an archive due to an inability to properly locate the files.

When creating your archive, you must decide whether to include a full recovery point chain in your archive. For more information, see [Recovery point chain options for archives](#).

1. On the button bar of the Rapid Recovery Core Console, click  **Archive**. The Archive Wizard opens.

2. On the *Archive Type* page of the wizard, select one of the following options:
 - One-time archive
 - Continual archive (by schedule)
3. Click **Next**.
4. On the *Location* page, select an option from the **Location type** drop-down list and then enter the information as described in the following table.

Table 154: Archive location type options

Option	Text Box	Description
Local	Location	Enter the local path where you want the archive to reside; for example, d:\work\archive.
Network	Location	Enter the network path where you want the archive to reside; for example, \\servername\sharename.
	User name	Enter the user name for the user with access to the network share.
	Password	Enter the password for the user with access to the network share.
Cloud	Account	Select an account from the drop-down list. <div>  NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account. </div>
	Container	Select a container associated with your account from the drop-down menu.
	Folder name	Enter a name for the folder in which the archived data is saved.
	Use Glacier storage	To store your archive in Amazon Glacier storage, select this option. This option is intended for long-term storage of archives. For more information, see Amazon storage options and archiving .

5. Click **Next**.
6. On the *Machines* page of the wizard, select the protected machine or machines you want to archive and then click Next.
7. Do one of the following:
 - If you chose to create a one-time archive, skip to [step 14](#).
 - If you chose to create a scheduled archive, continue to [step 8](#).
8. On the *Schedule* page, select one of the following options from the **Send data** drop-down list:
 - Daily
 - Weekly
 - Monthly

9. Enter the information described in the following table based on your selection from [step 8](#).

Table 155: Send data options

Option	Text Box	Description
Daily	At time	Select the hour of the day you want to create a daily archive.
Weekly	At day of week	Select a day of the week on which to automatically create the archive.
	At time	Select the hour of the day you want to create an archive.
Monthly	At day of months	Select the day of the month on which to automatically create the archive.
	At time	Select the hour of the day you want to create an archive.

10. Optionally, if you do not want the archive job to begin at the next scheduled time after you complete the wizard, select **Pause initial archiving**.

i **NOTE:** You may want to pause the scheduled archive if you need time to prepare the target location before archiving resumes. If you do not select this option, archiving begins at the scheduled time.

11. Click **Next**.
12. On the *Options* page for a continuous archive, select one of the recycle actions described in the following table.

Table 156: Continuous archive recycle options

Text Box	Description
Replace this Core	Overwrites any pre-existing archived data pertaining to this Core but leaves the data for other Cores intact.
Erase completely	Clears all archived data from the directory before writing the new archive.
Incremental	Lets you add recovery points to an existing archive. It compares recovery points to avoid duplicating data that exist in the archive.

13. On the *Options* page for a continuous archive, determine whether to include full recovery point chains in your archive. For more information about recovery point chain options, see [Recovery point chain options for archives](#). Do one of the following:
- Select **Build complete recovery point chains, including referenced base images outside the date range**, and then skip to [step 17](#).
 - Select **Include only the recovery points in the date range**. This option saves space, but you are responsible for archiving any needed base images, and then skip to [step 17](#).

14. On the *Options* page for a one-time archive, enter the information described in the following table.

Table 157: One-time archive options

Text Box	Description
Maximum size	<p>Large archives of data can be divided into multiple segments. Select the maximum amount of space you want to reserve for creating the archive by doing one of the following:</p> <ul style="list-style-type: none">• Select Entire Target to reserve all available space in the path provided on the destination provided in step 4. (For example, if the location is <code>D:\work\archive</code>, all the available space on the D: drive is reserved).• Select the text box, enter the amount of space, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve. <p>NOTE: Amazon cloud archives are automatically divided into 50 GB segments. Microsoft Azure cloud archives are automatically divided into 200 GB segments.</p>
Recycle action	<p>Select one of the following recycle action options:</p> <ul style="list-style-type: none">• Do not reuse. Does not overwrite or clear any existing archived data from the location. If the location is not empty, Rapid Recovery lets you select a different location.• Replace this Core. Overwrites any pre-existing archived data pertaining to this Core but leaves the data for other Cores intact.• Erase completely. Clears all archived data from the directory before writing the new archive.• Incremental. Lets you add recovery points to an existing archive. It compares recovery points to avoid duplicating data that already exist in the archive.
Comment	<p>Enter any additional information that is necessary to capture for the archive. The comment is displayed if you import the archive later.</p>
Build complete recovery points chains, including referenced base images outside the date range.	<p>Select this option to archive the entire recovery point chain. This option is selected by default.</p>
Include only the recovery points in the date range. This saves space, but you are responsible for archiving any needed base images.	<p>Select this option to archive only the incremental recovery points and no base images.</p>

Text Box	Description
	<p>i NOTE: This option results in an archive of orphaned recovery points. You will not be able to use them for recovery until you also archive their related base images.</p> <p>For more information about recovery point chain options, see Recovery point chain options for archives.</p>

15. If you selected the option to include only recovery points in the date range, click **Next**. Otherwise, click **Finish**.
16. On the *Date Range* page, either manually enter the start date and end date of the recovery points to be archived, or select the date time by clicking the calendar icon followed by the clock icon below the calendar window.
17. Click **Finish**.
The wizard closes.

Archiving to a cloud

When data reaches the end of a retention period, you may want to extend that retention by creating an archive of the aged data. When you archive data, there is always the matter of where to store it. Rapid Recovery lets you upload your archive to a variety of cloud providers directly from the Core Console. Compatible clouds include Microsoft Azure, Amazon S3, any OpenStack-based provider, Rackspace Cloud Files, and Google Cloud.

Exporting an archive to a cloud using Rapid Recovery involves the following procedures:

- Add your cloud account to the Rapid Recovery Core Console. For more information, see [Adding a cloud account](#).
- Archive your data and export it to your cloud account. For more information, see [Creating an archive](#).
- Retrieve archived data by attaching an archive, or importing it from the cloud location. For more information, see [Attaching an archive](#) or [Importing an archive](#), respectively.

Editing a scheduled archive



Rapid Recovery lets you change the details of a scheduled archive. To edit a scheduled archive, complete the steps in the following procedure.

1. In the Rapid Recovery Core Console, click the **☰ (More)** drop-down menu on the icon bar, and then select **Archives**.
2. On the *Archives* page, under Scheduled Archives, click the drop-down menu next to the archive you want to change, and then click **Edit**.
The Archive Wizard opens.

3. On the *Location* page of the Archive Wizard, select one of the following options from the **Location Type** drop-down list:
 - Local
 - Network
 - Cloud

4. Enter the details for the archive as described in the following table based on the location type you selected in [step 3](#).

Table 158: Archive details

Option	Text Box	Description
Local	Location	Enter the local path where you want the archive to reside; for example, <code>d:\work\archive</code> .
Network	Location	Enter the network path where you want the archive to reside; for example, <code>\\servername\sharename</code> .
	User name	Enter the user name for the user with access to the network share.
	Password	Enter the password for the user with access to the network share.
Cloud	Account	Select an account from the drop-down list. <div>  NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account. </div>
	Container	Select a container associated with your account from the drop-down menu.
	Folder name	Enter a name for the folder in which the archived data is saved.
	Use Glacier storage	To store your archive in Amazon Glacier storage, select this option. This option is intended for long-term storage of archives. For more information, see Amazon storage options and archiving . <div>  NOTE: This option will cause all new archives to be placed in Glacier storage. The existing archives will remain in S3 storage. </div>

5. Click **Next**.
6. On the *Machines* page of the wizard, select the protected machine or machines that contain the recovery points that you want to archive. Clear the machines that you do not want to archive.
7. Click **Next**.
8. On the *Schedule* page, select one of the following options from the **Send data** drop-down list:
 - Daily
 - Weekly
 - Monthly

9. Enter the information described in the following table based on your selection from [step 8](#).

Table 159: Send data options

Option	Text Box	Description
Daily	At time	Select the hour of the day you want to create a daily archive.
Weekly	At day of week	Select a day of the week on which to automatically create the archive.
	At time	Select the hour of the day you want to create a daily archive.
Monthly	At day of months	Select the day of the month on which to automatically create the archive.
	At time	Select the hour of the day you want to create a daily archive.

10. Optionally, to postpone archiving to resume at a later time, select **Pause initial archiving**.

i **NOTE:** You may want to pause the scheduled archive if you need time to prepare the target location before archiving resumes. If you do not select this option, archiving begins at the scheduled time.

11. Click **Next**.
12. On the *Options* page, use the **Recycle action** drop-down list to select one of the options described in the following table:

Table 160: Archive recycle options


Text Box	Description
Incremental	Lets you add recovery points to an existing archive. It compares recovery points to avoid duplicating data that already exists in the archive.
Replace this Core	Overwrites any pre-existing archived data pertaining to this Core but leaves the data for other Cores intact.
Erase completely	Clears all archived data from the directory before writing the new archive.

13. On the *Options* page of the wizard, determine whether to include full recovery point chains in your archive. For more information about recovery point chain options, see [Recovery point chain options for archives](#). Do one of the following:
- Select **Build complete recovery point chains, including referenced base images outside the date range**.
 - Select **Include only the recovery points in the date range**. This saves space, but you are responsible for archiving any needed base images.
14. Click **Finish**.
Rapid Recovery applies your changes to the archive.

Pausing or resuming a scheduled archive

If you have an archiving job scheduled to recur, you can pause or resume this action as necessary.


There may be times when you want to pause a scheduled archive job, such as if you need to change the destination archive location. Also, if you opted to initially pause archiving when you performed the [Creating an archive](#) procedure, you likely want to resume the scheduled archive later. Complete the steps in the following procedure to pause or resume scheduled archive.

1. From the Rapid Recovery Core Console, click the **⋮** (More) menu on the icon bar, and then click  **Archives**.
2. On the Archives page, under Scheduled Archives, do one of the following:
 - Select the preferred archive, and then click one of the following actions as appropriate:
 - Pause
 - Resume
 - Next to the preferred archive, click the drop-down menu, and then click one of the following actions as appropriate:
 - Pause
 - Resume

The status of the archive displays in the Schedule column.

Forcing an archive job

Using this procedure, you can force Rapid Recovery to perform the archive job on a scheduled archive at any time. To force an archive job, you must have an archive scheduled on the Core.

1. From the Rapid Recovery Core Console, in the icon bar, click the **⋮** (More) menu on the icon bar, and then click  **Archives**.
2. On the *Archives* page, under Schedule Archives, click the drop-down menu next to the archive you want to force, and then click **Force**.
Rapid Recovery archives the recovery points according to the settings you chose for that archive, regardless of the scheduled archive time you set.

Checking an archive

Checking an archive verifies whether an archive and its contents are healthy enough to be restored.

You can scan an archive for the integrity of its structure, data segments, and index files by performing an archive check. The archive check verifies the presence of all necessary files within the archive and that the files are healthy. To perform an archive check, complete the steps in the following procedure.





1. From the Rapid Recovery Core Console, in the icon bar, click the  (More) menu on the icon bar, and then click  **Archives**.
2. On the Archives page, click  **Check**.
The Check Archive dialog box appears.
3. For Location type, select one of the following options from the drop-down list:
 - Local
 - Network
 - Cloud
4. Based on the location type you selected in [step 3](#), enter the details for the archive as described in the following table.

Table 161: Archive details

Option	Text Box	Description
Local	Location	Enter the local path for the archive.
Network	Location	Enter the network path for the archive.
	User name	Enter the user name. It is used to establish logon credentials for the network share.
	Password	Enter the password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list.  NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account .
	Container	Select a container associated with your account from the drop-down menu.
	Folder name	Select the folder in which the archived data is saved; for example, Rapid-Recovery-7-Archive-[DATE CREATED]-[TIME CREATED].

5. Select or clear the checks described in the following table. All are selected by default.



NOTE: Do not clear all checks. You must select at least one option.

Option	Description
Index files mapping offsets	This option checks that all the data on the internal structure of the archive is in the correct location.
Structure integrity	This option verifies the presence of certain internal files and the folder structure of the archive. If any files or folders are missing, the check fails.
Checksum integrity	This option checks the integrity of the data segments in the archive to ensure that the segments are healthy.

6. Click **Check File**.
Rapid Recovery checks the archive according to your selections.

Attaching an archive

Attaching an archive lets you see recovery points from the archive.

You must have a pre-existing archive created in Rapid Recovery Core release 6.0.1 or later to complete this procedure. For more information, see [Creating an archive](#).

When you attach an archive, the archive name you provide appears as an archive menu in the left navigation menu of the Core Console. Each protected machine with recovery points in the archive is listed separately below the archive menu. You can click any machine name in the archive and browse its recovery points. You can then take the same actions as with any other recovery points currently visible in your Core.

Attaching the archive also caches the credentials for accessing the information. Unless you delete the attached archive permanently, you can easily detach and re-attach an archive, making its recovery points easily accessible without cluttering the Core Console. For information about detaching an archive, see [Detaching an archive](#).

Use this procedure to attach an archive.




1. On the Rapid Recovery Core Console, click the  **Archive** ▾ drop-down menu, and then select  **Attach Archive**.
The Attach Archive dialog box appears.
2. In the **Name** text box, enter a name for this attached archive.
The value you type in this field appears in the left navigation menu as the archive menu name.
Following best practice for display names, the archive name should contain between 1 and 64 alphanumeric characters, including spaces. Do not use [prohibited characters](#) or [prohibited phrases](#).
3. In the Location type drop-down list, select the location type of your archive from the following options:
 - Local
 - Network
 - Cloud
4. Enter the details for the archive as described in the following table based on the location type you selected in [step 3](#).

Table 162: Location type details

Option	Text Box	Description
Local	Location	Enter the path to the archive; for example, D:\Work\Archive.
Network	Location	Enter the path to the archive; for example, \\servername\sharename.
	User name	Enter user name for logging in to the network share.
	Password	Enter the password for logging in to the network share.
Cloud	Account	Select an account from the drop-down list. <div> NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account.</div>
	Container	Select the container of the archive associated with your account from the drop-down menu.
	Folder name	Enter the name of the folder of the archived data; for example, Rapid-Recovery-Archive-[DATE CREATED]-[TIME CREATED].

5. Click **Attach**.

The archive attaches to this Core and mounts the contents as a file system.

Detaching an archive

You must have an archive attached to your Core to perform this task.

When an archive is attached to your Core, you can navigate through the data in the archive as if the data were currently protected machines. Once you attach an archive, its credentials are cached and it remains visible in your Core until you detach it.

Perform this task after you have accomplished your objectives for attaching the archive, to remove it from view on your Core.



NOTE: Detaching the archive does not delete the data; it only removes the data from current view.

Use this procedure to detach an archive.

1. On the left navigation menu of the Rapid Recovery Core Console, locate the archive you want to detach.
2. Click on the ellipsis for the archive, and from the context-sensitive menu, select **Delete**.
3. In the resulting dialog box, confirm that you want to remove the selected archive.
The dialog box closes. After a brief pause, the attached archive is removed from the Core Console. This action is logged as an alert.

Importing an archive

You can use this procedure to import an archive one time, or schedule an archive to import on a recurring basis.

When you want to recover archived data, you can import the entire archive to a specified location.




CAUTION: Perform this step only after careful consideration. Importing an archive repopulates the repository with the contents of the archive, replacing any new data in the repository since the archive was captured.

To import an archive, complete the steps in the following procedure.

1. On the menu bar of the Rapid Recovery Core Console, click the **Archive** drop-down menu, and then select **Import Archive**.
The Import Archive Wizard opens.
2. On the Import Type page of the wizard, select one of the following options:
 - One-time import
 - Continual import (by schedule)
3. Click **Next**.

4. On the *Location* page, select the location of the archive you want to import from the drop-down list, and then enter the information as described in the following table:

Table 163: Imported archive location type options

Option	Text Box	Description
Local	Location	Enter the local path where you want the archive to reside; for example, d:\work\archive.
Network	Location	Enter the network path where you want the archive to reside; for example, \\servername\sharename.
	User name	Enter the user name for the user with access to the network share.
	Password	Enter the password for the user with access to the network share.
Cloud	Account	Select an account from the drop-down list. <div> NOTE: To select a cloud account, you must first have added it in the Core Console. For more information, see Adding a cloud account.</div>
	Container	Select a container associated with your account from the drop-down menu.
	Folder name	Enter a name for the folder in which the archived data is saved.

5. Click **Next**.
6. On the *Archive Information* page of the wizard, if you want to import every machine included in the archive, select Import all machines.
7. Complete one of the following options based on your selection:
 - If you selected **One-time import** in [step 2](#), you selected **Import all machines** in [step 6](#), and all the machines are present on the Core—as protected, replicated, or recovery points only machines—go to [step 12](#).
 - If you did not import all machines in [step 6](#), click **Next**, and then continue to [step 8](#).
8. On the *Machines* page, select the machines that you want to import from the archive.
 - If you selected **One-time import** in [step 2](#), and at least one machine is not present on the Core—as a protected, replicated, or recovery points only machine—use the drop-down lists to select a repository for each machine you want to import, and then go to [step 12](#).
 - If all machines are already present on the Core—as protected, replicated, or recovery points only machines—go to [step 12](#).
9. Click **Next**.

10. On the Repository page, complete one of the following options:

- If a repository is associated with the Core, select one of the options in the following table.

Table 164: Repository options

Option	Description
Use an existing Repository	Select a repository currently associated with this Core from the drop-down list.
Create a Repository	In the Server text box, enter the name of the server on which you want to save the new repository—for example, <code>servername</code> or <code>localhost</code> —and then see Creating a DVM repository .

- If no repository is associated with the Core, enter the name of the server on which you want to save the new repository—for example, `servername` or `localhost`—and then see [Creating a DVM repository](#) or [Connecting to an existing repository](#).

11. If you chose to **Continuous import (by schedule)** in [step 2](#), on the *Schedule* page, select the options described in the following table.

Table 165: Schedule import options

Option	Description
Daily	Click the clock icon and use the up and down arrows to select at what time you want to the archive job to begin. If you are using a 12-hour time system, click the AM or PM button to specify the time of day.
Weekly	Select the day of the week and then the time you want the archive job to begin. If you are using a 12-hour time system, click the AM or PM button to specify the time of day.
Monthly	Select the day of the month and the time you want the archive job to begin. If you are using a 12-hour time system, click the AM or PM button to specify the time of day.
Pause initial importing	Select this option if you do not want the import job to begin at the next scheduled time after you complete the wizard. i NOTE: You may want to pause the scheduled import if you need time to prepare the target location before importing resumes. If you do not select this option, importing begins at the scheduled time.

12. Click **Finish**.

Cloud accounts

Rapid Recovery lets you define connections between existing cloud storage or cloud service providers and your Rapid Recovery Core. Compatible cloud services include Microsoft Azure, Amazon Web Services (AWS), any OpenStack-based provider (including Rackspace), and Google Cloud. US government-specific platforms include AWS GovCloud (US) and Azure Government. You can add any number of cloud accounts to the Core Console, including multiple accounts for the same provider.

The purpose of adding cloud accounts to your Core Console is to work with them as described in the topic [About cloud accounts](#).

Once added, you can manage the connection between the Core and the cloud accounts. Specifically, you can edit the display name or credential information, configure the account connection options, or remove the account from Rapid Recovery. When you edit or remove cloud accounts in the Core Console, you do not change the cloud accounts themselves—just the linkage between those accounts and your ability to access them from the Core Console.

This section describes how to define links between existing cloud storage provider or cloud service provider accounts, and the Rapid Recovery Core Console. It also describes how to manage those cloud accounts in Rapid Recovery.

Topics include:

- [About cloud accounts](#)
- [Considering cloud storage options](#)
- [Adding a cloud account](#)
- [Editing a cloud account](#)
- [Removing a cloud account](#)

About cloud accounts

Rapid Recovery works with cloud accounts in the following ways:

- **Archive.** You can store a one-time archive or continual scheduled archive in the cloud. This feature is supported for all supported cloud account types. When you archive to the cloud, you can later access the information in the archived recovery points by attaching the archive (for archives created in release 6.0.1 or later). For all archives (including archives created prior to Rapid Recovery 6.0.1), you can import the archive (which then makes the imported data subject to your data retention policy). You can also perform bare metal restore from an archive.
- **Virtual export.** You can perform virtual export to an Azure cloud account. This includes one-time export of a virtual machine, or continual export for a virtual standby VM.

i **NOTE:** When archiving to Azure using Rapid Recovery release 6.8, use the cloud account type Microsoft Azure Service Management (for Archive). When exporting a VM to Azure, use the cloud account type Microsoft Azure Resource Management (for Virtual Export).

For conceptual information regarding various cloud accounts, see [Considering cloud storage options](#).

For information about configuring timeout settings between the Core and cloud accounts, see [Configuring cloud account connection settings](#).

For information about performing virtual export to the Azure cloud, see [Exporting data to an Azure virtual machine](#).

For information about adding a cloud account, see [Adding a cloud account](#).

Considering cloud storage options

This topic discusses support for US Government cloud storage accounts. It also discusses tradeoffs between cost and other factors when selecting cloud accounts for archiving.

Secure cloud accounts for US Government

United States federal, state, and local government agencies and their partners have access to increasing cloud account options. Rapid Recovery supports the following offerings for Government and related cloud accounts:

- **AWS GovCloud (US).** Amazon Web Services offers a service called AWS GovCloud (US). This is an isolated AWS region designed to meet specific regulatory and compliance requirements. Using this service lets United States government agencies and customers join private businesses in leveraging cloud accounts. Rapid Recovery supports archiving to Amazon S3 storage accounts in the AWS GovCloud.
- **Azure Government.** Azure Government is a United States government-only cloud platform exclusively for US federal, state, local, and tribal government agencies and their partners. Rapid Recovery supports Azure Government in the same manner that we offer standard Azure support. For example:
 - Rapid Recovery supports archiving to Azure Government and standard Azure storage accounts.
 - Rapid Recovery supports virtual export to an Azure virtual machine (one-time, or virtual standby) on public and Azure Government platforms.
 - Rapid Recovery supports running a Rapid Recovery Core in an Azure VM in the AWS GovCloud or in a standard Azure account.
 - Rapid Recovery supports replication from a source Core (running on-premises or in the AWS GovCloud) to an Azure VM target Core. If your source Core is located in the AWS GovCloud, your replication target must also run in the AWS GovCloud.

Balancing access time, cost and convenience for archiving to cloud accounts

To offer our users cost-effective cloud archiving and virtual export options, Rapid Recovery continues to expand support of cloud storage providers (and storage classes for leading providers that offer them). Educated users can leverage policies to balance data archive convenience, data access time, and cost.

When considering strategies for archiving or exporting to the cloud, Rapid Recovery users are encouraged to understand the tradeoffs between initial cost to store data, how frequently the data is expected to be used, the need to access that data within a prescribed period of time, and costs associated with retrieving the data.

Some providers (such as Amazon S3) offer different storage classes. Choosing the correct storage class can save you money if your assumptions regarding these factors are accurate. Quest recommends that Rapid Recovery

users review data storage policies at least once annually to ensure you are using your resources effectively. Similarly, administrators are cautioned to review the data being archived or exported to cloud accounts so you can update planning assumptions and migrate data accordingly.

The act of storing data, for some vendors, is extremely low or in some cases free. However, cloud service providers often apply charges to your account when you access or retrieve that data. There are often different fees based on how quickly you need to access the data. In some cases, using more expensive storage (such as Amazon S3 standard) is more cost effective if you plan to restore from recovery points than if you store data in Glacier and need to restore.

Amazon lets you define data life cycle policies that move data between Amazon S3 storage classes over time. For example, you could store freshly uploaded data using the Standard storage class, move it to Standard – Infrequent Access 30 days later, and then to Reduced Redundancy Storage after another 60 days have passed. You can also explicitly archive data for any type of Amazon S3 cloud account to Glacier, using the Archive Wizard. This is recommended if data recovery is expected very infrequently. Before selecting this option, familiarize yourself with fees related to access, storage age, and so on. See the topic [Amazon storage options and archiving](#).

Some Rapid Recovery features are designed specifically for the cloud. If performing virtual export to the cloud using Azure, consider virtual standby. This process lets you create a fully bootable virtual machine in the Azure cloud. The VM files are continually updated with newly captured recovery points. Unlike virtual standby performed on-premises, the VM files are not deployed into a bootable VM until or unless you need them. Your initial cost for virtual standby in Azure involve only storage. Compute costs (which in Azure can be considerable in the long term) are incurred only if the VM is deployed, which is required to spin up a VM and perform a restore.

You can run a Rapid Recovery Core in an Azure VM. You can also replicate an on-premises Core to a VM in the Azure cloud, or replicate a source Core in Azure to a target Core in Azure. Running a source or target Rapid Recovery Core in Azure uses compute resources for the active Core VM, and requires storage accounts to be created and associated with each Core VM for your repository, which incurs storage costs. For information about setting up a Core to run in Azure, see the *Rapid Recovery Azure Setup Guide*.

Users of Rapid Recovery that employ cloud storage options are encouraged to understand the tradeoffs between initial cost to store data, the need to access that data within a prescribed period of time, and costs associated with retrieving the data.



For example, the act of storing data, for some vendors, is extremely low or in some cases free. However, cloud service providers often apply charges to your account when you access or retrieve that data. There are often different fees based on how quickly you need to access the data. In some cases, using more expensive storage (such as Amazon S3 standard) is more cost effective if you plan to restore from recovery points than if you store data in Glacier and need to restore.

Amazon lets you define data life cycle policies that move data between Amazon S3 storage classes over time. For example, you could store freshly uploaded data using the Standard storage class, move it to Standard – IA 30 days later, and then to Amazon Glacier after another 60 days have passed.

Adding a cloud account

Before you can move data in either direction between a cloud account and your Core, you must add cloud provider account information to the Rapid Recovery Core Console. This information identifies the cloud account in the Core Console while caching the connection information securely. This process then lets Rapid Recovery Core connect to the cloud account to perform the operations you specify.

To add a cloud account, complete the steps in the following procedure.





1. On the Rapid Recovery Core Console icon bar, click the  (More) icon and then select  **Cloud Accounts**.
The *Cloud Accounts* page appears.
2. On the *Cloud Accounts* page, click **+ Add New Account**.
The *Add New Account* dialog box opens.
3. Select a compatible cloud provider from the Cloud type drop-down list.

4. Enter the details described in the following table based on the cloud type selected in step 3. Since the table is listed and organized by cloud type, the Cloud Type parameter is listed in the first column instead of the Text Box column.

i **NOTE:** In Rapid Recovery release 6.8, two Microsoft Azure cloud types are included on the *Cloud Accounts* page. Azure Service Management (ASM) specifically supports archiving to Azure. Azure Resource Management (ARM) supports virtual export to Azure. If using both archiving and virtual export to Azure in Rapid Recovery Core release 6.3, you are required to set up a cloud account for each.

Table 166: Cloud account details

Cloud Type	Text Box	Description
Amazon S3	Display name	Enter a display name for this cloud account to display on the Rapid Recovery Core Console; for example, Amazon S3 Cloud Account 1.
	Access key	Enter the access key for your Amazon cloud account.
	Secret key	Enter the secret key for this account.
	Service endpoint	Optionally, if using an S3-compatible storage account other than a standard Amazon S3 storage account, enter the fully qualified http or https URL for that storage account.
	Storage class	Select a storage class for the S3 account. You can choose from: <ul style="list-style-type: none">• Standard• Standard - Infrequent Access• Reduced Redundancy Storage <p>If you want to archive to Glacier, you can define your Amazon cloud account using any listed storage class. The option to select Glacier storage is accessible from the Archive Wizard.</p> <p>i NOTE: For more information on storage classes, see the topic Considering cloud storage options.</p>
Google Cloud	Display name	Enter a display name for this cloud account to display on the Rapid Recovery Core Console; for example, Google Cloud Account 1.
	Certificate file	Browse for and select your Google certificate file to authenticate this cloud account.
	Private key	Enter your private key for this account.
	Project ID	Enter the Project ID associated with this account.
	Service account email	Enter the email address registered with Google Cloud as the owner of this cloud service account.
Microsoft Azure Service	Display name	Enter a display name for this cloud account to display on the Rapid Recovery Core Console; for example, Azure Cloud Account 1.

Cloud Type	Text Box	Description
Management (for Archive)	Storage account name	<p>Enter the name of your Microsoft Azure storage account.</p> <p>NOTE: The name must match the storage account name in Azure precisely. It must contain lower case letters and numbers only, and be between 3 and 24 characters in length.</p>
	Access key	<p>Enter the access key for your account.</p> <p>NOTE: You can enter the primary or secondary key. To obtain the access key from your Azure account, check Keys under Settings.</p>
	Account type	Choose your Azure account type; for example Azure, Azure Germany, Azure China, or (US) Government.
	Use https protocol	Select this option to use the secure https protocol instead of the standard http protocol.
Microsoft Azure Resource Management (for Virtual Export)	Display name	Enter a display name for this cloud account to display on the Rapid Recovery Core Console; for example, Azure Cloud Account 1.
	Region	Select the appropriate region for your Azure account. For example, select from Azure Global Cloud, Azure China Cloud, Azure German Cloud, Azure US Government Cloud, and so on.
	Tenant ID	<p>Enter the tenant ID precisely. This is an alphanumeric string (also called the Directory ID) associated with your Azure Active Directory application.</p> <p>To obtain this value from the Azure UI, select Azure Active Directory > Properties > Directory ID.</p>
	Application ID	<p>Enter the application ID for your Azure AD application precisely.</p> <p>To obtain this value from the Azure UI, select Azure Active Directory > App registrations, select your application, and from the Settings pane, copy the Application ID.</p>
	Secret key	<p>Enter the secret key for this account.</p> <p>You must obtain this value from the Azure when you set up the key. If you do not record it, you must create a new secret key. From the Azure UI, to see or create secret keys, select Azure Active Directory > App registrations, select your application, click  Settings, and from the <i>Settings</i> pane, click  Keys.</p>
	Subscription ID	<p>Enter the subscription ID for your Azure account precisely.</p> <p>To obtain this value from the Azure UI, select  All services, click  Subscriptions, and from the appropriate subscription, copy the Subscription ID.</p>




Cloud Type	Text Box	Description
Powered by OpenStack	Display name	Enter a display name for this cloud account to display in the Rapid Recovery Core Console; for example, OpenStack Cloud Account 1.
	Region	Enter the region for your cloud account.
	User name	Enter the user name for your OpenStack-based cloud account.
	Password or API key	Select whether to use a password or an API key, and then enter your selection for this account.
	Tenant ID	Enter your tenant ID for this account.
	Authentication URL	Enter the authentication URL for this account. This the base URL for the cloud instance. If not provided, Rapid Recovery Core uses the default URL, https://identity.api.rackspacecloud.com .
Rackspace Cloud Files	Display name	Enter a display name for this cloud account to display on the Rapid Recovery Core Console; for example, Rackspace Cloud Account 1.
	Region	Use the drop-down list to select the region for your account.
	User name	Enter the user name for your Rackspace cloud account.
	Password or API key	Select whether to use a password or an API key, and then enter your selection for this account.
	Tenant ID	Enter your tenant ID for this account.
	Authentication URL	Enter the authentication URL for this account. This the base URL for the cloud instance. If not provided, Rapid Recovery Core uses the default URL, https://identity.api.rackspacecloud.com .
Backblaze B2	Display name	Enter a display name for this cloud account to display in the Rapid Recovery Core Console; for example, Backblaze B2.
	Key ID	Enter the Access Key ID or Access Key
	Application Key	Enter the Secret Access Key or Secret Key
	Service endpoint	Enter the S3 endpoint : https://s3.us-east-005.backblazeb2.com/


5. Click **Save**.

The dialog box closes, and your account appears on the Cloud Accounts page of the Core Console.

Editing a cloud account




If you need to change the information to connect to your cloud account, for example to update the password or edit the display name, you can do so on the Cloud Accounts page of the Rapid Recovery Core Console. Complete the steps in the following procedure to edit a cloud account.


1. On the Rapid Recovery Core Console icon bar, click  (More) icon and then select  **Cloud Accounts**. The Cloud Accounts page appears.
2. Next to the cloud account you want to edit, click the  drop-down menu, and then select **Edit**. The Edit Account window opens.
3. Edit the details as necessary, and then re-enter the password (or API key, secret key, private key, and so on) required to connect to the cloud account. Then click **Save**.

 **NOTE:** You cannot edit the cloud type.

Removing a cloud account

If you discontinue your cloud service, or decide to stop using it for a particular Core, you may want to remove your cloud account from the Core Console. Complete the steps in the following procedure to remove a cloud account.

1. On the Rapid Recovery Core Console icon bar, click the  (More) icon and then select  **Cloud Accounts**. The Cloud Accounts page appears.
2. Next to the cloud account you want to edit, click the  drop-down menu, and then select **Remove**.
3. In the Delete Account dialog box, click **Yes** to confirm that you want to remove the account.
4. If the cloud account is currently in use, a second dialog box prompts you to confirm that you still want to remove it. Click **Yes** to confirm.

 **CAUTION:** Removing an account that is currently in use causes all archive jobs scheduled for this account to fail.

The dialog box closes, and the credentials for accessing the specified cloud account are removed from your Core.

Core Console references

This appendix includes reference tables that describe many of the functions and icons available in the Rapid Recovery Core Console. It serves as a supplement to the [The Core Console](#) chapter of the *Rapid Recovery User Guide*.

Topics include:


[Viewing the Core Console user interface](#)








[Viewing protected machines](#)


Viewing the Core Console user interface

The Core Console is the main user interface (UI) through which users interact with Rapid Recovery. When you log into the Rapid Recovery Core Console, you see the following elements.

Table 167: UI elements included in the Core Console

UI Element	Description
Branding area	For typical environments, the top left side of the Core Console is branded with the full product name, including the Quest logo. Clicking anywhere on the branding area results in the opening of a new tab in the web browser, displaying product documentation on the Support website.
Button bar	The button bar, which appears to the right of the branding area, contains buttons accessible from anywhere in the Core Console. These buttons launch wizards to accomplish common groups of tasks such as protecting a machine; performing a restore from a recovery point; creating, attaching, or importing an archive; or replicating from this source Core to a target Core. Each button in the button bar is further described in the Button bar table below.
Running tasks count	Shows how many jobs are currently running. This value is dynamic based on the system state. When you click the drop-down menu, you see a status summary for all jobs currently running. By clicking the X for any job, you can choose to cancel that job.
Help drop- 	The Help menu includes the following options:

UI Element	Description	
down menu		Help. Links to in-product help, which opens in a separate browser window.
		Documentation. Links to Rapid Recovery technical documentation on the Quest website.
		Interactive training. Links to interactive web-based training on the Quest website.
		Support. Links to the product support page for Rapid Recovery on the Quest website.
		Quick Start Guide. Opens a guided flow of suggested tasks for configuring and using Rapid Recovery Core.
		Submit an idea. If you have suggestions for features or enhancements, select this option to open the Quest ideation portal, where you can quickly and easily describe your ideas and have them reviewed by product managers for potential inclusion in a future release of Rapid Recovery.
		About. Opens the <i>About Rapid Recovery</i> dialog box. Information here includes the Rapid Recovery component; version information; the Core ID; a hyperlink for third-party contributions; and a description of the software.
Server date and time	The current time of the machine running the Rapid Recovery Core service appears at the top right of the Core Console. When you hover your mouse over the time, the server date also appears. This is the date and time recorded by the system for events such as logging, scheduling, and reporting. For example, when applying protection schedules, the time displayed on the Core Console is used. This is true even if the time zone is different on the database server or on the client machine where the browser is running.	
Icon bar	<p>The icon bar includes a graphic representation for major functions accessible in the Core Console. It appears on the left side of the UI, directly below the branding area. Clicking the appropriate item in the icon bar takes you to the corresponding section of the UI where you can manage that function.</p> <p>Each of the icons in the icon bar is further described in the Icon bar table below.</p>	
Left navigation area	<p>The left navigation area appears on the left side of the user interface, below the icon bar.</p> <ul style="list-style-type: none"> The left navigation area contains the text filter and the Machines menu. 	


UI Element	Description
	<ul style="list-style-type: none"> If you have added replication to this Core, then this area contains a Replicated Machines menu. If you have any machines that were removed from protection but for which recovery points were saved, then this area contains a Recovery Points Only menu. If you added any custom groups, then this area contains a Custom Group menu. If you attached an archive, then this area contains an attached archives menu. <p>You can toggle the appearance of the left navigation area on and off. This is helpful when you need to see more content in the main navigation area of the UI. To hide this section, click the gray border between the left navigation and main navigation areas. To show this UI element once more, click the gray border again.</p> <p>Each of the elements in the left navigation area are further described in the Left navigation menu table below.</p>
Context-sensitive help	 <p>From the , each time you click the Help icon (a blue question mark), a resizable browser window opens with two frames. The left frame contains a navigation tree showing topics from the <i>Rapid Recovery User Guide</i>. The right frame displays content for the selected help topic. At any given time, the help navigation tree expands to show the location in its hierarchy for the selected topic. You can browse through all help topics using this context-sensitive help feature. Close the browser when you are done browsing topics.</p> <p>You can also open help from the Help option of the Help menu.</p>




Button bar

Details about the button bar appear in the following table.

i **NOTE:** Icons appear in the button bar only when the Theme general setting for the Core is set to **Hybrid** (the default display theme).

Table 168: Button bar buttons and menus








UI Element	Description
Button bar: Protect button and menu	 <p>The Protect button launches the Protect Machine Wizard, from which you can protect a single machine in the Rapid Recovery Core. Additionally, for other protection options, you can access the drop-down menu next to this button, which includes the following options.</p> <ul style="list-style-type: none"> The Protect Machine option is another method to launch the Protect Machine Wizard to protect a single machine.




UI Element	Description
	<ul style="list-style-type: none"> The Protect Cluster option allows you to connect to a server cluster. The Protect Multiple Machines option opens the Protect Multiple Machines Wizard to allow you to protect two or more machines simultaneously. The Deploy Agent Software option lets you install the Rapid Recovery Agent software to one or more machines simultaneously. This function uses the Deploy Agent Software Wizard.
Button bar: Restore button and menu	 <p>The Restore button opens the Restore Machine Wizard to allow you to restore data from recovery points saved from a protected machine. Additionally for other restore or export options, you can access the drop-down menu next to this button, which includes the following options.</p> <ul style="list-style-type: none"> The Restore Machine option is another method to launch the Restore Machine Wizard to restore data. The Mount Recovery Point option launches the Mount Wizard, which lets you mount recovery points from a protected machine. The VM Export option opens the Export Wizard. From this wizard you can create a virtual machine from recovery points saved in the Rapid Recovery Core. You have the option of creating a one-time export, or you can define parameters for a VM that is continually updated after every snapshot for a protected machine.
Button bar: Archive button and menu	 <p>The Archive button opens the Archive Wizard. From this wizard you can create a one-time archive from selected recovery points, or you can create an archive and continually save to that archive based on a schedule you define. Additionally for other archive options, you can access the drop-down menu next to this button, which includes the following options.</p> <ul style="list-style-type: none"> The Create Archive option is another method to launch the Create Archive Wizard to create a one-time archive or to archive continually. The Import Archive option launches the Import Archive Wizard, which lets you import an archive. The Attach Archive option mounts an archive so you can read the contents as a file system.
Button bar: Replicate button	 <p>The Replicate button opens the Replication Wizard. From this wizard you can specify a target Core, select machines protected on your source Core, and replicate recovery points from selected machines to the target Core in the repository you specify. You can pause replication when defining it, or you can have replication begin immediately. Additionally, you can specify whether a seed drive will be used to copy data for existing recovery points to the target Core.</p>













Icon bar

Details about the icon bar appear in the following table.

Table 169: Icon bar

UI Element	Description	
Icon bar	The icon bar includes a graphic representation for major functions accessible in the Core Console. Clicking the appropriate item takes you to the corresponding section of the user interface where you can manage that function. Icons in the icon bar include:	
Icon bar: Home icon		Home. Click the Home icon to navigate to the Core Home page.
Icon bar: Replication icon		Replication. Click the Replication icon to view or manage incoming or outgoing replication.
Icon bar: Virtual Standby icon		Virtual Standby. Click the Virtual Standby icon to export information from a recovery point to a bootable virtual machine.
Icon bar: Events icon		Events. Click the Events icon to view a log of all system events related to the Rapid Recovery Core.
Icon bar: Settings icon		Settings. Click the Settings icon to view or manage settings for the Rapid Recovery Core. You can back up or restore Core configuration settings. You can set general settings to control ports or display aspects. Additionally, you can configure settings in the following categories: automatic updates; nightly jobs; transfer queue settings; client timeout settings; DVM deduplication cache settings; Replay engine settings; and deploy settings. You can view or change database connections; SMTP server settings; cloud storage accounts; and change font settings for reports. You can set SQL attachability settings; Core job settings; license settings; SNMP settings; and vSphere settings.
Icon bar: More icon		More. Click the More icon to access other important features. Each has its own icon, listed below.
Icon bar: More icon	System Information	 System Information. Click System Information to display data about the Rapid Recovery Core server. You can see the host name, OS, architecture and memory for the Core. You can see the name displayed on the Core Console. You can also view the fully qualified domain name of the Core on your network, and the path for your cache metadata and deduplication caches. For more information about changing the display name, see Understanding system information for the Core . For more information about deduplication cache, see Understanding


UI Element	Description	
		<p>deduplication cache and storage locations.</p> <p>For information on adjusting the settings, see Configuring DVM deduplication cache settings.</p>
Icon bar: More icon	Virtual Environment	 <p>The Virtual Environment option appears in the ... (More) menu only if your Core is installed on a Hyper-V or vCenter/ESXi virtual machine. Clicking this option results in the <i>Virtual Environments</i> page, which provides options for managing credentials, storage locations, repositories, and volumes for Hyper-V and ESXi virtual environments. This page has three sub-pages (<i>Virtual Storage</i>, <i>Attached Disks</i>, and <i>Provisioning</i>), from which you can accomplish the following tasks, respectively:</p> <p>Manage hypervisor credentials. On the <i>Virtual Storage</i> sub-page, you can manage the credentials for Hyper-V or vCenter/ESXi hypervisor hosts added to or protected on your Core. If you add a hypervisor host and enter your credentials, Rapid Recovery caches them for future use.</p> <p>Manage hypervisor host storage locations. On the <i>Virtual Storage</i> sub-page, once one or more protected hypervisor hosts appear on this page, you can manage storage locations. You can see and set a path to a physical folder on your hypervisor (for Hyper-V), or a path for your data store (for ESXi). You can expand or contract the view of storage locations defined to see all disks; expand again to show snapshots on the base disks, and which is currently active and mounted.</p> <p>Monitor virtual disks. On this page, you can see disks attached to your virtual environments. From the <i>Attached Disks</i> sub-page, you can view and monitor all virtual disks currently attached to the virtual machine, including the subset of disks not specifically defined as storage locations in the <i>Virtual Storage</i> sub-page.</p> <p>i NOTE: For this reason, the number of disks listed in the <i>Attached Disks</i> sub-page may exceed the number of volumes shown on the <i>Virtual Storage</i> sub-page.</p> <p>Define repositories or virtual volumes. On the <i>Provisioning</i> sub-page, you can create a new repository for your Hyper-V or ESXi protected machines. You can also add an empty volume for your virtual environments.</p> <p>i NOTE: Creation of either a repository or a virtual volume requires a storage location to be defined on the <i>Virtual Storage</i> sub-page as a prerequisite.</p>
Icon bar: More icon	Archives	 <p>Archives. Rapid Recovery lets you manage archives of information from the Core. You can view information about scheduled or attached archives, and you can add, check, or import archives.</p>
Icon bar:	Mounts	 <p>Mounts. Lets you view and dismount local mounts, and view and</p>

UI Element	Description	
More icon	disconnect remote mounts.	
Icon bar: More icon	Boot CDs 	Boot CDs. Lets you manage boot CDs, typically used for a bare metal restore (BMR). You can create a boot CD ISO image, delete an existing image, or click the path for the image to open or save it.
Icon bar: More icon	Repositories 	Repositories. Lets you view and manage repositories associated with your Core.
Icon bar: More icon	Encryption Keys 	Encryption Keys. Lets you view, manage, import, or add encryption keys that you can apply to protected machines. If not being used, you can delete encryption keys.
Icon bar: More icon	Cloud Accounts 	Cloud Accounts. Lets you view and manage connections between your Core and Cloud storage accounts.
Icon bar: More icon	File Search 	File Search. Lets you search through recovery points for specific files that you can then restore to a local disk.
Icon bar: More icon	Retention Policy 	Retention Policy. Lets you view and modify the Core retention policy, including how long to keep recovery points before rolling them up and eventually deleting them.
Icon bar: More icon	Credentials Vault 	Credentials Vault. Lets you access and manage user accounts used within the Rapid Recovery Core Console. Easily update credentials or save time entering them when working with your Core.
Icon bar: More icon	Notifications 	Notifications. Lets you configure notifications about Core events, define SMTP server settings to email notifications, and set repetition reduction to suppress repeated notifications about the same event.
Icon bar: More icon	Mail Restore 	Mail Restore. Lets you search through Exchange server databases to locate and restore mail messages.
Icon bar: More icon	Downloads 	Downloads. You can download the Agent software web installer, the Local Mount Utility, or MIB files containing event information to use in an SNMP browser.
Icon bar: More icon	Reports 	Reports. Lets you access Core reports or schedule reports to generate on an ongoing basis.
Icon bar: More icon	Core Log 	Core Log. Lets you download Core log file for diagnostic purposes.

Left navigation menu


The full set of menus that may appear in the left navigation area are described in the following table:












Table 170: Left navigation menu options



UI Element	Description
Protected Machines menu	<p>The PROTECTED MACHINES menu appears as the first menu in the left navigation area, if one or more machines is protected in your Core.</p> <p>If you click a specific machine name shown in this pane, a Summary page appears, showing summary information for the selected machine. For more information on what you can accomplish on the Summary page, see Viewing summary information for a protected machine.</p>
Replicated machines menu	<p>If you see the name of another Rapid Recovery Core as a top-level navigation menu, then the Core on which you are viewing the Core Console is a target Core. The menu is named after the source Core, and each machine listed under it represents a machine from that source Core that is replicated on this target.</p> <p>If this target Core replicates recovery points from more than one source Core, each source Core appears as its own navigable menu in the left navigation area.</p> <p>If you click a specific machine name shown in a replicated machines menu, a Summary page appears, showing summary information for the selected replicated machine.</p> <p>For more information about replication, see Replication.</p>
Recovery Points Only menu	<p>If you see a RECOVERY POINTS ONLY menu, your Core retains recovery points for a machine it once protected or replicated. While that machine is no longer continuing to capture new snapshots, the recovery points previously captured on your Core remain. These recovery points can be used for file-level recovery, but cannot be used for bare metal restore, for restoring entire volumes, or for adding snapshot data.</p>
Custom groups menu	<p>If you created any custom groups, a custom group menu appears in the navigation menu. Custom groups are logical containers used to group machines together (for example, by function, or organization, or by geographic location). Custom groups can contain heterogeneous objects (protected machines, replicated machines, and so on). You can define the label for a custom group; like other menus, the name appears in the menu in all upper-case letters.</p> <p>You can perform actions for like items in a custom group by clicking the arrow to the right of the custom group title. For example, you can force a snapshot for every protected machine in a custom group.</p> <p>For more information about creating and managing custom groups, see Understanding custom groups.</p>
Attached archives menu	 <p>If you attach any archives to your Core, each archive is listed in the left navigation menu. Its label is the name of the archive. Contained in this list is each machine included in the archive.</p>

Details about the elements in the left navigation area appear in the following table.

Table 171: Left navigation area and menus

UI Element	Description
Machines menus text filter	 <p>The text filter is a text field that lets you filter the items displayed in the Machines, Replicated Machines, and Recovery-Point Only machines menus. If you type your criteria in this filter, then only the machines that meet your criteria display in the appropriate menus.</p>

UI Element	Description
Expand and contract details	 Click the arrow to the right of the text filter to expand and contract detail for the Machines, Replicated Machines, and Recovery-Point Only machines menus.
Protected Machines menu	<p>The PROTECTED MACHINES menu appears in the left navigation area of the UI. In this menu, you can view any protected machines, protected clusters, or replicated machines configured in your Core. If you have any protected groups or recovery point-only machines, these also appear as part of this menu.</p> <p>If you click the Protected Machines menu label, a Protected Machines page appears, showing on one page all of the machines protected on this Core. For more information, see Viewing the Protected Machines menu.</p> <p>The default view for this menu is expanded. You can contract or expand the view for any of the protected machines in your Core by clicking the  [Contract menu] and  [Expand menu] arrows on the left side of this menu label, respectively. The following list shows the various icons that can appear in the Protected Machines menu by machine type:</p> <div>  A simple machine icon portrays a physical machine or a protected VM with Rapid Recovery Agent software installed. </div> <div>  A hollow triple-machine icon portrays an VMware vCenter/ESXi host. <div>  NOTE: If a vSphere/ESXi hypervisor has machines under protection, the appears in this menu, with its VMs as children, even if the host is not protected. </div> </div> <div>  A hollow double-machine icon portrays a protected VM guest machine on a VMware vCenter host. </div> <div>  A multi-machine icon portrays a protected cluster. </div> <div>  A single-machine icon with a small horizontal line above portrays a single node in a cluster. </div> <div>  A double-machine icon filled with two short horizontal lines and a dot portrays a Scale-Out File Server (SOFS) cluster protected in your Core. </div> <div>  A single-machine icon filled with two short horizontal lines and a dot portrays a node in a protected SOFS cluster. </div>
Replicated Machines menu	<p>If replicating machines from another Rapid Recovery Core, the name of that Core appears as a separate menu under the Machines menu. Each machine replicated on this target Core from the listed source Core appears under this menu.</p> <p>For each replicated machine, the icon indicates the type of machine being replicated. For example, if replicating a single machine, the icon shows one machine. If replicating a server cluster, the icon represents a cluster.</p> <p>You can collapse or expand the view for any of these replicated machines by clicking the arrow on the left side of this menu label.</p> <p>From the Replicated Machines menu, you can perform actions on all replicated machines. If you click the Replicated Machines menu, the Machines page appears. This page shows all machines protected on another (source) Core that are replicated to this target Core. For</p>

UI Element	Description
	more information, see Viewing replicated machines from the navigation menu .
Recovery Points Only menu	 <p>If any machines previously protected on the Core were removed from protection, but the recovery points were not deleted, then the Recovery Points Only menu appears. There is no menu icon. Each of the formerly protected machines with retained recovery points displays in this list. The recovery points-only machine show a standard protected machine icon, with no status icon.</p> <p>You can collapse or expand the view for any of the recovery points-only machines by clicking the arrow on the left side of this menu label.</p> <p>From the Recovery Points Only menu, you can remove the recovery points for all the recovery-points only machines on this Core.</p> <p>If you click the Recovery Points Only menu, the Machines page appears, showing the machines from which the recovery points were saved. For more information, see Viewing the Recovery Points Only menu.</p>
Custom Groups menu	<p>If your Core includes any custom groups, then the left navigation area includes a Custom Group menu. Each of the objects in that custom group displays in this list.</p> <p>You can collapse or expand the view for any of the custom groups in your Core by clicking the arrow on the left side of this menu label.</p> <p>From the Custom Groups menu, you can perform actions for like items in the group.</p> <p>If you click the Custom Groups menu, the Machines page appears, showing a pane for each of the Rapid Recovery objects that appear in your group: protected machines, replicated machines, and recovery points-only machines. For more information, see Viewing the Custom Groups menu.</p>
Attached archives menu	 <p>If you attached any archives to your Core, then the left navigation area includes a menu for each attached archive. Each of the protected machines included in the archive displays in this list. The menu label uses the name specified when the archive was saved.</p> <p>You can collapse or expand the view for any of the archives attached to your Core by clicking the arrow on the left side of this menu label.</p> <p>From the attached archives menu, you can perform actions for like items in the group.</p> <p>If you click the attached archives menu, the Machines page appears, showing a pane for each of the Rapid Recovery objects that appear in your group: protected machines, replicated machines, and recovery points-only machines. For more information, see Viewing the Custom Groups menu.</p>

Viewing protected machines

From the Home page on the Rapid Recovery Core Console, when viewing the Summary Tables view, you can see summary information for any machines protected by the Core in the Protected Machines pane.

i **NOTE:** A software agent acts on behalf of the user to take specific actions. Protected machines are sometimes referred to as agents, since they run the Rapid Recovery Agent software to facilitate data backup and replication on the Rapid Recovery Core.

You can view the status, the display name for each machine, which repository it uses, the date and time of the last snapshot, how many recovery points exist in the repository for the machine, and the total amount of storage space the snapshots use in the repository.

To manage aspects of any protected machine, start by navigating to the machine you want to view, configure, or manage. From the Home page, there are three ways to navigate to a protected machine:

- You can click on the IP address or display name of any protected machine from the Protected Machines pane. This takes you to the *Summary* page for the selected protected machine.
- In the left navigation area, you can click on the title of the Protected Machines menu. The Protected Machines page appears. On this page, you can see summary information about each machine. For a detailed description of this page, see [Viewing summary information for a protected machine](#).
- In the left navigation area, under the Protected Machines menu, you can click any protected machine IP address or display name. This takes you to the Summary page for the selected protected machine. For a detailed description of this page, see [Viewing summary information for a protected machine](#).

Viewing events for a protected machine


On the *Events* page, you can view the jobs that occurred or are in progress for the protected machine you selected. Buttons at the top of the page let you navigate to lists of jobs in each of the three categories of activities:

- **Tasks.** A job that the Rapid Recovery Core must perform to operate successfully.
- **Alerts.** A notification related to a task or event that includes errors and warning.
- **Journal.** A composite of all protected machine tasks and alerts.

The following table includes descriptions of each element on the *Events* page.


Table 172: *Events* page elements

UI Element	Description
Search keyword	Lets you search for a specific item within each category. Available for tasks only.
From	To narrow your results, you can enter a date at which to begin searching. Available for tasks only.
To	To narrow your results, you can enter a date at which to stop searching. Available for tasks only.
Status icons	<p>Each icon represents a different job status. For alerts and tasks, clicking one of the icons lets you filter the list by that status, essentially generating a report. Clicking the icon a second time removes the filter for that status. You can filter by more than one status. Statuses include:</p> <ul style="list-style-type: none">• Active. A job that is in progress.• Queued. A job that is waiting for another job to complete before it can initiate.• Waiting. A job waiting for your approval or completion, such as a seed drive. (For more information about seed drives, see Replication.)• Complete. A job that completed successfully.• Failed. A job that failed and did not complete.
Service icon	This button adds services jobs to the list of jobs. When you click this icon, a smaller service icon appears on each status icon, which lets you filter by service jobs that have those statuses (if any exist). Examples of services jobs include deleting index files or removing a machine from protection.

UI Element	Description
Export type drop-down list	<p>The drop-down list includes the formats to which you can export the event report. Available for tasks only. It includes the following formats:</p> <ul style="list-style-type: none"> • PDF • HTML • CSV • XLS • XLSX
 (Export icon)	Converts the event report to the format you selected. Available for tasks only.
Page selection	Event reports can include several jobs across multiple pages. The numbers and arrows at the bottom of the <i>Events</i> page let you navigate the additional pages of the report.

The *Events* page displays all events in a table. The following table lists the information shown for each item.






Table 173: Detailed information for the Event summary table

UI Element	Description
Status	Shows the status for the task, alert, or journal item. Available for alerts or journal items, click the header to filter the results by status.
Name	Name is available for tasks only. This text field lists the task type that completed for this protected machine. Examples include transfer of volumes, maintaining repository, rolling up, performing mountability checks, performing checksum checks, and so on.
Start Time	Available for tasks, alerts, and journal items. Shows the date and time when the job or task began.
End Time	Available for tasks only. Shows the date and time when the task completed.
 Job Details	Available for tasks only. Opens the <i>Monitor Active Task</i> dialog box, so you can view details of the specific job or task. These details include an ID for the job, rate at which the Core transferred data (if relevant), elapsed time for the job to complete, total work in amount of gigabytes, and any child tasks associated with the job.
Message	Available for alerts and journal items. This text field provides a descriptive message of the alert or journal item.

Viewing the More menu for a protected machine

The *More* menu offers additional options to help manage the selected a protected machine. To access these tools, click the More drop-down menu and select from one of the options described in the following table.

Table 174: Tools accessible from the *More* option for a protected machine

Icon	UI Element	Description
	System Information	<p>Shows information about the protected machine, system information, volumes, processors, network adapters, and IP addresses for this machine.</p> <p>For more information, see Viewing system information for a protected machine.</p>
	Mounts	<p>From the Local Mounts pane, you can view or dismount volumes mounted locally. From the Remote Mounts pane, you can view or dismount volumes mounted using the Local Mount Utility.</p> <p>For information on dismounting volumes, see Dismounting recovery points.</p> <p>For information on mounting a recovery point locally, see Mounting a recovery point or Mounting a recovery point volume on a Linux machine, respectively.</p>
	Retention Policy	<p>Lets you specify a retention policy for the selected machine. You can choose to use the Core's default policy, or you can differentiate the retention policy for this machine only. For more information, see Customizing retention policy settings for a protected machine.</p>
	Notifications	<p>Lets you specify a custom notification group for events pertaining to the selected machine. This does not change the notifications already set on the Core. For more information, see Configuring notification groups.</p>
	Agent Log	<p>Lets you download and view the log file for a machine protected using the Rapid Recovery Agent software. For more information, see Downloading and viewing the log file for a protected machine.</p>

REST APIs

The purpose of this section is to provide an introduction and overview of the Rapid Recovery Representational State Transfer (REST) Application Program Interfaces (APIs), their use, and their function.

The Rapid Recovery Web Service APIs are RESTful and let you automate and customize certain functions and tasks within the Rapid Recovery software solution to assist you with meeting your business objectives.

These APIs are accessible from the *Downloads* page of the Rapid Recovery License Portal.

Topics include:

[Intended audience](#)

[Working with Rapid Recovery REST APIs](#)

[Downloading and viewing Core and Agent APIs](#)

[Recommended additional reading](#)

Intended audience

Rapid Recovery APIs are intended for use by application developers who want to integrate and extend Rapid Recovery in their application, as well as administrators who want to script interactions with the Rapid Recovery Core server.

Working with Rapid Recovery REST APIs

The Rapid Recovery APIs are REST-style APIs, which means that they use HTTP requests to provide access to resources (data entities) through URI paths. Rapid Recovery APIs use standard HTTP methods such as GET, PUT, POST, and DELETE. Because REST APIs are based on open standards, you can use any language or tool that supports HTTP calls.

There are two ways that application developers and administrators can work with Rapid Recovery APIs. They are:


- Using C# or other .NET languages to directly use Rapid Recovery .NET client DLL files.
- Communicate directly with the HTTP endpoint to generate your own XML.

The first approach is recommended. The client DLLs are included in the Rapid Recovery SDK. The method for calling Rapid Recovery APIs is consistent with the way you would consume any .NET 4.5X Windows Communication Foundation (WCF) service.

Downloading and viewing Core and Agent APIs

The Rapid Recovery *Software Developer Kit (SDK)* includes REST APIs for the Rapid Recovery Core and Rapid Recovery Agent components, and samples and supporting files. These contents are contained in the following folders and then compressed as an archive that includes the following components:

Table 175: Components included in the SDK archive

Folder name	Contents	Description
Core.Contracts	Rapid Recovery Core APIs	<p>Contains APIs to assist developers or administrators to script functions in Rapid Recovery Core. There are 2 sets of service contracts.</p> <ol style="list-style-type: none">1. Open the CoreWeb.Client HTML file in a web browser to view information for general REST standards. The service contracts are listed. When you click any corresponding hyperlinked uniform resource identifier (URI), the browser opens information in the Core.Contracts/docWeb/ directory. The resulting page shows information for general REST service operations, including methods and descriptions.2. Open the Core.Client HTML file in a web browser to view detailed C# information. When you click any hyperlinked service contract (class), the browser opens information in the Core.Contracts/doc/ directory. The resulting page shows detailed information for all C# methods in the selected class.
Agent.Contracts	Rapid Recovery Agent APIs (deprecated)	<p>Contains APIs that developers or administrators can use to manipulate Rapid Recovery Agent on protected machines.</p> <div> CAUTION: Agent APIs are deprecated and will be removed from a future version of the SDK. Direct manipulation of the Agent APIs is not recommended. Use of these APIs is considered customization and will not be supported. The information is provided in documentation for historical purposes.</div> <ol style="list-style-type: none">1. Open the AgentWeb.Client HTML file in a web browser to view information for general REST standards.2. Open the Agent.Client HTML file in a web browser to view detailed C# information.
AppRecoveryAPISamples	Code samples and dynamic	AppRecoveryAPISamples contains code samples that are written in C# programming language. These files represent

Folder name	Contents	Description
	link libraries	<p>a good starting point to view code snippets if using the APIs to customize your GUI, management systems, and so on.</p> <p>AppRecoveryAPISamples\Dependencies contains dynamic link library (DLL) files that Rapid Recovery uses. The DLLs contain data contracts (types the Core is familiar with) and service contracts (management methods and operations that can be used to force Core do something). If you want to customize your own graphic user interface, or use a management system to work with Rapid Recovery, these DLLs are required.</p> <p>NOTE: The DLL version used must match the version of the Core.</p>

You can download the SDK as archive (API-Reference-x.x.x-xxxx). Each x represents a digit in the build number for the relevant release.

Complete the steps in this procedure to obtain the SDK, download it to your specified destination, and decompress the files in preparation for using Core and Agent APIs.

1. Log in to the Rapid Recovery License Portal at <https://licenseportal.com>.
2. In the left navigation menu of the license portal, click **Downloads**.
The Downloads page of the license portal appears.
3. On the Downloads page, in the Windows-Based Applications section, scroll down to the description for the SDK and click **Download**.
4. Save the downloaded archive to your preferred location.
5. Decompress the archive.
In the new API-Reference-x.x.x-xxxx folder, you see the separate sets of files described in the preceding table.
6. Open the key HTML files described in the preceding table in a web browser to see guidance about the APIs.

Recommended additional reading

The *Rapid Recovery Installation and Upgrade Guide* provides an overview of the Rapid Recovery architecture; describes the steps necessary for installing Rapid Recovery components and for upgrading the Core or Agent components from earlier versions. It also describes licensing steps required from the Rapid Recovery Core.

You can view or download this guide from <https://support.quest.com/rapid-recovery/technical-documents/>.

Glossary

Agent

The Rapid Recovery Agent is a software agent that can be installed on a physical or virtual machine to protect it in the Rapid Recovery Core. When enabled in the Core at a specified interval, the Agent software tracks changed blocks on volumes of protected disks, and captures snapshots of the changed blocks based on the protection interval.

base image

The first backup transfer saved to the Core is called a base image snapshot. All data on all specified volumes (including the operating system, applications, and settings), are saved to the Core. For more information, see [snapshot](#).

checksum

A checksum is a function that creates blocks of data that are used for the purpose of detecting accidental errors that are created during transmission or storage.

cluster

See [Windows failover cluster](#).

cluster node

An individual machine that is part of a [Windows failover cluster](#).

compression

The Storage Networking Industry Association (SNIA) defines compression as the process of encoding data to reduce its size.

Core

The Rapid Recovery Core is the central component of the Rapid Recovery architecture. The Core provides the essential services for backup, recovery, retention, replication, archiving, and management. In the context of outgoing replication, the Core is also called a source Core. When a Rapid Recovery Core is the recipient or destination of recovery points from another Core, it is called a target Core.

Core Console

The Rapid Recovery Core Console is a Web-based interface that lets you fully manage the Rapid Recovery Core.

database availability group (DAG)

A set of up to 16 Microsoft Exchange Server 2010 Mailbox servers that provide automatic, database-level recovery from a database, server, or network failure. DAGs use continuous replication and a subset of Windows failover clustering technologies to provide high availability and site resilience. Mailbox servers in a DAG monitor each other for failures. When a Mailbox server is added to a DAG, it works with the other servers in the DAG to provide automatic, database-level recovery from database failures.

encryption

Data is encrypted with the intent that it is only accessible to authorized users who have the appropriate decryption key. Data is encrypted using 256-bit AES in Cipher Block Chaining (CBC) mode. In CBC, each block of data is compared with the previous ciphertext block before being encrypted, this way each new ciphertext block depends on all preceding plaintext blocks. A passphrase is used as an initialization vector.

event

An event is a process that is logged by the Core. Events can be viewed within the Core Console by clicking the (Events) icon from the icon bar. The default view when you click this icon shows the Tasks page. This view shows events related to a job. Priority events about which you are notified can be viewed on the Alerts page. A log of all events appears in the Journal page. By setting up or modifying existing notification groups, you can customize notification for any event. This action raises the priority of the event by displaying it on the Alerts page. Members of a notification group are notified of events using the notification methods set in the notification options for the group.

global deduplication

The Storage Networking Industry Association (SNIA) defines data deduplication as the replacement of multiple copies of data—at variable levels of granularity—with references to a shared copy to save storage space or bandwidth. The Rapid Recovery Volume Manager performs global data deduplication within a logical volume. The granularity level of deduplication is 8 KB. The scope of deduplication in Rapid Recovery is limited to protected machines using the same repository and encryption key.

incremental snapshot

Incremental snapshots are backups consisting only of data changed on the protected machine since the last backup. They are saved to the Core regularly, based on the interval defined (for example, every 60 minutes). For more information, see [snapshot](#).

license key

A license key is one method used to register your Rapid Recovery software or appliance. (You can also use a license file.) You can obtain license keys or files when you register on the Rapid Recovery License Portal for an account. For more information, see [License Portal](#).

License Portal

The Rapid Recovery License Portal is a Web interface where users and partners can download software, register Rapid Recovery appliances, and manage license subscriptions. License Portal users can register accounts, download Rapid Recovery Core and Agent software, manage groups, track group activity, register machines, register appliances, invite users, and generate reports. For more information, see the *Rapid Recovery License Portal User Guide*.

Live Recovery

Rapid Recovery Live Recovery is an instant recovery technology for VMs and servers. It provides near-continuous access to data volumes in a virtual or physical server, letting you recover an entire volume with near-zero RTO and a RPO of minutes.

Local Mount Utility

The Local Mount Utility (LMU) is a downloadable application that lets you mount a recovery point on a remote Rapid Recovery Core from any machine.

log truncation

Log truncation is a function that removes log records from the transaction log. For a SQL Server machine, when you force truncation of the SQL Server logs, this process simply identifies free space on the SQL server. For an Oracle database server, when you truncate logs (on a schedule, or manually on demand), archive logs are deleted, freeing up space. Likewise, when you force truncation of the Exchange Server logs, this action frees up space on the Exchange server.

management roles

From the *Roles* page, the QorePortal enables the assignment of management roles from which you can divide administrative responsibility among trusted data and service administrators, as well as access control to support secure and efficient delegation of administration.

mountability

Exchange mountability is a corruption detection feature that alerts administrators of potential failures and ensures that all data on the Exchange servers is recovered successfully in the event of a failure.

Object File System

The Rapid Recovery Scalable Object Store is an object file system component. It treats all data blocks, from which snapshots are derived, as objects. It stores, retrieves, maintains, and replicates these objects. It is designed to deliver scalable input and output (I/O) performance in tandem with global data deduplication, encryption, and retention management. The Object File System interfaces directly with industry standard storage technologies.

passphrase

A passphrase is a key used in the encryption of data. If the passphrase is lost, data cannot be recovered.

prohibited characters

Prohibited characters are characters that should not be used when naming an object in the Rapid Recovery Core Console. For example, when defining a display name for a protected machine, do not use any of the following special characters:

Table 176: Prohibited characters

Character	Character name	Prohibited from
?	question mark	machine display name, encryption key, repository, path description
	pipe	machine display name, encryption key, repository, path description
:	colon	machine display name, encryption key, repository Use of this symbol is supported when specifying a path; for example, c:\data.
/	forward slash	machine display name, encryption key, repository, path description
\	back slash	machine display name, encryption key, repository Use of this symbol is supported when specifying a local or network path; for example, C:\data or \\ComputerName\SharedFolder\

Character	Character name	Prohibited from
*	asterisk	machine display name, encryption key, repository, path description
"	quotation mark	machine display name, encryption key, repository, path description
<	open angle bracket	machine display name, encryption key, repository, path description
>	close angle bracket	machine display name, encryption key, repository, path description

prohibited phrases

Prohibited phrases are phrases (or sets of characters) that should not be used as the name for any object in the Rapid Recovery Core Console, because they are reserved for the use of operating systems. It is best practice is to avoid using these phrases at all if possible. For example, when defining a display name for a protected machine, do not use any of the following phrases:

Table 177: Prohibited phrases

Character	General use	Prohibited from
con	console	machine display name, encryption key, repository, path description
prn	printer port	machine display name, encryption key
aux	auxiliary port	machine display name, encryption key
nul	null value	machine display name, encryption key
com1, com2, ... through com9	communication port	machine display name, encryption key
lpt1, lpt2, ... through lpt9	asterisk	machine display name, encryption key, repository, path description

protected machine

A protected machine—sometimes called an "agent"—is a physical computer or virtual machine that is protected in the Rapid Recovery Core. Backup data is transmitted from the protected machine to the repository specified in the Core using a predefined protection interval. The base image transmits all data to a recovery point (including the operating system, applications, and settings). Each subsequent incremental snapshot commits only the changed blocks on the specified disk volumes of the protected machine. Software-based protected machines have the Rapid Recovery Agent software installed. Some virtual machines can also be protected agentlessly, with some limitations.

Rapid Recovery

Rapid Recovery sets a new standard for unified data protection by combining backup, replication, and recovery in a single solution that is engineered to be the fastest and most reliable backup for protecting virtual machines (VM), as well as physical and cloud environments.

quorum

For a failover cluster, the number of elements that must be online for a given cluster to continue running. The elements relevant in this context are cluster nodes. This term can also refer to the quorum-capable resource selected to maintain the configuration data necessary to recover the cluster. This data contains details of all of the changes that have been applied to the cluster database. The quorum resource is generally accessible to other cluster resources so that any cluster node has access to the most recent database changes. By default there is only one quorum resource per server cluster. A particular quorum configuration (settings for a failover cluster) determines the point at which too many failures stop the cluster from running.

recovery points

Recovery points are a collection of snapshots of various disk volumes. For example, C:\, D:\, and E\.

recovery points-only machine

A recovery points-only machine is the representation on the Core of recovery points from a machine that was previously protected on the Core, and since removed. If you remove replication but retain the recovery points, this also results in a recovery points-only machine. Information can be viewed and recovered at a file level. You cannot use a recovery points-only machine to perform BMR or to restore full volumes, nor can you add more data to a recovery points-only machine.

remote Core

A remote Core represents an Rapid Recovery Core that is accessed by a non-Core machine using the Local Mount Utility or the Central Management Console.

replication

Replication is the process of copying recovery points from one Rapid Recovery Core and transmitting them to another Rapid Recovery Core for disaster recovery purposes. The process requires a paired source-target relationship between two or more Cores. Replication is managed on a per-protected-machine basis. Any machine (or all machines) protected or replicated on a source Core can be configured to replicate to a target Core. It is the recovery points that are copied to the target Core.

repository

A repository is a collection of base image and incremental snapshots captured from the machines protected on a Rapid Recovery Core. Repositories must be created on fast primary storage devices. The storage location for a DVM repository can be local to the Core machine (in which case it is hosted on a supported Windows OS only). It can use direct-attached storage, a storage area network, or an appropriately rated network-attached server.

REST APIs

Representational State Transfer (REST) is a simple stateless software architecture designed for scalability. Rapid Recovery uses this architecture for its Application Program Interfaces (APIs) to automate and customize certain functions and tasks. There is a separate set of REST APIs for Core functionality and for protected machine (agent) functionality.

restore

The process of restoring one or more storage volumes on a machine from recovery points saved on the Rapid Recovery Core is known as performing a restore. This was formerly known as rollback.

retention

Retention defines the length of time the backup snapshots of protected machines are stored on the Rapid Recovery Core. Retention policy is enforced on the recovery points through the rollup process.

rollup

The rollup process is an internal nightly maintenance procedure that enforces the retention policy by collapsing and eliminating dated recovery points. Rapid Recovery reduces rollup to metadata operations only.

seeding

In replication, the initial transfer of deduplicated base images and incremental snapshots of protected agents, which can add up to hundreds or thousands of gigabytes of data. Initial replication can be seeded to the target Core using external media, which is useful for large sets of data or sites with slow links.

server cluster

See [Windows failover cluster](#).

SharePoint backup

A SharePoint backup is a copy of data that is used to restore and recover that data on a SharePoint server after a system failure. From the SharePoint backup, you can perform recovery of the complete SharePoint farm, or one or more components of the farm.

Smart Agent

See [agent](#).

snapshot

A snapshot is a common industry term that defines the ability to capture and store the state of a disk volume at a given point, while applications are running. The snapshot is critical if system recovery is needed due to an outage or system failure. Rapid Recovery snapshots are application aware, which means that all open transactions and rolling transaction logs are completed and caches are flushed prior to creating the snapshot. Rapid Recovery uses Microsoft Volume Shadow Services (VSS) to facilitate application crash consistent snapshots.

SQL attachability

SQL attachability is a test run within the Rapid Recovery Core to ensure that all SQL recovery points are without error and are available for backup in the event of a failure.

SQL backup

A SQL backup is a copy of data that is used to restore and recover that data on a SQL server after a system failure. From the SQL backup, you can perform recovery of the complete SQL database, or one or more of the components of the SQL database.

SQL differential backup

A differential database backup is a cumulative copy of all changes in data since the last full backup of the SQL database. Differential backups are typically faster to create than full database backups, and reduce the number of transaction logs required to recover the database.

target Core

The target Core, which is sometimes referred to as replica Core, is the Rapid Recovery Core receiving the replicated data (recovery points) from the source Core.

Transport Layer Security

Transport Layer Security (TLS) is a modern cryptographic network protocol designed to ensure communication security over the Internet. This protocol, defined by the Internet Engineering Task Force, is the successor to Secure Sockets Layer (SSL). The SSL term is still generally used, and the protocols are interoperable (a TLS client can downgrade to communicate to an SSL server).

True Scale

True Scale is the scalable architecture of Rapid Recovery.

Universal Recovery

Rapid Recovery Universal Recovery technology provides unlimited machine restoration flexibility. It enables you to perform monolithic recovery to- and from- any physical or virtual platform of your choice as well as incremental recovery updates to virtual machines from any physical or virtual source. It also lets you perform application-level, item-level, and object-level recovery of individual files, folders, email, calendar items, databases, and applications.

Verified Recovery

Verified Recovery technology is used to perform automated recovery testing and verification of backups. It supports various file systems and servers.

virtual standby

Virtual standby is a process that creates a clone virtual machine of a protected machine. The original source machine can be physical or virtual, but the product is always virtual. You can create a virtual standby one time on demand, or you can define requirements to create the bootable VM, and continually update it after each snapshot is captured on the original protected machine.

Volume Manager

The Rapid Recovery Volume Manager manages objects and then stores and presents them as a logical volume. It leverages dynamic pipeline architecture to deliver TruScale scalability, parallelism, and asynchronous inputand-output (I/O) model for high throughput with minimal I/O latency.

white labeling

Rapid Recovery provides the ability for certain specific providers of backup and disaster recovery services to white label (or rebrand) Rapid Recovery with their own logo, and then to market it as their own product or service. The option to white label Rapid Recovery is only available to organizations whose Quest Software license agreements explicitly grant rebranding privileges. Options to license, sub-license and/or re-distribute this software is also only granted to organizations with whom a specific license agreement granting these privileges is in effect. Quest Software reserves all rights to Rapid Recovery. The white labeling option does not grant anyone any intellectual property rights to Rapid Recovery. To request this functionality for your organization, please contact your Quest Software representative.

Windows failover cluster

A group of independent computers that work together to increase the availability of applications and services. The clustered servers (called nodes) are connected by physical cables and by software. If one of the cluster nodes fails, another node begins to provide service (a process known as failover). Users experience a minimum of disruptions in service. Rapid Recovery supports the protection of a number of SQL Server and Exchange Server cluster types.

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product