

**CARBONITE**<sup>®</sup>  
an **opentext**™ company

# Carbonite Server Backup Director 8.7

User Guide



© 2023 Open Text. All rights reserved.

This product may be protected by one or more US patents. See <https://www.opentext.com/patents> for details.

For terms and conditions, see <https://www.carbonite.com/terms-of-use/carbonite-general-enterprise-terms-of-service>.

Carbonite makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Carbonite reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of Carbonite to notify any person of such revision of changes. All companies, names and data used in examples herein are fictitious unless otherwise noted.

No part of this document may be reproduced, transmitted, transcribed, stored in a retrieval System or translated into any language including computer language, in any form or by any means electronic, mechanic, magnetic, optical, chemical or otherwise without prior written permission of:

Carbonite LLC  
251 Little Falls Drive  
Wilmington, DE 19808  
[www.carbonite.com](http://www.carbonite.com)

Carbonite and the Carbonite logo are trademarks of Carbonite, LLC. Product names that include the Carbonite mark are trademarks of Carbonite, LLC. All other products or company names mentioned in this document are trademarks or registered trademarks of their respective owners.

## Document History

Version	Date	Description
1	April 2023	Initial guide provided for Director 8.7x.

## Contents

<b>1</b>	<b>Introduction to Carbonite Server Backup Director .....</b>	<b>9</b>
1.1	Support for hourly backups .....	9
1.2	Director processes .....	12
1.3	Data replication between vaults .....	17
1.4	Companion products .....	24
<b>2</b>	<b>View and manage vaults using the Director UI.....</b>	<b>25</b>
2.1	Start and view the Director UI.....	25
2.2	View worker node information .....	27
2.3	Add a vault connection.....	28
2.4	Set workspace options .....	30
2.5	Encrypt a workspace .....	31
2.6	Add a workspace .....	32
2.7	Open a workspace .....	32
2.8	Rename a workspace .....	32
2.9	Delete a workspace .....	33
<b>3</b>	<b>Manage licensing.....</b>	<b>34</b>
3.1	View licensing information.....	35
3.2	Add a license key .....	35
3.3	Activate a license online.....	36
3.4	Activate a license manually .....	36
3.5	Request a license activation code .....	36
3.6	Enter a license activation code.....	37
3.7	Remove a license key .....	37
3.8	Use customer license quotas .....	38
3.9	Set license quotas for a customer .....	38
3.10	View the license quota for a computer .....	39
3.11	Reclaim a license .....	39
<b>4</b>	<b>Manage primary storage.....</b>	<b>40</b>
4.1	Add storage locations to the primary storage group .....	40

4.2	Add and edit policies for primary storage locations .....	41
4.3	Assign a location policy to a storage location .....	42
4.4	Change credentials for UNC storage locations.....	42
4.5	Change the maximum pool file size for the primary storage group .....	43
4.6	Set the minimum pool usage parameter for the primary storage group .....	44
4.7	Enforce retention settings in primary storage .....	44
4.8	Optimize pool files.....	46
4.9	Deduplicate data .....	48
4.10	Verify the integrity of backup files .....	48
4.11	Copy and clone data .....	49
4.12	Retire primary storage locations.....	54
4.13	Remove storage locations from the primary storage group.....	55
4.14	Delete a location policy.....	55
<b>5</b>	<b>Manage secondary storage .....</b>	<b>56</b>
5.1	Add secondary storage groups.....	56
5.2	Add storage locations to a secondary storage group.....	56
5.3	Assign secondary storage groups to tasks.....	57
5.4	Migrate safesets to secondary storage .....	58
5.5	Change the maximum pool file size for a secondary storage group.....	60
5.6	Apply safeset retention settings to the secondary storage pool .....	60
5.7	Secondary storage safeset deletion .....	61
5.8	Apply retention settings to safesets in primary and secondary storage.....	61
5.9	Secondary storage maintenance.....	62
5.10	Close secondary storage pools .....	62
5.11	Detach secondary storage pools .....	62
5.12	Attach secondary storage pools.....	63
5.13	Delete storage locations from a secondary storage group .....	63
5.14	Delete secondary storage groups.....	64
<b>6</b>	<b>Manage archive storage .....</b>	<b>65</b>
6.1	Add archive locations .....	65
6.2	Assign archive storage locations to tasks.....	65
6.3	Archive safesets from a task.....	66

6.4	Archive safesets to disk .....	66
6.5	Recall safesets .....	67
6.6	Remove an archive location .....	67
<b>7</b>	<b>Manage customers, locations, accounts and users.....</b>	<b>68</b>
7.1	Add a customer, location, account and user.....	68
7.2	Filter the customer list .....	69
7.3	Add a location, account and user .....	69
7.4	Add a location (billing) code.....	70
7.5	Edit a location (billing) code .....	70
7.6	Delete an unassociated location (billing) code .....	71
7.7	Add an account and user.....	71
7.8	Change account properties .....	72
7.9	Add a user.....	72
7.10	Change user properties .....	73
7.11	Disable a user .....	73
7.12	Enable a user .....	74
7.13	Delete a customer .....	75
7.14	Delete a location .....	75
7.15	Delete an account.....	75
7.16	Delete a user .....	76
<b>8</b>	<b>Manage computers, tasks and safesets .....</b>	<b>77</b>
8.1	View registered computers and tasks .....	77
8.2	Change a registered computer's Agent type.....	77
8.3	View safesets for a task.....	78
8.4	View and change safeset properties .....	80
8.5	Change the retention settings of multiple safesets .....	81
8.6	Enabled, disabled and suspect tasks .....	81
8.7	Enable a disabled or suspect task.....	82
8.8	Disable a task.....	83
8.9	Export safesets .....	84
8.10	Import safesets.....	86
8.11	Delete a registered computer .....	87

8.12	Delete a task.....	88
8.13	Delete safesets .....	88
<b>9</b>	<b>Manage replication between Satellite and Base vaults .....</b>	<b>90</b>
9.1	Schedule replication from Satellite vaults.....	90
9.2	Replicate data from Satellite vaults .....	91
9.3	Disable and enable N:1 replication services on a Base vault .....	91
9.4	Disable over-the-wire encryption for replication from earlier Satellite vault versions	92
9.5	Implement bandwidth throttling for replication from Satellite vaults .....	92
9.6	Set the replication policy for a Satellite vault .....	93
9.7	Set the retention policy for a Satellite vault .....	93
9.8	Set the operating mode for a Satellite vault .....	93
9.9	Set the operating mode for a customer, location or computer on a Satellite vault.....	94
9.10	Set the operating mode for a task on a Satellite vault.....	95
9.11	Set the heartbeat interval for a Satellite vault.....	96
9.12	Assign settings control to a Satellite vault .....	96
9.13	View replication activity between Satellite and Base vaults.....	97
9.14	Stop Base vault replication processes and services .....	97
9.15	Run Satellite to Base vault replication reports.....	98
9.16	Replace a failed Satellite vault .....	100
9.17	Set up dual network connections.....	100
<b>10</b>	<b>Manage replication between Active and Passive vaults .....</b>	<b>102</b>
10.1	Schedule replication between Active and Passive vaults.....	102
10.2	Replicate data from an Active vault to a Passive vault .....	103
10.3	Allow or pause replication from Active to Passive vaults for customers, locations or computers .....	103
10.4	Allow or pause replication from Active to Passive vaults for a task .....	104
10.5	Disable and enable replication services on an Active or Passive vault .....	104
10.6	Disable over-the-wire encryption for replication from earlier Active vault versions	105
10.7	Implement bandwidth throttling for replication from an Active vault .....	105
10.8	View replication activity between Active and Passive vaults.....	105
10.9	Run Active to Passive vault replication reports.....	106
10.10	Fail over from an Active vault to a Passive vault.....	108

10.11	Fail back to a formerly Active vault .....	109
<b>11</b>	<b>Configure and run reports.....</b>	<b>110</b>
11.1	Run a Storage, Storage Location or Last Backup Status report.....	110
11.2	Create a Storage report.....	111
11.3	Create a Storage Location report .....	113
11.4	Create a Vault Storage report .....	114
11.5	Create a Last Backup Status report .....	115
11.6	Create a Late Server Status report .....	116
11.7	Create a Missed Backups report.....	117
11.8	Create a Storage Pool Summary report .....	119
11.9	Select the destination for a report .....	121
<b>12</b>	<b>Automate maintenance processes .....</b>	<b>124</b>
12.1	Scheduled maintenance operations.....	124
12.2	Enable or disable automated maintenance .....	125
12.3	Verify that automated maintenance is running.....	125
12.4	Stop automated maintenance.....	125
12.5	Run a maintenance operation on demand .....	126
12.6	Change the time of a scheduled maintenance operation.....	126
12.7	Enable automated secondary storage maintenance .....	126
12.8	Verify maintenance performance .....	127
12.9	View maintenance logs .....	127
12.10	Create a custom scheduled operation .....	128
12.11	Disable a scheduled operation .....	128
12.12	Enable a scheduled operation.....	128
12.13	Remove a scheduled operation.....	129
<b>13</b>	<b>Monitor and manage vaults.....</b>	<b>130</b>
13.1	View and stop vault processes in the Job Monitor .....	130
13.2	View, start and stop vault services.....	131
13.3	View log files.....	132
13.4	Select log file types to display .....	134
13.5	View replication logs .....	134
13.6	Log message codes.....	135

13.7	Purge activity records and logs .....	136
13.8	Schedule a log file purge .....	137
13.9	Configure email notifications .....	137
13.10	Set primary storage thresholds .....	138
13.11	Modify advanced vault settings .....	139
13.12	View, add and remove a vault's alternate vaults .....	140
13.13	Back up a vault database.....	140
<b>14</b>	<b>Command Reference .....</b>	<b>142</b>
14.1	agentinfoports .....	143
14.2	dbbackup .....	143
14.3	migratesecondary.....	144
14.4	replvault .....	147
14.5	secondaryop .....	150
14.6	vaultop.....	155
14.7	Vault Settings.....	168
14.8	vanalyz .....	171
14.9	varchive .....	177
14.10	vvexport.....	178
14.11	vvimport .....	179
14.12	vvmigrat.....	179
14.13	vvmove .....	181
14.14	vvpoolop.....	183
14.15	vvpurge.....	186
14.16	vvrecall .....	188
<b>15</b>	<b>Carbonite Server Backup Support .....</b>	<b>189</b>
15.1	Contacting Carbonite.....	189



# 1 Introduction to Carbonite Server Backup Director

Carbonite Server Backup Director is an online data vault that securely receives and stores backup data from servers where agents are installed. The Director application manages the data, monitors vault activities, provides data for restores, and deletes data in response to requests from Carbonite Server Backup Portal. For more information, see [Director processes](#).

To ensure that data is always available, data can be backed up to one vault and replicated to another vault. Data can then be backed up to and restored from the second vault if the first vault is unavailable. See [Data replication between vaults](#).

To manage Director vaults, you can use the Director UI: the graphical user interface (GUI) for managing vaults, scheduling automated tasks, and monitoring activities. For more information, see [View and manage vaults using the Director UI](#). To perform tasks using the command line interface (CLI), see [Command Reference](#).

You can automate vault management and monitoring using Vault API to create, delete, update, and enumerate customers, locations, computers, tasks, accounts, and users. For more information, see the *Vault API Guide*.

## 1.1 Support for hourly backups

To help customers meet their recovery point objectives (RPOs), backups to a vault can be scheduled to run multiple times per day, as often as hourly. Support for hourly backups was added in Director 8.60.

To ensure that replication and maintenance operations run on tasks with multiple backups per day, Director includes:

- [Backup load management](#)
- [Replication load management and deferral](#)
- [Triggered maintenance and replication](#)

*Note:* Backup performance is affected by many factors. To meet your RPOs, you might need to increase vault resources or change backup jobs and schedules. For information about changing backup jobs and schedules, see the [Server Backup online help](#) or Agent guides.

### 1.1.1 Backup load management

When an agent that supports hourly backups contacts a Director 8.60 or later vault to start a backup, Director indicates whether the vault is busy with high-priority maintenance for the job data.

If the vault is busy and the backup is scheduled by an intra-daily schedule, the agent skips the backup. Intra-daily schedules, where backups are scheduled to run multiple times per day, are created using Portal 8.88 or later.

If the vault is busy and the backup is scheduled to run daily, weekly or monthly, or is an ad hoc backup (i.e., not scheduled), the backup is delayed for five minutes. After this delay, the backup starts and interrupts any maintenance that is running for the task.

The following table summarizes how new backups are handled when a backup or maintenance is already running for the task.

<b>Current vault activity for the task</b>	<b>Schedule of new backup</b>	<b>Result</b>
Backup	Intra-daily	New backup is skipped.
	Daily or less often, or ad hoc (not scheduled)	New backup starts when the current backup is finished.
High-priority maintenance	Intra-daily	New backup is skipped.
	Daily or less often, or ad hoc (not scheduled)	New backup is delayed for five minutes. After this delay, the backup starts and interrupts any maintenance that is running for the task.
Other maintenance	Any	New backup interrupts maintenance.

In previous Director versions, a new backup was always queued if a backup was already running for the task, and always interrupted maintenance on a task.

### 1.1.2 Replication load management and deferral

Beginning in Director 8.60, replication is deferred or delayed if the target vault is busy with high-priority maintenance for the task, if maintenance on the source vault has not run in the past 24 hours, or if replication for a task fails.

If the target vault (e.g., the Base vault in N:1 replication or Passive vault in 1:1 replication) is busy with high-priority maintenance for the task and data for the task was replicated successfully within the past 12 hours, incoming replication for a task is deferred until a later time. If this occurs, the incoming replication is deferred for 30 minutes and then retries.

If replication for a task has been deferred for more than 12 hours, the incoming replication is performed immediately without checking whether the vault is busy. If multiple replication sessions have been deferred for the task, when replication resumes after the 12-hour deferral period or when maintenance is finished, only the last replication session resumes. All existing safesets are replicated in this replication session.

The following table summarizes how, beginning in Director 8.60, a target vault handles incoming replication for a task when it is busy with maintenance for the task.

Target vault condition		Result
Vault busy with high-priority maintenance for the task	Task replicated successfully within past 12 hours	Incoming replication is deferred for 30 minutes and then retries.
	Replication for the task has been deferred for more than 12 hours	Replication is performed immediately without checking whether vault is busy. The maintenance process is interrupted.

Replication for a task is also delayed or deferred if:

- On a source vault where maintenance is enabled, replication is triggered by a backup or incoming replication (see [Triggered maintenance and replication](#)) but maintenance has not run successfully on the task in the past 24 hours and a successful replication has run in the past 24 hours.
- Replication for a task fails (e.g., due to network failure). If this occurs, the replication is deferred for 5 minutes and then retries multiple times.

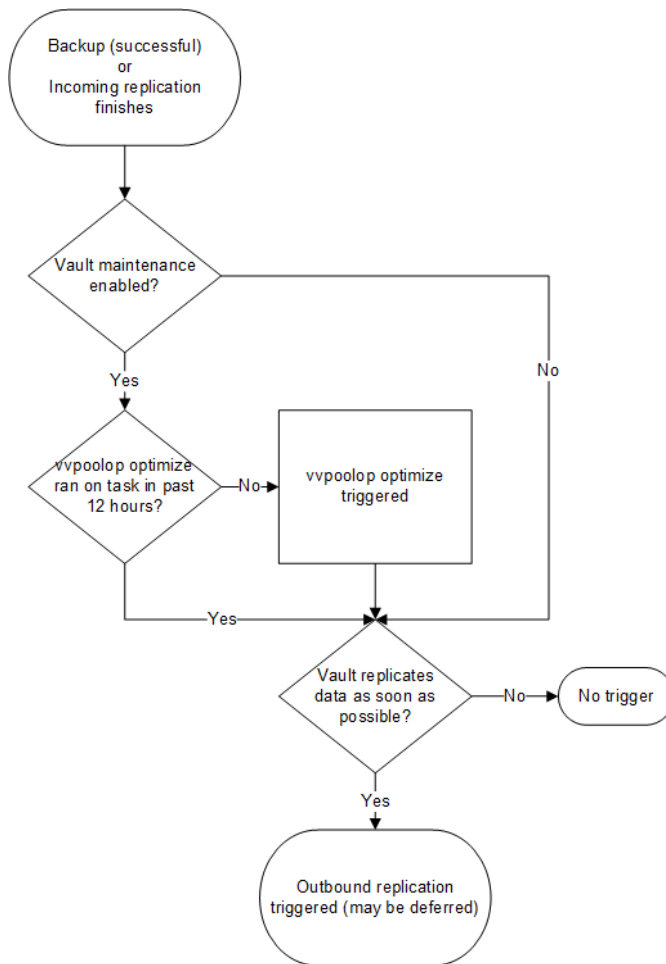
### 1.1.3 Triggered maintenance and replication

Beginning in Director 8.60, maintenance and replication operations are sometimes triggered to run for a task when:

- A successful backup finishes for the task.
- Incoming replication (successful or not) finishes for the task.

As shown in the following diagram, after incoming replication or a successful backup finishes for a task, 'vvpoolop optimize' is triggered for the task if maintenance is enabled on the vault and 'vvpoolop optimize' did not run on the task in the past 12 hours. This maintenance operation removes pool system data that is no longer referenced by safesets and defragments the pool. Migration also runs as part of this process, to identify safesets that should be kept based on retention settings and mark other safesets for deletion. See [Migration and safeset retention management](#).

Outbound replication is triggered for the task if the vault is a source vault (i.e., Satellite, Active or Active Base vault) that is set to replicate data as soon as possible after changes are made. However, the triggered replication might not run immediately. If maintenance is enabled on the vault, 'vvpoolop optimize' has not run successfully on the task in the past 24 hours and replication has run successfully on the task in the past 24 hours, or if the target vault is busy with maintenance, the outbound replication will be deferred.



On a source vault that is set to replicate data as soon as possible after changes are made, outbound replication is triggered whenever 'vvpoolop optimize' finishes for a task, regardless of whether it was triggered by a backup or incoming replication or not.

## 1.2 Director processes

Director receives data from servers where Agents are installed, and stores and manages the data using the following processes:

- [Backup](#)
- [Migration and safeset retention management](#)
- [Optimization](#)
- [Diagnostics](#)
- [Replication](#)
- [Restore](#)
- [Data deletion](#)

### 1.2.1 Backup

Agents send backup data to the Director vault. A Director Listener service listens for new agent connections. When an agent connects to the vault, the Listener starts a server process. The server process authenticates the agent credentials, and receives, synchronizes and commits the backup.

When an agent sends backup data to a vault, the data is saved in primary storage and pooled with other safesets from the same task. Each backup is saved as a safeset. A safeset is an instance of backup data with defined retention policies and other properties. See [View safesets for a task](#).

The first backup from an agent is called a full seed, and includes all data. In subsequent backups, the Director indexes changed data blocks against previous safesets to create delta backups that require less disk space on the vault. With Delta technology, only data that has changed from the previous backup is sent to the vault. Deltizing is Agent technology that tracks the changes within a backup file and transmits only those changes to the vault. On the vault, a corresponding mechanism reconstructs the deltized files. Director creates a database that indexes changes across versions of data, allowing you to reconstruct the data upon request.

Director uses a specialized storage pooling system to manage large volumes of backup data. When data is backed up from a server, the data is deduplicated so that only one physical copy of each common segment of data is stored. The data is also compressed, ensuring that it requires a small amount of storage space on the vault.

Director security features protect your data from unauthorized access. Over-the-wire (OTW) encryption is used for all communications to and from the vault. To ensure the integrity of your backup data, Director verifies each backup to detect processing or transmission errors. In addition, scheduled analysis and verification tasks maintain the validity of backup data.

To enhance security, agents can encrypt backup data. When data is encrypted, the entire cryptographic process and key management is unknown to Director, and independent of Director processes. When running a new backup job, data must be encrypted using AES-256. When running a backup job that was originally sent to a vault version earlier than 7.11, data can be encrypted with its original encryption type.

Beginning in Director 8.60, if agent-vault certificate pinning is enabled in a vault, when an agent that supports this security feature tries to connect to the vault (e.g., to run a backup or restore), it checks whether the public key of the vault's TLS certificate is the same as when the agent previously connected to the vault. If the public key of the vault certificate is different, the agent reports a certificate failure and will not connect to the vault unless you re-pin the certificate using Portal. For more information, see the [Server Backup online help](#).

Backup and Listener log files are created during the backup phase. Backup log files are stored in the task's Logs folder. Listener log files are stored in the global Logs folder.

## 1.2.2 Migration and safeset retention management

Migration is a Director maintenance process that identifies which safesets in primary storage should be kept based on retention settings and retention groups. Other safesets are expired and marked for deletion. See [Enforce retention settings in primary storage](#).

When running or scheduling a backup, the user specifies retention settings for the resulting safeset. Retention settings for a safeset specify the number of days the safeset should be kept online on the vault, how many safesets in the task's retention group should be stored online, and whether/how long the safeset should be stored offline.

Based on its retention settings, each safeset is assigned to a retention group. A retention group is a group of safesets for a task that are treated as a group during migration. For example, in a single task, safesets that have the default "Daily" retention settings would be in one retention group, safesets with "Weekly" retention settings would be in a second retention group, and safesets with "Monthly" retention settings would be in a third retention group. For information about viewing a safeset's retention group, see [View and change safeset properties](#).

Beginning in Director 8.60, 'vvpoolop optimize' is sometimes triggered after backups and replications. Migration runs as part of this process. See [Triggered maintenance and replication](#). In previous Director versions, migration was scheduled to run on the entire vault once a day.

A migrate log file is created during migration. The log file is stored in the Global Logs folder when the migration runs on multiple tasks or in the Task Logs folder when the migration runs on a single task.

## 1.2.3 Optimization

Pool file optimization reclaims unused storage on the vault. The optimization process ensures that the vault uses the smallest amount of storage possible. It removes pool system data that is no longer referenced by any safesets, defragments the pool, and uses the Minimal Pool Usage setting to combine valid data into new pool files.

The optimization phase starts after the migration phase identifies expired data. This process ensures that the vault retains only valid safesets and that the vault only uses the space required by valid safesets. Running regular migrations and optimizations removes unnecessary safesets and conserves storage space.

Beginning in Director 8.60, pool file optimization ('vvpoolop optimize') is sometimes triggered after backups and replications. See [Triggered maintenance and replication](#).

By default, deduplication runs once per month. The deduplication process deletes all data blocks that are the same (except for one), and updates references to point to the same common block. Use vault storage reports to review what data blocks deduplication removed.

Data for a backup job cannot be deduplicated when the encryption type, compression type or encryption password changes. If parameters are not changed, deduplication can run on the next backup of the same job.

The optimizer log file is created during the optimization phase. The log file is stored in the Global Logs folder when the optimization runs on multiple systems or in the Task Logs folder when the optimization runs on a single task.

## 1.2.4 Diagnostics

To ensure that safesets are valid and that client data can be retrieved, you can use Director diagnostic tools to verify the integrity of backup data and troubleshoot issues. You can also use the diagnostic tools to check physical objects within a Pool System as well as the logical integrity of the Pool System data. See [Verify the integrity of backup files](#).

## 1.2.5 Replication

To ensure that data is always available for restore, even if a particular vault is offline or unavailable, backup data can be replicated from one vault to another. See [Data replication between vaults](#).

This Director version uses dual-stage replication. In this replication method, data deduplication is performed during replication. As a result, the pool size on the target vault can be smaller than the pool size on the source vault.

If a replication session is interrupted (e.g., due to power or network loss) before a safeset is completely replicated, the next replication session begins where the last session ended. Data that has already been replicated does not have to be sent again.

*Note:* You cannot restore data from a partial safeset on a target vault. A safeset must be completely replicated before it can be restored from the target vault.

### 1.2.5.1 Certificate verification and pinning for vault-to-vault communications

To improve the security of vault-to-vault communications, vault certificates are verified when a Director 8.7 source vault connects to a Director 8.7 target vault to replicate data, copy or clone data, or run a replication report. For the operation to proceed:

- The source vault certificate must be within its validity period (i.e., not expired).
- The target vault certificate must be within its validity period (i.e., not expired).
- If the target vault certificate is CA-signed:
  - The subject name or subject alternative name of the target vault certificate must match the FQDN of the target vault server. If the target vault has a wildcard certificate, the domain name of the certificate must match the domain name of the target server.
  - The CA-signing authority certificate must be in the Trusted Root Certification Authorities store on both the source vault server and the target vault server.

If the certificate verification fails, the source vault does not connect to the target vault and the replication, copy, clone or replication report fails.

If the certificate verification passes and the target vault certificate is self-signed, the source vault pins the target vault's certificate when it first connects to the target vault. When the source vault tries to connect to the target vault again, it checks whether the target vault certificate has changed. If the vault certificate has changed to another self-signed certificate, the source vault will not connect to the target vault unless you delete the pinned certificate from the source vault. See [vaultop delete pinned certificate](#). If the vault certificate has changed to a CA-signed certificate that passes verification, the source vault will connect to the target vault without manual intervention.

If you replace the CA-signed certificate on a target vault with a self-signed certificate, a source vault will not connect to the target vault again unless you delete the recorded certificate from the source vault. See [vaultop delete pinned certificate](#).

**IMPORTANT:** Do not delete a pinned or recorded vault certificate unless you are sure that there is no security risk. Please contact your IT security staff or service provider to determine whether the certificate change on the target vault was expected or whether further investigation is required.

Certificate verification and pinning for vault-to-vault communications is always performed when the source and target vaults are version 8.7. You do not have to manually enable this security feature.

*Note:* Vault certificates are not verified or pinned when a Director version 8.6x or earlier vault connects to a Director 8.7 vault.

## 1.2.6 Restore

When an agent submits a restore request, the system uses the pool index to rebuild a virtual copy of the data. The pool index contains pointers to pool files that contain the data.

Director handles two types of restores:

- When a user submits a request to restore some or all data from a safeset, the vault streams the data to the agent.
- When a user mounts a volume from a safeset, the agent makes small calls for parts of the data as needed when a user browses, copies and reads the data.

The restore process transmits backup data from the vault to the agent. You restore data from an online safeset in primary or secondary storage at any time after a backup completes.

## 1.2.7 Data deletion

When deleting a job or computer in Carbonite Server Backup Portal, a Portal Admin user can request that backup data for the job or computer be deleted from all vaults. The data deletion is scheduled for 72 hours after the request is made.

If a data deletion request is not canceled during the 72-hour waiting period, the deletion request is sent to Director vaults that are registered to Carbonite Server Backup API – Monitoring. In response to the request, Director deletes the data from any standalone, Base or Active vault where the data is stored. Replication processes then delete the data from any associated Satellite or Passive vault.



If a deletion request fails, an email notification is sent to a vault administrator whose email address is specified in Portal. The vault administrator can then manually delete the data.

To enable this data deletion functionality:

- The Reporting service must be installed with each standalone, Base and Active vault and registered to Carbonite Server Backup API – Monitoring. For data deletion, the Reporting service does not have to be installed with each Satellite and Passive vault; replication processes delete data from these vaults after a data deletion request. However, we recommend installing the Reporting service with each Passive vault and registering it to the API so that it is available if you have to fail over to the (formerly) Passive vault.

*Note:* Although it is not required for data deletion, the Reporting service must be installed with Satellite and Passive vaults and registered to API – Monitoring to provide data through API – Monitoring calls. For more information, see the *API – Monitoring Installation Guide*.

The Reporting service can be installed using the Director installation kit. For more information, see the *Director Installation Guide*.

- Portal must be registered to the same Carbonite Server Backup API – Monitoring instance as the Director Reporting service.
- The data deletion feature must be enabled in Portal.

For more information about data deletion, see the Portal documentation.

## 1.3 Data replication between vaults

To ensure that data is available for restore even if one vault is offline or unavailable, backup data can be replicated from one vault to another.


Three replication scenarios are available:

- One-to-one (1:1) replication. In this configuration, which is typically used for Offsite Replication Services (ORS), data is replicated from an Active vault to a Passive vault. See [One-to-one \(1:1\) replication](#).
- Many-to-one (N:1) replication. In this configuration, which is typically used for Managed Service Providers (MSPs), data is replicated from multiple Satellite vaults to one Base vault. See [Many-to-one \(N:1\) replication](#).
- Many-to-one-to-one (N:1:1) replication. In this configuration, which is typically used for Cloud-Connected Service Providers (CCSPs), data is replicated from multiple Satellite vaults to an Active Base vault and then to a Passive Base vault. See [Many-to-one-to-one \(N:1:1\) replication](#).


### 1.3.1 Many-to-one (N:1) replication

In Many-to-one (N:1) replication, which is typically used for Managed Service Providers (MSPs), Satellite vaults are installed at customer locations to allow for quick local backups.

As shown in the following diagram, the data is then replicated to a Base vault in the cloud, or in a secondary location in the customer's environment. A standard vault that receives data from Satellite vaults in N:1 replication is called a Base vault (also known as a BAV).



As shown in the following diagram, if a Satellite vault fails or becomes unavailable for any reason, agents send backups directly to the Base vault and restore data from the Base vault.



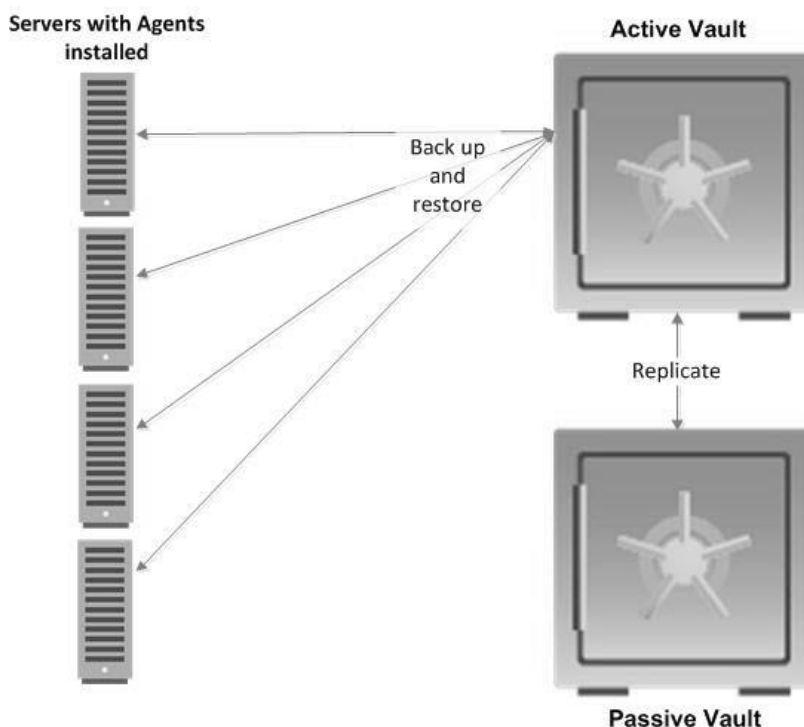
For N:1 replication, you must use customer license quotas to limit the amount of storage or number of agents and plug-ins that a specific customer can use. When you create customers, locations, accounts and users on the Base vault, they are replicated on the Satellite vaults.

You cannot perform administration tasks on Satellite vaults. Instead, you can manage Satellite vaults on the Base vault.

### 1.3.2 One-to-one (1:1) replication

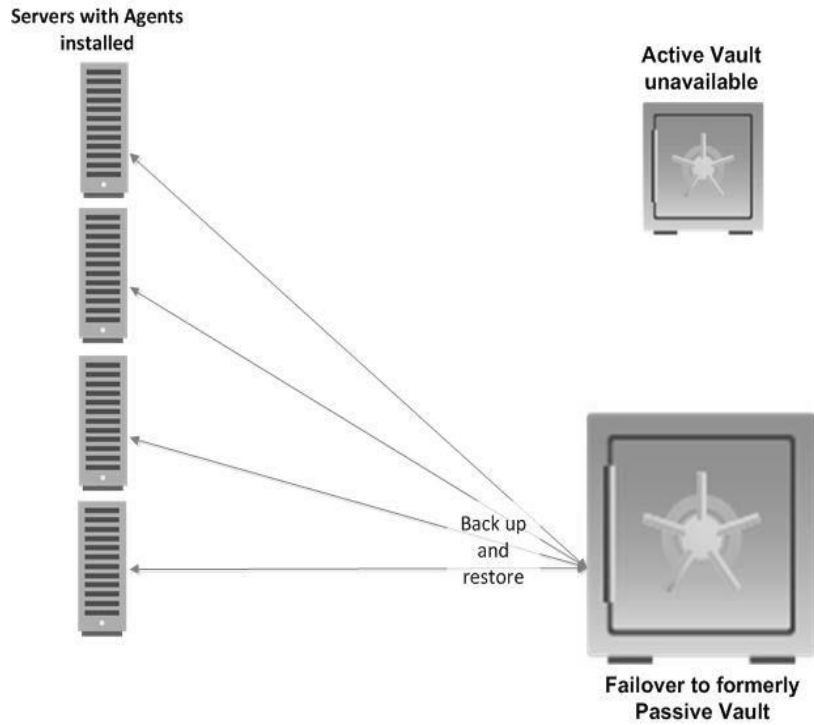
In 1:1 replication, safeset data is automatically and continually replicated from an Active vault to a Passive vault. You can then back up and restore data even if one vault is unavailable, by “failing over” to the other vault.

As shown in the following diagram, agents send backup data to the Active vault and can restore data from the Active vault. Backup data is replicated from the Active vault to the Passive vault.



If the Active vault fails or becomes unavailable for some reason, you can “fail over” to the Passive vault so that the formerly Passive vault becomes the new Active vault. Failover does not happen automatically; you must manually fail over to the Passive vault. See [Fail over from an Active vault to a Passive vault](#). As shown in the following diagram, agents then bypass the formerly Active vault, send backups directly to the formerly Passive vault, and restore data from the formerly Passive vault.

**Warning:** Do not register an agent directly to a Passive vault. If an agent exists on the Passive vault but not on the Active vault, the agent will be removed from the Passive vault when the vaults are synchronized. To restore data from a Passive vault, fail over to the Passive vault.

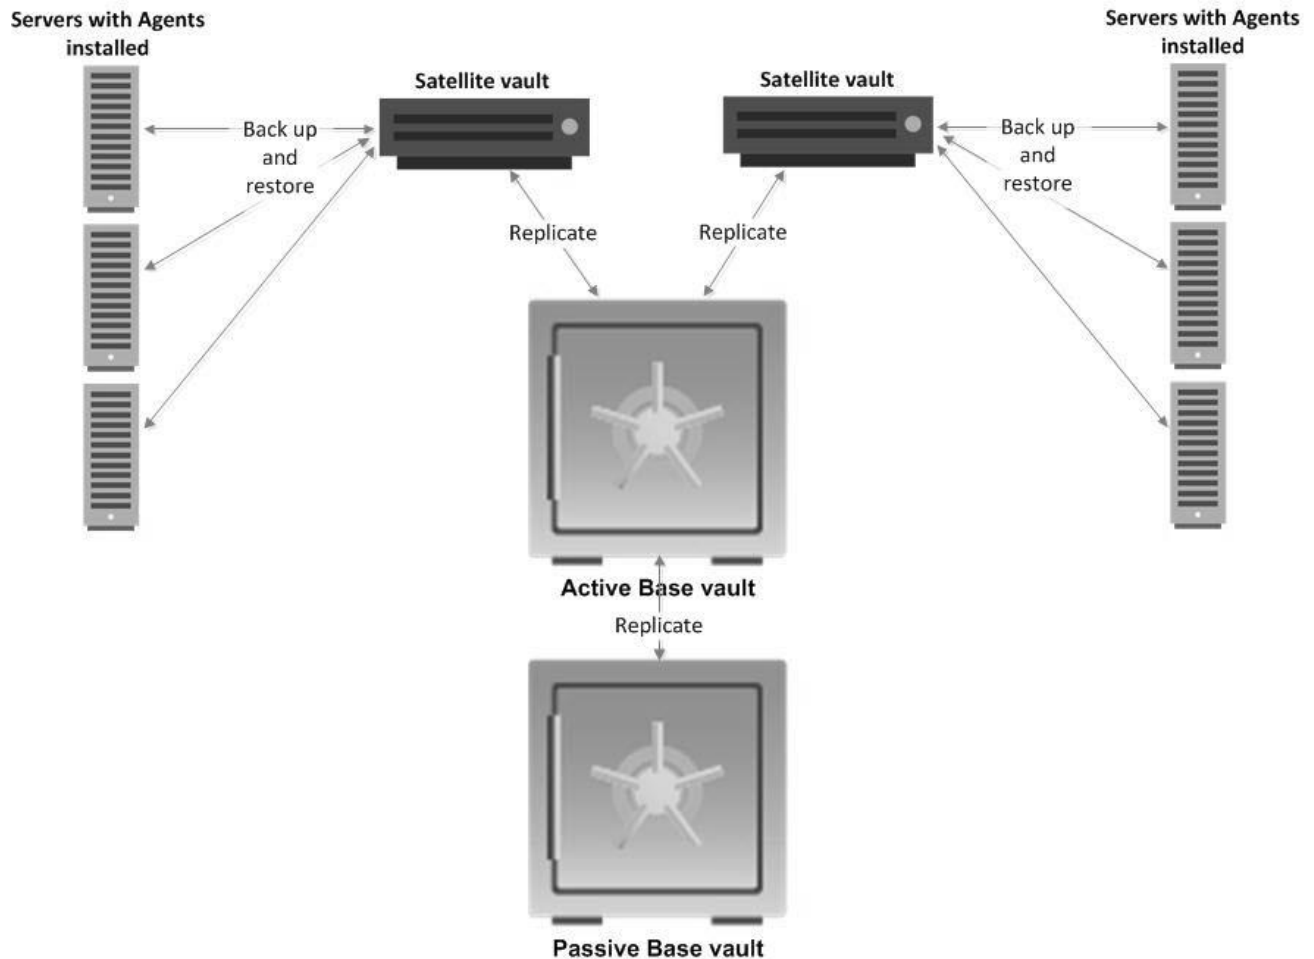


In 1:1 replication, safesets are replicated if they are online and not expired. Data stored in archives and secondary storage pools are not replicated. Safesets on the Active vault can be replicated to the Passive vault when a backup completes successfully, or you can schedule them to replicate on specific days at a specific time.

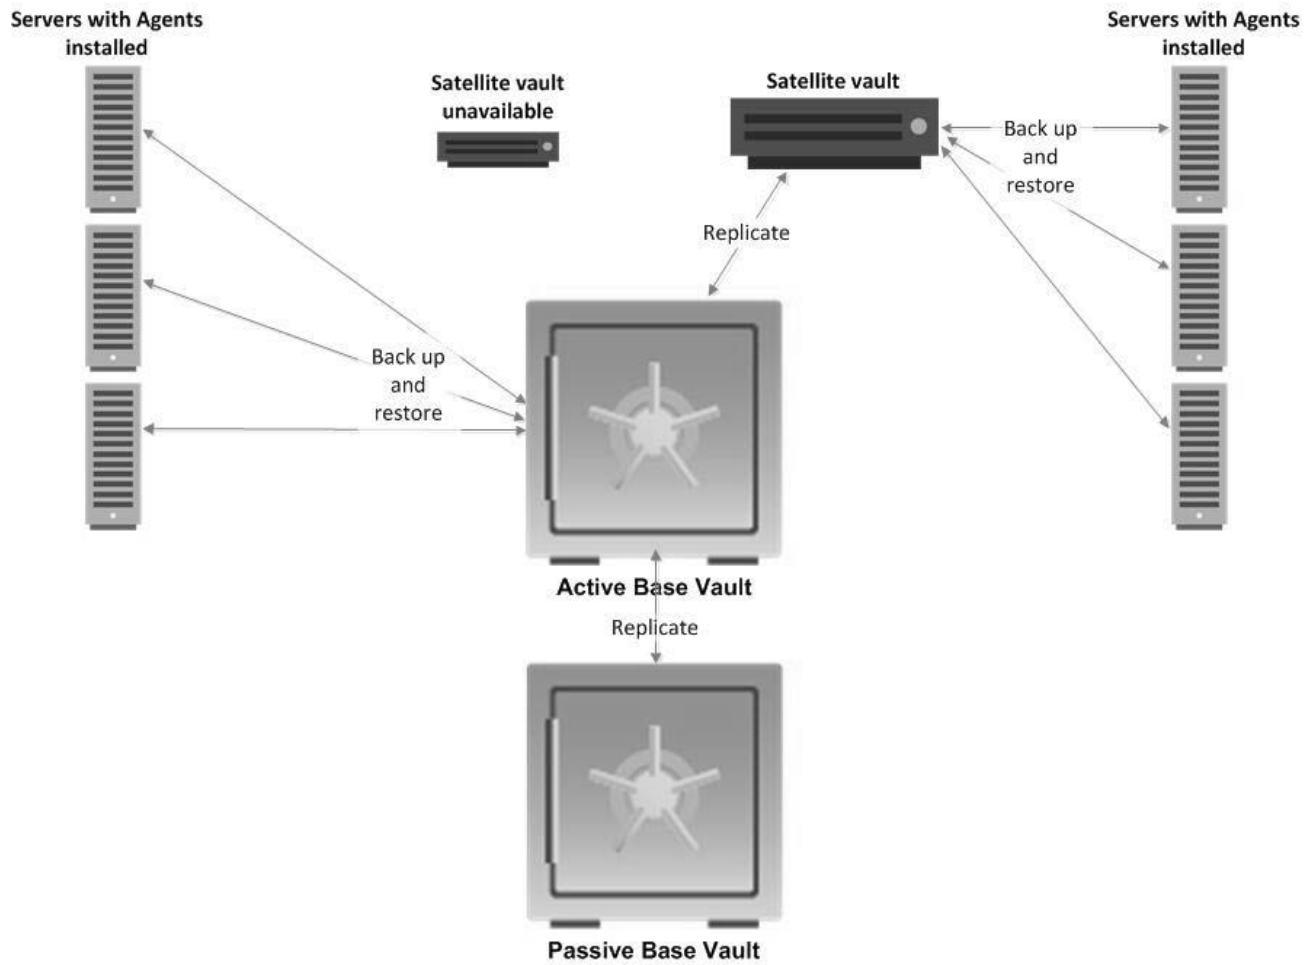
### 1.3.3 Many-to-one-to-one (N:1:1) replication

Many-to-one-to-one (N:1:1) replication, which is typically used for Cloud-Connected Service Providers (CCSPs), combines the benefits of N:1 and 1:1 replication and provides the highest level of data protection available.


Satellite vaults are installed at customer locations to allow for quick local backups. As shown in the following diagram, data is replicated to a vault called an “Active Base vault” and then to a second vault called a “Passive Base vault”. The Active Base vault could be in a secondary location in the customer’s environment, and the Passive Base vault could be in the cloud.



As shown in the following diagram, if a Satellite vault fails or becomes unavailable for some reason, agents can send backups directly to the Active Base vault, and data can be restored from the Active Base vault.



As shown in the following diagram, if the Active Base vault fails, you can fail over from the Active Base vault to the formerly Passive Base vault. Data from the Satellite vaults is then replicated directly from the Satellite vaults to the new Active Base vault.



When the previously Active Base vault becomes available again, you must run replication from the current Active Base vault to synchronize the data. After synchronization, you can leave the vault assignments as they are, or change the current Passive Base vault into the Active Base vault again. If you change a Passive Base vault into an Active Base vault without synchronizing and replicating the data, all unreplicated data is permanently deleted.

In N:1:1 replication, you must use customer license quotas to limit the amount of storage, number of agents, and number of plug-ins that a specific customer can use.

## 1.4 Companion products

This section describes optional products that are available for use with Director.

### Reports Extractor

Carbonite Server Backup Reports Extractor is an application that gathers vault information (e.g., activity, storage, task, etc.) and sends it to a web server where it can be displayed in various reports.

The Extractor is installed on the server where the vault is installed.

Reports Extractor is the first part of the Web reporting process that produces reports from a Web browser using data collected from the vault.

### Secondary Restore Server

Carbonite Server Backup Secondary Restore Server is used to restore backup data from a detached secondary storage device to an agent computer. The original vault is not required to complete the restore.



## 2 View and manage vaults using the Director UI

You can view and manage vaults using the Director UI. You can add connections for one or more vaults that you want to manage, view information for each vault, and manage the vaults. Communication between the Director UI and vault is encrypted, to allow remote vault management.

You can organize vault connections in one or more workspaces, and save the workspaces so that you do not have to re-enter vault information each time you manage the vaults.

You cannot use an earlier Director UI version to manage a later version vault. The UI must be the same Director version as the vault or a later Director version.

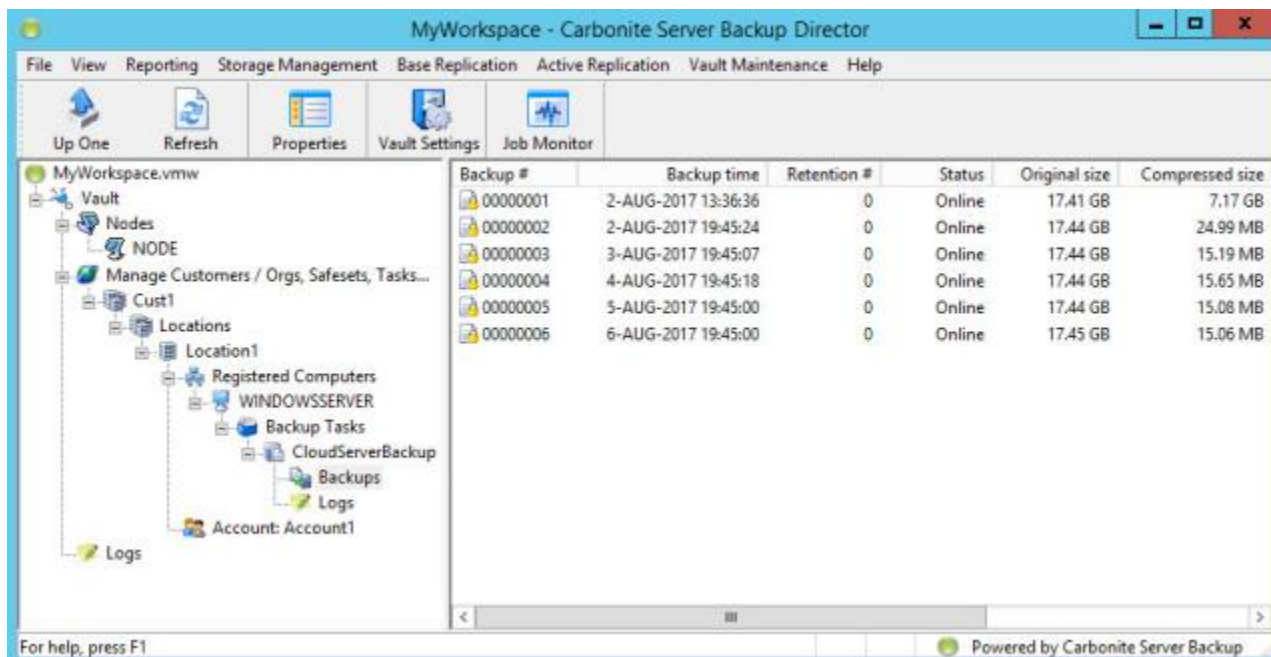
You do not need Administrator permissions to run the Director UI and save workspaces.

### 2.1 Start and view the Director UI

To start the Director UI, do one of the following:

- On the desktop, double-click the Carbonite Server Backup Director shortcut.
- In the Director UI installation directory, double-click the VMAdmin.exe file.

If vault connections have been added to the workspace, the left pane of the Director UI shows vaults, vault nodes, registered computers, and backup tasks in a hierarchy. The right pane shows detailed information about items you select in the left pane.























Backup #	Backup time	Retention #	Status	Original size	Compressed size
00000001	2-AUG-2017 13:36:36	0	Online	17.41 GB	7.17 GB
00000002	2-AUG-2017 19:45:24	0	Online	17.44 GB	24.99 MB
00000003	3-AUG-2017 19:45:07	0	Online	17.44 GB	15.19 MB
00000004	4-AUG-2017 19:45:18	0	Online	17.44 GB	15.65 MB
00000005	5-AUG-2017 19:45:00	0	Online	17.44 GB	15.08 MB
00000006	6-AUG-2017 19:45:00	0	Online	17.45 GB	15.06 MB

For help, press F1

Powered by Carbonite Server Backup

The following table describes items in the hierarchy, and shows the icon that represents each item.

Item	Icon	Description
Workspace		A workspace (saved as a .vmw file) with vault connections. You can create different workspaces to group different vaults. You can password protect a workspace.
Vault		Connection to a standalone vault.
		Connection to an Active vault in one-to-one (1:1) replication.
		Connection to a Passive vault in one-to-one (1:1) replication.
		Connection to a Base vault in many-to-one (N:1) replication.
		Connection to a Satellite vault in many-to-one (N:1) or many-to-one-to-one (N:1:1) replication.
		Connection to an Active Base vault in many-to-one-to-one (N:1:1) replication.
		Connection to a Passive Base vault in many-to-one-to-one (N:1:1) replication.
Nodes		Worker node container.
Node		A worker node in a vault. This icon represents the server where all vault components are installed.
Customer		Organization that owns the protected systems and data. There may be only one, or many different customers on a vault.
Location		Physical location (e.g., office building) of the customer.
Account/User	 	Needed to authenticate systems when they connect to the vault.
Registered computer		System protected by the vault.
Registered cluster of VMware ESX and ESXi hosts		Cluster of VMware ESX and ESXi hosts protected by the vault.
Task		The backup job defined to protect selected data.

Item	Icon	Description
Safeset in primary storage		Data that has been backed up. The data can be encrypted so that no one but the owner can restore the data, and no one can read the data when it stored on a vault, or is transmitted to or from a vault.
Safeset in secondary storage		
Logs		Logs for vault processes. Logs are available at the vault and job levels in the hierarchy.

## 2.2 View worker node information

In the Director UI, you can view information for the server where all vault components are installed (also known as a worker node).

You can view the state and status, internal and external address, operating system and software version of a worker node. Because information for a worker node is saved in the SQL database, the information can appear in the Director UI even if the node is offline.

A worker node can have any of the following states:

- Online. The node is available to perform Director processes.
- Rampdown. The node is finishing the current processes before it goes offline. No new processes will start on the node before it goes offline.
- Offline. The node is not available to perform Director processes, and will not accept requests.

If a node is changing between states, both states are shown in the Director UI. For example, if a node is changing from Offline to Online, the following information appears in the State column for the node.

Hostname	State
 VaultNode	 Offline >> Online

A worker node can have any of the following statuses:


- Healthy. All Director services are running on the worker node.
- Not healthy. Some Director services are not running on the worker node.
- Service not running. A specific Director service is not running on the worker node.
- Not responding. The worker node has not recently contacted the SQL node.

To view worker node information:

1. In the left pane of the Director UI, click the + sign beside the vault.




The vault expands to show the following items:



2. Click **Nodes**.

The right pane of the screen shows information about the node where all vault components are installed.



To determine which services are running, point to the node's Status column.

A tooltip indicates which Director services are running and stopped.

## 2.3 Add a vault connection

In a workspace, you can add connections to one or more vaults that you want to manage. A vault connection includes the IP address or host name of the computer where the vault is running, and credentials for connecting to the vault computer.

When vault connections are saved in a workspace, you do not have to re-enter vault connection information each time you want to access or manage the vaults.

You can add vault connections in the default workspace (named MyWorkspace) that is opened when you first start the Director UI. You can also create workspaces. See [Add a workspace](#).

When you create a vault connection to a standard vault (i.e., not a Satellite vault), the credentials that you provide for connecting to the vault computer determine your level of access to the vault. The following table lists and describes available access levels, and shows the credentials required for accessing the vault.

*Note:* Operator access is not allowed for Satellite vaults. Satellite vault management operations are performed via the Base vault.

Vault access level	Description	Credentials required
Operator	Operator permissions allow you to access all functionality for managing the vault.	Valid credentials in one of the following groups: <ul style="list-style-type: none"> <li>local VaultAdmins group. The VaultAdmins group must be created manually.</li> <li>local Administrators group.</li> </ul>
SalesPerson	Salesperson permissions provide limited, and primarily read-only, access to the vault.	Valid credentials in a local Sales group. The Sales group must be created manually.
No vault access	No permission to access the vault.	Credentials that do not belong to any of the groups listed above.

To add a vault connection:

1. From the File menu, choose **New Vault Connection**.

The Vault Connection Properties dialog box appears.

2. Enter the following vault connection information:

- In the **Description** field, enter a name for the vault connection.
- In the **Network Address** field, enter the IP address or hostname (DNS) for the vault.

3. Specify credentials for connecting to the vault by doing one of the following:

- To connect to the vault using the current Windows user's credentials, select **Use Windows Authentication**.

*Note:* When you are logged in to a vault server as a member of the local Administrators group, to connect to the local vault using Windows Authentication, you must either run the Director UI as an Administrator or add the user to a local VaultAdmins group. The VaultAdmins group must be created manually.

- To connect to the vault using a default domain account that can be used to connect to multiple vaults in the domain, select **Use Default Credentials**.
- To enter a domain account and password to use as default credentials, click **Edit** in the dialog box or, from the **File** menu, choose **Set Default Credentials**. In the Default Credentials dialog box, specify the domain account information. To save the password so that the vault can be accessed without re-entering the password, select **Save Password**. Click **OK**.

**IMPORTANT:** If you save the password for the default or custom credentials, be sure to encrypt the workspace and pick a unique password so that an unauthorized person cannot use the workspace to gain vault access. See [Encrypt a workspace](#).


*Note:* When **Save Password** is selected, the password is not stored in the workspace file. If you copy the workspace to another computer, you must re-enter the credentials. You can have more than one workspace with saved credentials on the same system.

- To specify a custom account for connecting to the vault, select **Use Custom Credentials**. In the **Domain\Username** and **Password** fields, specify the account information. To save the password so that the vault can be accessed without re-entering the password, select **Save Password**. Click **OK**.

- Click **Get Status** to test the vault connection.

If the authorization information is incorrect, a *Could not authorize the request* message appears.

If the DNS or IP information is incorrect, the following message appears:



If the vault connection information is correct, the Vault Status dialog box shows the following information:

Information	Description
Type	Name of the Director software.
Version	Director vault version and build number.
Operating system	Operating system of the machine where the vault is running.
System name	Name of the machine where the vault is running.
Address	Host name or IP address of the machine where the vault is running.
Build date	Date and time of the Director software build.
Unique identifier	System-generated identifier for the vault.
Access level	<p>Access level for connecting to the vault with the specified credentials. The access level can be:</p> <ul style="list-style-type: none"> <li>Operator. Operator permissions allow you to access all functionality for managing the vault and are granted for valid credentials in the local Administrators group.</li> <li>Salesperson. Salesperson permissions provide limited, and primarily read-only, access to the vault, and are granted for valid credentials in a local Sales group. The Sales group must be created manually.</li> </ul>

- Click **OK**.

## 2.4 Set workspace options

By setting workspace options, you can specify which logs you can view in the workspace and the frequency of refreshing vault data in the Director UI.

To set workspace options:

1. In the left pane of the Director UI, select a workspace.
2. From the View menu, choose **Options**.
3. In the Options dialog box, change one or more of the following options:

Option	Description
Automatically reload last workspace on startup	If selected, the workspace that is loaded when you close the Director UI is opened the next time you start the Director UI.
Auto-refresh tree for selected vault every	Specifies how often in minutes that data is updated in the left pane of the Director UI. You can enter values from 60 to 32767. The default value is 60.
Update executing jobs information every	Specifies how often in seconds that information about running jobs is updated in the Director UI. You can enter values from 60 to 32767. The default value is 30.
Show Spawn Logs and License Logs	If selected, Spawn Logs and License Logs are listed for each vault.
Show Replication Logs	If selected, Replication Logs are listed for each vault.
Show Server Logs	If selected, Server Logs are listed for each vault.
Default text viewer	Specifies the default text viewer for opening log files and other reports.

4. Click **OK**.

## 2.5 Encrypt a workspace

Workspaces contain system names, usernames and passwords that allow vault access. You can encrypt a workspace to prevent unauthorized access.

You must enter a password each time you open an encrypted workspace in order to protect all other passwords and access to user information on the vault configured in that workspace. The password is set on a workspace in the Director UI. Users that have access through another workspace may also have access to the same passwords on the same vaults.

*Note:* Encrypting a workspace does not affect backup data in any way.

*Note:* If you forget a workspace password, you will need to re-create the workspace. You do not need the encryption password to delete the workspace.

To encrypt a workspace:

1. Right-click a workspace in the left pane of the Director UI, and choose **Encrypt Workspace** from the menu.

2. Enter your existing password in the **Old Password** field.
3. Select an encryption type in the **Encryption Type** list.
4. Enter a new password in the **New Password** field. If you forget this password, you must recreate the workspace.
5. Enter the new password in the **Confirm Password** field.
6. Click **OK**.

## 2.6 Add a workspace

The default workspace in the Director UI is named “MyWorkspace”. In addition to adding vault connections in this default workspace., you can create other workspaces with vault connections. For example, if you manage vaults for more than one company, you could create a separate workspace for each company’s vaults. However, you can only have one workspace open at a time.

After creating a workspace, you can add vault connections in the workspace. See [Add a vault connection](#).

To add a workspace:

1. From the **File** menu, choose **New Workspace**.
2. Select the workspace in the left pane of the Director UI.
3. From the **File** menu, choose **Save Workspace As**.
4. Enter a name for the workspace in the **File Name** field.
5. Click **Save**.

## 2.7 Open a workspace

To open a workspace:

1. From the **File** menu, choose **Open Workspace**.
2. Browse to the location of the workspace (.vmw) file. By default, workspace files are saved in the user’s Documents folder (e.g., \Users\*user*\Documents\Carbonite Server Backup).
3. Click **Open**.

## 2.8 Rename a workspace

To rename a workspace:

1. Select a workspace in the left pane of the Director UI.
2. From the **File** menu, choose **Save Workspace As**.
3. Enter a name for the workspace in the **File Name** field.
4. Click **Save**.



## 2.9 Delete a workspace

To permanently delete a workspace, you must remove the workspace file. Workspace files are saved with the .vmw (Vault Manager Workspace) extension.

To delete a workspace:

1. In Windows Explorer, browse to the location of the workspace (.vmw) file. By default, workspace files are saved in the user's Documents folder (e.g., \Users\*user*\Documents\Carbonite Server Backup).
2. Right-click the workspace file and select **Delete**.
3. Click **Yes**.

### 3 Manage licensing

Each Director vault must have a vault license. Additional licenses are required for Agents, Plug-ins, replication and other features.

When an agent connects to a vault, Director automatically supplies the license. To run a backup, there must be enough licenses on the vault. If an agent connects to a vault and a license is unavailable, the backup for the agent fails.

A quota system is used to control Agent licensing. You can use a quota to limit a customer's access to features and functionality. See [Set license quotas for a customer](#).

If you have insufficient licensing on the vault, the following error messages might appear when you connect an agent to a vault:

*Vault storage limit exceeded.*

*Vault limit for Agent type exceeded, or type not found.*

*Vault limit for Plug-In type exceeded, or type not found.*

*Customer Quota for Plug-In type exceeded.*

The following table shows which processes are allowed and disallowed when a vault license is invalid or when a limit or customer quota is exceeded:

Process	If vault license is invalid	If vault storage limit is exceeded	If vault limit for Agent type is exceeded	If vault limit for plug-in type is exceeded	If customer quota for plug-in type is exceeded
Agent Registration	Allow	Allow	Disallow	Allow	Allow
Job Creation	Disallow	Allow	Warn	Warn	Warn
Backup	Disallow*	Allow	Disallow*	Disallow*	Disallow*
Restore	Disallow	Disallow	Disallow*	Disallow*	Disallow*

\* If the Agent already has a claim on the necessary licenses (it has previously done a backup of that type), the backup or restore operation is allowed.

If the Agent count exceeds the activation limit, tasks created by new agents cannot be run. Existing backups, restores, imports, and exports are allowed. Agents registered on computers without tasks, are not considered in the total number of Agents calculation.

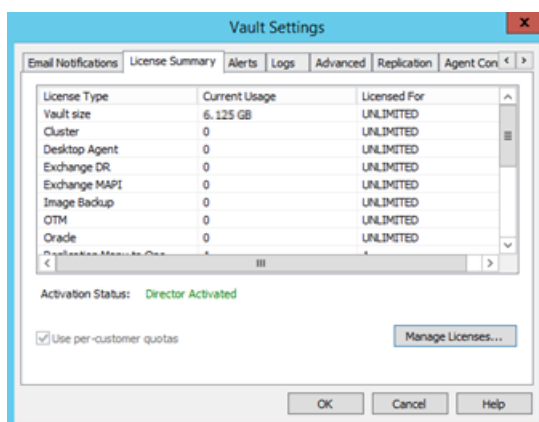
### 3.1 View licensing information

Each vault has a vault license. Additional licenses are required for Agents, Plug-ins, replication and other features. You can use a quota to limit a customer's access to features and functionality. See [Set license quotas for a customer](#).

To view license information:

1. From the **Vault Maintenance** menu, choose **Vault Settings**.
2. In the Vault Settings dialog box, click the **License Summary** tab.

If you have not activated a license, a warning message appears in the **Activation Status** field. See [Activate a license online](#) or [Activate a license manually](#).



3. Click **OK**.

### 3.2 Add a license key

To add a license key:

1. From the **Vault Maintenance** menu, choose **Vault Settings**.
2. In the Vault Settings dialog box, click the **License Summary** tab.
3. Click **Manage Licenses**.
4. Enter or paste the license key in the top pane.
5. Click **Add License Key(s)**.
6. Click **OK**. The license key is now added to the list of available licenses that can be activated. The message will appear, "The License Keys have not been activated yet. Do you wish to activate now?"
7. To activate the newly added license, click **Yes**. All licenses in the License Key list will be activated.

*Note:* If you click **No** when asked "The License Keys have not been activated yet. Do you wish to activate now?", the newly added license key will be removed from the License Key list.

### 3.3 Activate a license online

To activate a license online, the computer must be connected to the Internet.

To activate a license online:

1. From the **Vault Maintenance** menu, choose **Vault Settings**.
2. In the Vault Settings dialog box, click the **License Summary** tab.
3. Click **Manage Licenses**.

The Licenses dialog box appears.

4. Make sure the license key you want to activate is in the License Key list. You do not need to select a license key as all licenses in the License Key list will be activated.
5. Click **Online Activation**. If your activation is unsuccessful, contact your service provider.
6. Click **OK**.

### 3.4 Activate a license manually

To activate a license manually:

1. From the **Vault Maintenance** menu, choose **Vault Settings**.
  2. On the License Summary tab, click **Manage Licenses**.
- The Licenses dialog box appears.
3. Select a license key in the **License Key** list.
  4. Click **Manual Activation**.
  5. Click **No**.
  6. Click **Save to Text File** to save the activation request as a text file.
  7. Click **OK**.

8. Open your Internet browser and navigate to the Carbonite Server Backup Licensing site.
9. Click **Browse** and browse to the location of your activation request text file.
10. Click **Activate License**.
11. Click **Download Activation Key**.
12. Activate the license manually and select **Yes** when the **Do you have an activation code for this specific computer** prompt appears. See [Enter a license activation code](#).

### 3.5 Request a license activation code

After requesting a license activation code from your service provider, you can enter the code. For more information, see [Enter a license activation code](#).

To request a license activation code:

1. From the **Vault Maintenance** menu, choose **Vault Settings**.
2. In the Vault Settings dialog box, click the **License Summary** tab.
3. Click Manage Licenses.  
The Licenses dialog box appears.
4. Click Manual Activation.  
A message asks if you have an activation code.
5. Click **No**.  
The Manual License Activation dialog box shows activation request text.
6. Click **Save To Text File**. In the Save As dialog box, choose a location for the text file and click **Save**.
7. Click **OK**.
8. Send your activation request to your service provider to get an activation code.

### 3.6 Enter a license activation code

This procedure describes how to enter a license activation code. You must request the license activation code from your service provider. For more information, see [Request a license activation code](#).

To enter a license activation code:

1. From the Vault Maintenance menu, choose Vault Settings.
2. On the License Summary tab, click Manage Licenses.  
The Licenses dialog box appears.
3. Click **Manual Activation**. A message asks if you have an activation code.
4. Click **Yes**. The Manual License Activation dialog box shows activation request text.
5. Click **Open From File** to paste the activation code from a text file.
6. Click Activate Now.
7. Click **OK**.

### 3.7 Remove a license key

Before you remove a license key, verify that it is not in use. See [View licensing information](#).

To remove a license key:

1. From the Vault Maintenance menu, choose Vault Settings.
2. On the License Summary tab, click Manage Licenses.  
The Licenses dialog box appears.

3. Select a license key in the **License Key** list.
4. Click Remove License Key(s).
5. Click **Yes**.
6. Click **OK**.

### 3.8 Use customer license quotas

You can use quotas to limit the amount of storage and number of Agent and Plug-in licenses that a customer can use. If a vault will be used as a Base vault in N:1 replication, or as an Active Base vault in N:1:1 replication, you must set customer quotas.

If you are using Satellite vaults, customer quotas need the correct Agent type. If the customer quota does not allow sufficient Desktop Agents, the license claim fails even if there is a customer quota available for Server Agents. To adjust customer quotas on satellites, you must make the adjustments on the Base vault (BAV).

To set quotas for a specific organization, see [Set license quotas for a customer](#).

To use customer license quotas:

1. From the Vault Maintenance menu, choose Vault Settings.
2. Click the **License Summary** tab.
3. Select a license type in the **License Type** list.
4. Click Use per-customer quotas.

If you are using N:1 or N:1:1 replication, this option is enabled by default and cannot be changed.

5. Click **Yes**.
6. Click **OK**.

### 3.9 Set license quotas for a customer

You can use customer license quotas to limit the amount of storage or number of Agents and plug-ins that a specific customer can use. If a vault will be used as a Base vault in N:1 replication, or as an Active Base vault in N:1:1 replication, you must set customer quotas.

To set license quotas for a customer:

1. Expand a vault in the left pane of the Director UI.
2. Expand Manage Customers/Orgs, Safesets, Tasks.
3. Right-click a customer and select **Properties**.
4. Click the **Quotas** tab.
5. Select a feature (Storage, or a type of Agent or plug-in) and click **Set Quota**.

6. Select **Unlimited** or enter a quota number in the **Set quota** area.

You can assign a quota number that exceeds the available licenses. If the license quota is reached, the following message appears: *The total quota for this feature exceeds the amount licensed for the vault. Agents may fail even though customer level quota is not exceeded, if total usage exceeds the licensed amount.*

7. Click **OK**.
8. Click **OK** again.

### 3.10 View the license quota for a computer

To view the license quota for a computer:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Expand a **Location**.
6. Expand the Registered Computers list.
7. Right-click a registered computer and select **Properties**.
8. Click the **License Quotas** tab.
9. Click **OK**.

### 3.11 Reclaim a license

A reclaimed license is assigned back to the vault license pool and can be claimed by another computer.

To reclaim a license:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Expand a **Location**.
6. Expand the Registered Computers list.
7. Right-click a registered computer and select **Properties**.
8. Click the **License Quotas** tab.
9. Select a quota and click **Reclaim**.
10. Click **OK**.

## 4 Manage primary storage

When an agent sends backup data to a vault, the data is saved in primary storage. Primary storage is typically the fastest and most active storage in a vault.

Data is added to primary storage by backups and removed through optimization and migration as data expires based on retention settings. The constant addition and removal of data can result in fragmented data pool files and inefficient disk use. To improve disk storage usage and reduce fragmentation, the vault performs maintenance, including expiring safesets, removing unused data and preparing optimized indexes for better performance. See [Automate maintenance processes](#).

To improve backup performance and storage efficiency, consider the following suggestions:

- Add storage space before insufficient space causes your backup to fail.
- Verify that there is enough storage space to reseed your largest backup.
- Specify preferred storage locations in the primary storage group for index, data and log files. Faster storage for indexes will significantly improve Director performance. See [Add and edit policies for primary storage locations](#).

### 4.1 Add storage locations to the primary storage group

You can add multiple storage locations in the primary storage group. You can add local and UNC storage locations.

*Note:* Only one primary storage group is allowed. You can add primary storage locations but you cannot add a primary storage group.

When primary storage group locations are on different volumes, Director distributes files across storage locations based on available free space, and adds pool files to locations with more free space.

To add a storage location to the primary storage group:

1. Select a vault in the left pane of the Director UI.
2. In the Storage Management menu, click Storage Groups and Locations.
3. In the Storage Locations dialog box, on the **Primary Storage Groups** tab, select a storage group.
4. Click Add Location.
5. Enter or select a path in the **Path** field.

The path can be a maximum of 128 characters in length, including spaces and slashes (\). Enter a UNC path as `\\hostnameOrIPAddress\share`.

We recommend specifying the hostname in the UNC path rather than the IP address. This will allow DNS to handle IP address changes.

6. Select a storage policy for the storage location from the **Storage Policy** list.



A storage policy specifies the amount of disk space to use in the storage location, and whether the storage location is preferred for index, data or log files. See [Add and edit policies for primary storage locations](#).

7. Click **OK**.
8. Click **OK** again.

## 4.2 Add and edit policies for primary storage locations

For primary storage locations, you can create policies that specify:

- the amount of free disk space to leave in the primary storage location
- the maximum percentage of disk space to use in the primary storage location

If vault storage usage is above this threshold before a backup, or reaches this threshold during a backup, the backup fails or is terminated. A cleanup operation then verifies safesets for the task and removes orphaned files.

*Note:* You can set up email notifications to warn you when storage utilization reaches specific thresholds, and run the Storage Pool summary report to determine whether space can be reclaimed through optimization. See [Set primary storage thresholds](#) and [Create a Storage Pool Summary report](#).

- whether the storage location is a preferred location for index, data or log files

To improve performance, you can specify that faster storage is preferred for index files, while slower storage is preferred for data files. Index files require 5-10% of the storage in a vault. Faster storage for indexes will significantly improve Director performance.

After creating a location policy, you can assign it to one or more locations in the primary storage group. For more information, see [Assign a location policy to a storage location](#).

To add or edit a location policy for primary storage groups:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.  
The Storage Locations dialog box appears.
3. On the Primary Storage Groups tab, click Locations Policies.  
The Policies dialog box appears.
4. Do one of the following:
  - To add a location policy, click **Add**.
  - To edit a location policy, select a location policy in the list, and click **Edit**.The Storage Policy dialog box appears.
5. Complete the following fields:
  - **Policy Name** — A name for the policy. You cannot edit this name after you create the policy.

- **Stop using this volume if free disk space on this volume falls below x GB** — The amount in GB of free disk space to leave on the volume. Typically, the free space is left for the operating system to use. You can enter values from 0 to 4,194,303.
  - **Use no more than x % of disk space on this volume** — The maximum percentage of disk space to use on the volume.
  - **Use for creation of pool index files** — If selected, the storage location is a preferred location for saving index files. Index files have the following extensions: .file-names, .files, .segments, .stream-names, .streams, .names, .dat, .mir, .pfblocksizecache
  - **Use for creation of pool data files** — If selected, the storage location is a preferred location for saving pool data files. Pool data files have the following extensions: .pf, .pf-map
  - **Use for creation of log files** — If selected, the storage location is a preferred location for saving log files. Log files have the .log extension.
6. Click **OK**.

### 4.3 Assign a location policy to a storage location

After creating a location policy, you can add the policy to one or more locations in the primary storage group.

For information about creating location policies, see [Add and edit policies for primary storage locations](#).

To assign a location policy to a storage location:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.  
The Storage Locations dialog box appears.
3. On the **Primary Storage Groups** tab, select a storage location.
4. Click Location Settings.  
The Location dialog box appears.
5. In the **Storage Policy** box, select a storage policy.
6. Click **OK**.

### 4.4 Change credentials for UNC storage locations

You can change the credentials used by the vault for accessing a UNC storage location.

When a connection to a UNC location has already been established, the user credentials cannot be verified because a cached connection is used. By default, Microsoft Windows caches network connections.

To change network storage credentials:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Network Storage Credentials.

3. On the first page of the Change User Credentials wizard, click **Next**.
4. Do one of the following:
  - To change the username and password at multiple locations, select **Replace Credentials associated with a Specific Location**, and then click **Next**.
  - To change the username and password for a specific user at multiple locations, select **Replace Credentials associated with a Specific Username**, and then click **Next**. The new credentials should allow the same access to the secondary pools.
  - To replace the default username and password for a UNC path, select **Replace Default Credentials**, and then click **Next**.
5. Complete any of the following fields that appear:

Field	Description
Select ALL	Applies the new settings to all locations.
Select one or more locations where you want to change the credentials	Applies the new settings to specific locations that you select.
Use Default Credentials	Applies the default username and password to the locations you specify.
Create New User Credentials	Creates a new username and password for the locations you specify.
Username	A name that identifies the account user.
Password	The password for the account user.
Test connection under these credentials to	The UNC path to test.
Test Credentials	Tests the UNC path credentials.

6. Click Finish.

## 4.5 Change the maximum pool file size for the primary storage group

Director saves backup data in pool files. You can change the maximum pool file size for the primary storage group. When the pool file size value is met, Director creates a new pool file to accept safeset data.

The pool file size setting must be less than your primary location volume size. If your volumes are smaller than your pool file size, backups can fail when the maximum volume size is exceeded.

To change the maximum pool file size for the primary storage group:

1. Select a vault in the left pane of the Director UI.

2. From the Storage Management menu, choose Storage Groups and Locations.

The Storage Locations dialog box appears.

3. On the **Primary Storage Groups** tab, select the storage group.

4. Click Group Settings.

The Edit Location Settings dialog box appears.

5. In the **Pool file size** field, enter the maximum pool file size.

You can enter values from 10 to 2000. The default is 512 MB. When the pool file size value is met, the Director creates a new pool file to accept safeset data.

The pool file size setting must be less than the primary location volume size. If volumes are smaller than the pool file size, backups can fail when the maximum volume size is exceeded.

6. Click **OK**.
7. Click **OK** again.

## 4.6 Set the minimum pool usage parameter for the primary storage group

The optimizer uses the Minimum Pool Usage parameter to determine whether to defragment pool files. Director defragments a pool file when storage usage in the file is lower than the specified value. For example, if the minimum pool usage value is 75, the system defragments a pool file and recovers unused storage space in the file when its space usage falls below 75%.

To set the Minimum Pool Usage parameter for the primary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.

The Storage Locations dialog box appears.

3. On the **Primary Storage Groups** tab, select the storage group.

4. Click Group Settings.

The Edit Location Settings dialog box appears.

5. In the **Minimum pool usage** field, enter the minimum percentage of storage used in a pool file. You can enter values from 5 to 95. The default is 75.
6. Click **OK**.

## 4.7 Enforce retention settings in primary storage

Migration is a Director maintenance process that identifies which safesets in primary storage should be kept based on retention settings and retention groups. Other safesets are expired, and marked for deletion.

Beginning in Director 8.60, 'vvpoolop optimize' is sometimes triggered after backups and replications. Migration runs as part of this process. See [Triggered maintenance and replication](#). In previous Director versions, migration was scheduled to run on the entire vault once a day.

When running or scheduling a backup, a user specifies retention settings for the resulting safeset. Retention settings for a safeset include the following:

1. **Online days** — Number of days that the safeset should be kept online on the vault
2. **Online safesets** — Number of safesets in the task's retention group that should be kept online
3. **Archive days** — Indicates whether/how long the safeset should be stored in offline storage. A value of zero (0) indicates that that the safeset will not be archived or stored offline.

Based on its retention settings, each safeset is assigned to a retention group in Director. A retention group includes safesets for a task that have the same retention settings enforced during migration. For example, in a single task, safesets that have the default "Daily" retention settings would be in one retention group, safesets with "Weekly" retention settings would be in a second retention group, and safesets with "Monthly" retention settings would be in a third retention group. You can view a safeset's retention group number in the **Safeset Properties** dialog box. See [View and change safeset properties](#).

Director uses the following logic in each retention group to determine which safesets in primary storage to keep, and which safesets to mark for deletion:

1. Director obtains the **Online safesets** value from the most recent safeset in the retention group. Director uses this value as the **Online safesets** value for the retention group. For example, if the **Online safesets** value is 7 for safesets #1 to #7 in a retention group, and 5 for safesets #8 to #10, Director uses 5 as the **Online safesets** value for the retention group.
2. Director compares the number of safesets in the retention group to the group's **Online safesets** value. If the number of safesets in the retention group is less than or equal to the **Online safesets** value, Director keeps all of the online safesets. If the number of safesets in primary storage in the retention group is greater than the **Online safesets** value, Director continues to step 3.
3. Director finds the oldest safeset in the retention group. If the safeset is older than its **Online days** value, and its **Archive days** value equals zero (0), the safeset is marked for deletion and no longer appears in the Director UI. If the safeset is not older than its **Online days** value, or the **Archive days** value is greater than zero, the safeset is kept in primary storage.

*Note:* If the **Archive days** value for a safeset is greater than zero, the safeset will never be deleted by the migration process. The safeset will remain in primary storage unless you archive the safeset or move it to secondary storage.

4. If, after Step 3, the number of safesets in the retention group is greater than the group's **Online safesets** value, Director repeats step 3 with the next oldest online safeset. This process is repeated until the number of safesets in the retention group is equal to the group's **Online safesets** value, or no safesets in the retention group meet the criteria for deletion described in Step 3.

This migration logic is applied separately to safesets in each retention group. For example, if a task has safesets in three retention groups, Director follows this migration logic separately for each retention group.

Running regular migrations and optimizations removes unnecessary safesets and conserves storage space. It is recommended that you do not disable the Maintenance Host.

During the migration, a message is added to the log when a safeset with a backup time that is earlier than the previous entry's backup time is encountered. This entry and all subsequent entries are processed as if they were backed up in sequential order.

*Note:* You can also run an online migration using the `vvmigrat online` command. See [vvmigrat](#).

*Note:* To enforce retention policies in secondary storage, use the `secondaryop expire` command. See [secondaryop](#) and [Enable automated secondary storage maintenance](#).

You can run a migration on the entire vault, or on an individual task. If you run a migration on the entire vault, and a single task fails, the migration does not stop. When the migration finishes, it reports one of these statuses:

- Completed
- Completed with errors
- Failed

If a task fails during the migration, a Completed with Errors status message appears in the log and the email notification.

To enforce retention settings in primary storage:

1. Select an object in the left pane of the Director UI.
2. From the Storage Management menu, choose Migrate Tasks.
3. On the first page of the Migrate wizard, click **Next**.
4. On the Migrate page, select **Online**, and then click **Next**.
5. On the Action Scope page, do one of the following:
  - To run the migration on safesets for all tasks in the vault, select **Vault**, and then click **Next**.
  - To run the migration on safesets for a specific task, select **Task**, and then click **Next**. On the following pages, select the customer, computer, and task for running the migration, and click **Next**.
6. On the "When would you like to start the job" page, do one of the following:
  - To run the migration immediately, select **Submit job immediately**, and then click **Next**.
  - To schedule the migration to run, select **Schedule job**, click **Next**, and then specify a schedule for running the migration.
7. Click **Finish**.

## 4.8 Optimize pool files

When the vault receives data from its customers' agent computers, it "pools" the data with other safesets from the same task. This data pool can become fragmented as data is written and deleted and safesets are

migrated. Pool files are left with data that is no longer referenced but still occupies space in the pool system. Over time, this fragmented data grows, wasting storage space.

Vault optimization removes unreferenced data and defragments the pool system. In addition, optimization ensures that the vault retains only the required safesets. A migration is included as part of an optimization, even if you do not schedule a migration.

Scheduled maintenance processes are recommended for optimizing pool files. You can also use the Pool Optimization wizard to manage the pool files and reduce data fragmentation. This reconstructs the agent's data using better space conservation.

Pool optimization maintains a minimum storage footprint, and allows fast restores. Optimization maintains the size of the pool by removing non-essential expired or duplicate data. Every optimization tunes the database to allow optimal restore performance.

Optimize does not check to see if all the files exist, and does not verify indexes. To verify the physical status of the pool system, run the command line utility CheckCRC. If you want to check the data integrity, run the command line utilities SSIVerOne or SSIVerAll. For more information, see the [Command Reference](#).

To optimize pool files:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Optimize.
3. Complete fields in the Pool Optimization wizard:

Field	Description
Remove duplicate and expired data from the pool	Detects duplicate pool system data, adjusts references, and then uses optimization to remove unused data. This process is longer than optimization.
Remove expired data only	Removes unused data from the pool system, and allocates the recovered space. <i>Note:</i> After running an optimization with deduplication, run a Storage Pool Summary report to view the amount of space that was reclaimed.
By computer	Selects tasks on a specific computer.
Entire Vault	Selects all accounts and computers for the online storage location. Generally, this option takes longer to complete than selecting the tasks on a single computer.
Selected Organization/Customer	Specifies a customer. This option is available when you select By computer.
Selected location	Specifies a customer location. This option is available when you select By computer.

Field	Description
Selected Computer	Specifies a customer computer. This option is available when you select By computer.
Selected Task	Specifies a specific task on a customer computer. This option is available when you select By computer.
Submit job immediately	Submits the job immediately.
Schedule job	Creates a schedule to complete the job weekly or monthly.

4. Click **Finish**.

## 4.9 Deduplicate data

Deduplication (common block pooling) inside the primary pool helps to reduce the storage requirements for the vault. The process results in different disk space savings for different backup jobs. In the case of a re-seed, it can significantly reduce the size of the pool system.

If the agent changes the encryption type, or compression type, or encryption password of a safeset, then all pool blocks are different on the vault. A re-seed like this will not see any common block pooling (deduplication). The next backup of the same job with the same parameters, though, will benefit from deduplication.

The vault calculates a Vault Hash Value (VHV) for each data block, and updates the indexes to the data. Any duplicated blocks are marked for removal on the next Optimization.

The vault can detect and deduplicate situations such as the same file in different directories, renamed files, and renamed parent directories.

If deduplication starts before any commands such as Backup, Restore, Import or Export, then deduplication detects the lock request, stores its current state, releases the pool system, and retries the command later.

If you want the deduplication to run without interruption, you can run the `vvpoolop` command from the command line (CLI) with a "nointerrupt" option. Only one `vvpoolop` command can run against a pool system at any given time. For more information, see [vvpoolop](#).

## 4.10 Verify the integrity of backup files

To verify the integrity and consistency of a vault pool, and troubleshoot errors, you can create a backup verification job.

To verify the integrity of backup files:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Diagnose.
3. Complete fields in the Pool System Diagnostic wizard:



Field	Description
Check CRC integrity of the Pool System	Validates the physical integrity of files in the pool system. The CRC integrity check does not parse the data, but just checks that the pool files and indexes are intact and uncorrupted.
Verify all safeset generations	Verifies the logical integrity of all safesets. Verifying safesets reads the System Independent Data Format (SIDF) and verifies the integrity of a safeset as if it were getting ready to do a restore. A user's private information is not read. Only a user with the encryption key can read encrypted data.
Verify the last committed safeset generation	Verifies the logical integrity of the last committed safeset.
Verify a specific safeset generation	Verifies the logical integrity of a specific safeset.
Entire Vault	Selects all accounts and computers for the online storage location.
By Task	Selects a specific task on a specific computer.
Selected Organization/Customer	Specifies a specific customer. This option is available when you select Verify all safeset generations, Verify the last committed safeset generation, Verify a specific safeset generation, or By Task.
Selected Location	Specifies a specific customer location. This option is available when you select Verify all safeset generations, Verify the last committed safeset generation, Verify a specific safeset generation, or By Task.
Selected Computer	This option is available when you select Verify all safeset generations, Verify the last committed safeset generation, Verify a specific safeset generation, or By Task.
Selected Task	This option is available when you select Verify all safeset generations, Verify the last committed safeset generation, Verify a specific safeset generation, or By Task.
Submit job immediately	Submits the job immediately.
Schedule job	Creates a schedule to complete the job weekly or monthly.

4. Click **Finish**.

## 4.11 Copy and clone data

### IMPORTANT:

- To improve the security of vault-to-vault communications, vault certificates are verified when you copy or clone data between Director 8.7 vaults or copy data within a Director 8.7 vault. If the

certificate verification fails, the copy or clone operation does not proceed. For more information, see [Certificate verification and pinning for vault-to-vault communications](#).

- You can copy or clone data from Director 8.7, 8.62, 8.61 and 8.56 vaults to Director 8.7 vaults. By default, you cannot copy or clone data from other vault versions to Director 8.7 vaults. If you need to copy or clone data from another vault version, please contact Support.
- You cannot copy or clone data from a Director 8.7 vault to an earlier vault version.

Three methods are available for copying vault items and data, including customers, locations, computers, tasks and safesets:

- **Copy.** A Copy operation creates duplicate items with different globally unique identifiers (GUIDs). You can copy items and data to a destination in the same vault or in another vault. Items are disabled on the destination while the copy is in progress, but are enabled after the copy is finished.
- **Clone (File by File).** A Clone (File by File) operation creates identical items with identical GUIDs by physically copying pool files. You can only clone items to a destination in another vault. You cannot clone items to a destination in the same vault. Entire tasks are cloned when you run a Clone (File by File) operation. You cannot clone a task if it already exists on the destination vault. This clone method is faster than Clone (Safeset by Safeset), but backups, restores and replication cannot run while this operation is running. Items on the source are disabled while they are being cloned using this method, and remain disabled unless you enable them. Items on the destination are disabled while this operation is in progress, but are enabled after the operation is finished.
- **Clone (Safeset by Safeset).** A Clone (Safeset by Safeset) operation creates identical items with identical GUIDs by logically copying safesets. Backups, restores and replication can run while this operation is running. You must clone entire customers when you use Clone (Safeset by Safeset). You can only clone customers to a destination in another vault. You cannot clone customers to a destination in the same vault.

If you clone a task that exists on both the source and destination vault, but the destination vault has safesets that do not exist on the source vault, the following warning appears: *This Clone operation could result in data loss*. If you continue with the clone operation, safesets on the target vault are marked for deletion if they do not exist on the source vault. These safesets are deleted when migration next runs on the task.

Because operations such as backups can run concurrently with Clone (Safeset by Safeset), backups can arrive in the source vault while the clone operation is running. To ensure that all safesets are cloned to the destination, you might have to run the operation multiple times. A difference report runs automatically after a Clone (Safeset by Safeset) operation to indicate whether items are the same on both the source and destination, or whether some items are missing. The report is available in the VVCopy log on the source vault.

When copying or cloning data, you might have to resolve name or code conflicts. A conflict can occur, for example, if the source and destination have customers with the same name. If you are copying items, you can resolve conflicts using the Vault Copy wizard. If you are cloning items, you might need to resolve conflicts outside of the Vault Copy wizard before trying to clone the items again.

You cannot copy or clone a task if it has safesets in both secondary storage and archive storage. Using secondary and archive storage for the same task is not supported, but is not prevented by Director.

If a copy or clone is successful, fast index verification is performed on the last safeset copied to ensure the data integrity of the pool.

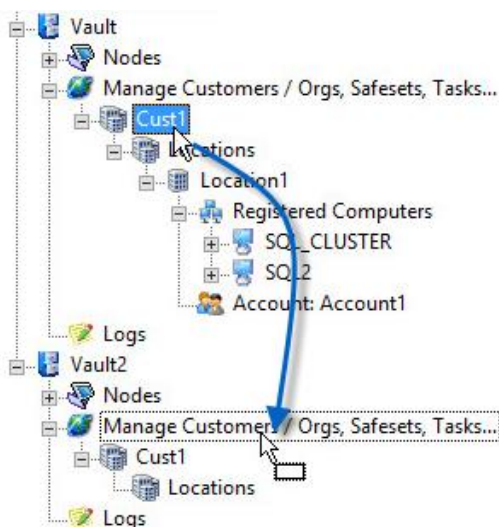
**IMPORTANT:** To ensure that copy operations succeed, be sure that your vault server system clocks are synchronized with a Network Time Protocol (NTP) server.

To copy or clone data:

1. In the left pane of the Director UI, do one of the following:
  - Drag the item that you want to copy or clone to its destination. For example, you can drag a location to a **Locations** level, or a computer to a **Registered Computers** level. On the Welcome page of the Vault Copy wizard, click **Next**.

When you drag an item to its destination, items lower in the vault hierarchy will also be copied or cloned. For example, when you copy a location, the location's computers, tasks, safesets and accounts are also copied.

*Note:* If you want to run a Clone (safeset by safeset) operation, you must drag a customer to the destination.

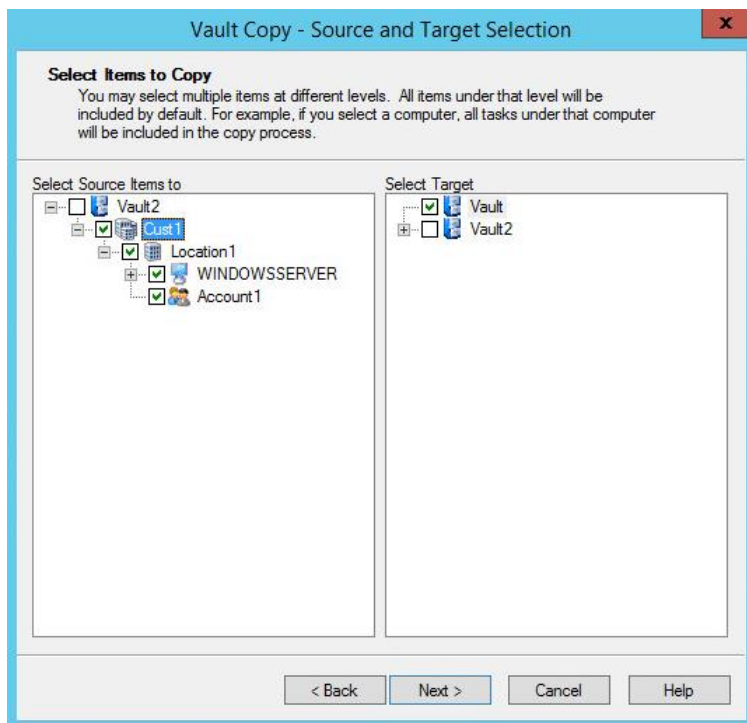


- Select an item that you want to copy or clone. From the **Storage Management** menu, choose **Copy**. On the Welcome page of the Vault Copy wizard, click **Next**. On the Select Items to Copy page, do the following:
  - In the **Select Source Items** area, select the check box for each item that you want to copy or clone. You can select multiple items at the same level in the hierarchy (e.g., multiple customers or multiple computers).

*Note:* If you want to run a Clone (safeset by safeset) operation, you must select one or more customers to clone.

- In the **Select Target** area, select the check box for the destination. The destination must be the appropriate level for the item in the vault hierarchy. For example, you can copy or clone a customer to a vault, or a computer to a location.

2. Click **Next**.



3. On the Connection Information page, complete the fields shown in the following table and then click **Next**.

Host Name	The host name of the destination vault server.
Port	The port used by the destination vault server.
User Name	The user name used to access the destination vault server.
Password	The password used to access the destination vault server.
Confirm password	The password used to access the destination vault server.
Domain	The domain of the destination vault server.

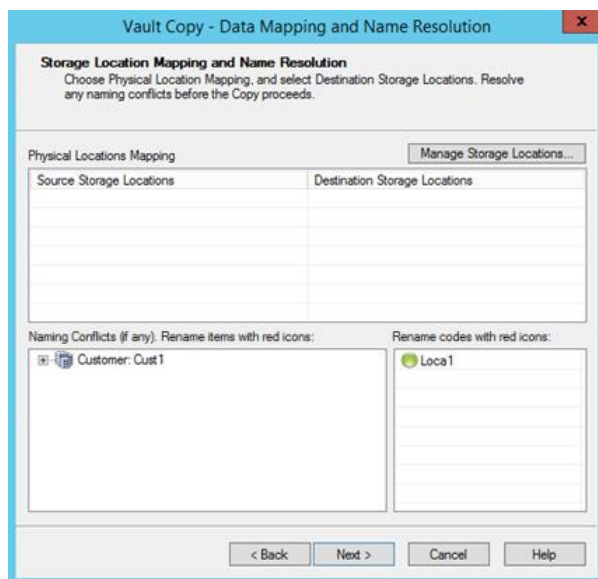
4. On the Copying Options page, select one of the following options:


- **Copy.** Select this option to create duplicate items. The resulting items are not identical.
- **Clone (file by file).** Select this option to create identical items by physically copying pool files. This clone method is faster than Clone (safeset by safeset) but backups, restores and replication cannot run while this operation is running.

- **Clone (safeset by safeset).** Select this option to create identical items by logically copying safesets. This clone method is slower than Clone (file by file), but backups, restores and replication can run concurrently with this operation. This option is disabled if the source item is not a customer.
  - I have both pool data files and meta information already in their respective locations and I only want to adjust the target files to the proper locations. Select this option to update the file paths and names of copied pool files. This advanced option is only used when pool indexes and data were manually moved from the source to the destination.
5. In the Data options area, select one or more of the following:
- **Copy pool data files.** Select this option to copy or clone backup data (safesets).
  - **Copy meta information (does not include pool data files).** Select this option to copy or clone item information, but not backup data (safesets).
  - **Do not copy pool data files from secondary pools, just link.** Safesets in the secondary pools will be marked as detached until you manually attach the secondary storage. See [Attach secondary storage pools](#).
6. Click **Next**.

*Note:* Director might appear unresponsive while gathering information and checking for name conflicts. This process can take several minutes, depending on the data being processed.

The Storage Location Mapping and Name Resolution page lists items that you are copying or cloning.



7. If a Stop icon (  ) appears beside an item, a name or code conflict exists. For example, if a Stop icon appears beside a customer, the source and destination have customers with the same name. Do one of the following:
- If you are copying items, resolve name and code conflicts on the page. Select the item with the conflict, and press **F2** to enter a new name or code.
  - If you are cloning items and need to resolve conflicts, click **Cancel** in the Vault Copy wizard. Resolve conflicts outside of the Vault Copy wizard before trying to clone the items again.

*Note:* If you are cloning items using the Clone (safeset by safeset) method and have already cloned some of a customer's tasks and safesets, the Stop icon appears beside the customer. However, you do not need to resolve the name conflict. You can continue the operation and clone the customer's remaining safesets.

8. If the Physical Locations Mapping area shows source storage locations, you are copying a task that has safesets in secondary or archive storage. Click **Manage Storage Locations**. In the Storage Locations dialog box, create or select a destination storage location for each source storage location, and then click **OK**.

9. Click **Next**.

If you are cloning items, and a message box indicates that there are billing code, customer or account conflicts, you might have to exit from the Vault Copy wizard and resolve the conflicts before starting the clone operation again.

10. Click **Next**.

11. Click **Finish**.

To determine whether the copy or clone was successful, review the VVCopy log. If you ran a Clone (safeset by safeset), the VVCopy log on the source vault includes a difference report. This report indicates whether items are the same on both the source and destination, or whether some items are missing.

## 4.12 Retire primary storage locations

To move pool data from a drive that is old, small or slow to other locations in the primary storage group, you can retire the primary storage location. Unless you specify a destination for the files, the system moves files to the storage location with the most available space.

*Warning:* Do not use this procedure to move pool files to the same physical storage location with a new UNC path, or your pool system will be corrupted.

To retire a primary storage location:

1. Set the location as read-only, using the Storage Locations dialog box. This prevents new information from being written to the retired drive.
2. Stop or pause the Maintenance Host on the vault. If the Maintenance Host runs during the drive retirement process, maintenance processes could fail or cause problems.
3. Select an object in the left pane of the Director UI.
4. From the **Storage Management** menu, choose **Move to Another Location**.
5. Complete the Move Data/Retire a failed storage location wizard.
6. Click **Finish**.
7. Restart the Maintenance Host on the vault.

## 4.13 Remove storage locations from the primary storage group

You can remove a storage location from the primary storage group if it does not have associated data. You must first move all data from the storage location. You cannot delete the last location assigned to the last primary storage group.

To remove a storage location from the primary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. Click the Primary Storage Groups tab.
4. Click Remove Location.
5. Click **Yes**.
6. Click **OK**.

## 4.14 Delete a location policy

*Note:* You cannot delete a location policy if it is assigned to a storage location with associated tasks.

To delete a location policy:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. On the Primary Storage Groups tab, click Locations Policies.
4. Select a location policy.
5. Click **Delete**.
6. Click **Yes**.

## 5 Manage secondary storage

In addition to primary storage, vaults can include secondary storage. Typically, secondary storage is cheaper and slower than primary storage, and is used for long-term data storage. Secondary storage can be online, detached and stored (offline), or attached to another system.

A secondary pool is self-contained. It includes both pool data and pool indexes in one directory. The first safeset in the secondary pool is a seed, and the rest of the safesets (for that task) are deltas.

When you create a new secondary pool, the existing secondary pool becomes inactive. However, the safesets in the original secondary pool have a Secondary status and remain online.

You can use the Secondary Restore Server application to read or export a safeset directly from a secondary pool, without the participation of the primary vault.

Data stored in secondary storage cannot be replicated to another vault.

### Manage safesets and pools in secondary storage

You can select aggressive settings for migration rules that require migration to secondary storage. This could result in a significant increase in the amount of disk space required for secondary storage.

After a Migrate to Secondary, on a large seed safeset with many files, the Secondary Pool size may have a size increase. If you have a large safeset with few files (such as a database) you will not notice much difference. However, if you have many files (such as a MAPI backup) you may notice an increase in size.

### 5.1 Add secondary storage groups

To add a secondary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. Click the Secondary Storage Groups tab.
4. Click New Group.
5. Enter a name for the group.
6. Click **OK**.
7. Add a location to the group. See [Add storage locations to a secondary storage group](#).
8. Click **OK**.

### 5.2 Add storage locations to a secondary storage group

You can add multiple locations to a secondary storage group.

When you first migrate safesets to secondary storage, Director saves the safesets in the storage location with the most available space. Director then saves safesets in this location until the location is closed.



To add a storage location to a secondary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. Click the Secondary Storage Groups tab.
4. Select a storage group.
5. Click Add Location.
6. Enter, select, or browse to a path in the **Path** field. You can enter a Universal Naming Convention (UNC) path for the new location. The path must meet these requirements:
  - The path must be absolute, not relative.
  - The path cannot contain an invalid name.
  - The path cannot contain invalid characters.
  - UNC paths must be available at all times.
7. Click **OK**.
8. Click **OK** again.

### 5.3 Assign secondary storage groups to tasks

Before data from a task can be moved from primary to secondary storage in a vault, a secondary storage group must be assigned to the task.

You can assign a secondary storage group to a specific task. You can also assign the same secondary storage group to all tasks for a vault, customer or location. The secondary storage group is then automatically assigned to subsequent tasks for the vault, customer or location.

This procedure describes how to assign a secondary storage group to an existing vault, customer, location or task. You can also assign a secondary storage group to a customer, location or account when you create it. See [Add a customer, location, account and user](#).

To assign a secondary storage group to a task:

1. From the **Storage Management** menu, choose **Secondary Storage Pools**, and then click **Assign**.  
The Global Scope Configuration wizard appears.
2. Click **Next**.

3. Complete fields in the Global Scope Configuration wizard:

Field	Description
Change Locations	Add locations and groups.
Override the Secondary Storage assignment for existing tasks which already have a Secondary Storage assignment	Overwrites the existing secondary storage settings for existing tasks. Changing the settings causes a re-seed when the next backup runs. Selecting this option can use more disk space.
Assign this Secondary Storage Group only to existing tasks, which currently do not use the Secondary Storage	Assigns the secondary storage settings to existing tasks that do not currently use secondary storage.
Assign this Secondary Storage Group to all new and existing tasks which currently do not have a Secondary Storage assignment	Assigns the secondary storage settings to all new and existing tasks that do not currently use secondary storage. This option is available when you select a vault and all customers, computers, and tasks.

4. Click **Finish**.
5. Right-click a task in the Director UI to which you applied secondary pool settings, and select **Properties** to see if the changes were applied.

## 5.4 Migrate safesets to secondary storage

You can move safesets from a primary storage location to a secondary storage location. When you migrate the first safeset for a task into a new pool in a secondary storage location, the data is a full seed and could result in a significant increase in the amount of disk space required for secondary storage.

After you move a safeset from within a group of safesets, you cannot move previous safesets to secondary storage.

You cannot migrate safesets from a secondary pool to a primary pool or from one secondary pool to another.

*Note:* Do not move safesets for the same task to both secondary storage and archive storage. Using secondary and archive storage for the same task is not supported, but is not prevented by Director.

To migrate safesets to secondary storage:

1. In the left pane of the Director UI, select the vault with safesets that you want to migrate to secondary storage.
2. From the Storage Management menu, choose Migrate to Secondary Storage Pool.

## 3. Complete fields in the Migrate to Secondary Storage Pool wizard:

Field	Description
Use Rules to Determine Which Safeset To Migrate	Migrates safesets that meet specific retention settings.
Explicitly Select Online Safesets To Migrate To Secondary Storage	Migrates safesets that you select to secondary storage.
Use the Archive Retention settings	Migrates safesets when the archive retention settings for online copies or days are reached.
Migrate all the safesets that are older than x days	Migrates safesets that are older than a specified number of days. The default is 180 days.
Keep up to x of the most recent safesets online. Migrate the rest.	The number of safesets to keep in the primary pool. When the number is exceeded, the older safesets are moved to secondary storage. The default is 100.
Detach the migrated safesets	Identifies migrated safesets as offline to the agent.
I want to preview which safesets are going to be migrated	Applies the retention rules to the safesets in the primary pool so you can see which safesets would be migrated.
Also preview the storage size required on the secondary storage pool	Applies the retention rules to the safesets in the primary pool so you can see how much storage would be required for the migrated safesets.
Deduplicate data. Leave this option selected to reduce secondary pool bloat.	Migrating an already deduplicated pool with the Deduplicate data checkbox selected will allow the pool size to remain the same. Unchecking the Deduplicate data checkbox may cause the migrated data to dramatically increase in size.
Verify the first and then every x	Verifies safesets that are migrated to a secondary pool. Verification ensures that the data in the safeset is valid and can be used for a restore. The default value 1 indicates that every safeset is verified. It is recommended that you keep the default value.
Submit job immediately	Submits the job immediately.
Schedule job	Creates a schedule to complete the job weekly or monthly.

4. Click **Finish**.

5. Review the MigrateSecondary log to determine if the move was successful. The status of migrated backups changes to Secondary in the Director UI.

## 5.5 Change the maximum pool file size for a secondary storage group

You can change the maximum pool file size for a secondary storage group.

After you change the group settings of a secondary storage group, all affected tasks start with a new secondary pool when the next secondary migration runs. This could significantly increase the amount of space that is required for secondary storage.

To change the maximum pool file size for a secondary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. On the Secondary Storage Groups tab, select a storage group.
4. Click Group Settings.

The Edit Location Settings dialog box appears.

5. In the **Pool file size** field, enter the maximum pool file size. When the pool file size value is met, Director creates a new pool file to accept safeset data. You can enter values from 10 to 2000 MB. The default is 512 MB.

The pool file size setting must be less than the primary location volume size. If volumes are smaller than the pool file size, backups can fail when the maximum volume size is exceeded.

6. Click **OK**.
7. Click **OK** again.

## 5.6 Apply safeset retention settings to the secondary storage pool

You can apply primary pool safeset retention settings to the secondary storage pool.

To apply safeset retention settings to the secondary storage pool:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Expand Registered Computers.
7. Expand a computer.

8. Expand Backup Tasks.
9. Expand a backup task.
10. Select **Backups**.
11. Right-click a backup in secondary storage in the right pane and select **Properties**.
12. Click **Also apply setting to secondary location** on the **General** tab. This change applies only to safesets that are not detached.
13. Click **OK**.

## 5.7 Secondary storage safeset deletion

You can use the Director UI or run the `secondaryop delete` command to delete a safeset from attached secondary storage.

When you delete a safeset from secondary storage, two entries appear in the Job Monitor and log files. One entry is the primary storage deletion, the second entry is the command that deleted the safeset from secondary storage:

```
secondaryop delete taskid=<t> safeset=<s>
```

When you use the Director UI to delete a safeset and the secondary storage is detached (offline), only the safeset information on the vault is deleted. The safeset data on the secondary storage is unaffected.

When you use the Director UI to delete a safeset and the secondary storage is attached, but not connected, you must connect the secondary storage before you can delete the safeset.

## 5.8 Apply retention settings to safesets in primary and secondary storage

Automatic retention management is disabled in the secondary pool by default.

To expire safesets in primary storage according to retention settings, specify **Online** in the Migrate wizard. Safesets in secondary storage are unaffected. For more information, see [Enforce retention settings in primary storage](#).

To expire safesets in secondary storage according to retention settings, run the following `secondaryop expire` command:

```
secondaryop expire taskid=<t>
```

You can also run the `secondaryop expire` command for particular storage locations:

```
secondaryop expire storage_location=<path>
```

If a path is specified, `secondaryop` scans all secondary pools in that path.

The secondaryop expire is similar to the vvmigrat online command for safesets in primary storage, but applies retention settings across safesets in both primary and secondary storage and only expires safesets from secondary storage. For more information, see [secondaryop](#).

*Warning:* Running secondaryop expire followed by secondaryop optimize can result in the deletion of all safesets in secondary storage. For example, if a task has one backup in primary storage and nine in secondary storage, and the retention settings are Online safesets=1, Online days=0 and Archive days=0, running the secondaryop expire command marks all safesets in secondary storage for deletion.

## 5.9 Secondary storage maintenance

Use the secondaryop command to perform maintenance functions on secondary storage pools. For more information, see [secondaryop](#).

When the secondary storage pool is attached to its originating vault, you must specify the task ID on the command line. To find a task ID, right-click the task in the left pane of the Director UI, and then select **Properties**. The task ID appears in the **Task Name** field.

When the secondary storage pool is maintained by a foreign vault you must specify the physical storage location on the command line.

You can use scheduled entries to automate secondary storage pool maintenance. Some maintenance schedule entries that are provided with Director are disabled by default, and can be enabled.

## 5.10 Close secondary storage pools

If you close a secondary pool for writing, or set a secondary storage location as read only, the affected tasks start with a new secondary pool the next time that secondary migration occurs. This results in a full seed of the data in the new secondary storage location, and could result in a significant increase in the amount of disk space required for secondary storage.

You cannot migrate safesets to a closed secondary pool because it is read-only.

To close a secondary storage pool:

1. From the Storage Management menu, choose Secondary Storage Pools, and then click Close.
2. Complete the Close Secondary Pool for Writing wizard fields.
3. Click **Finish**.

## 5.11 Detach secondary storage pools

Safesets in a detached secondary storage pool are offline and you cannot use them for restores or exports. Detaching a secondary storage pool allows you to remove storage media. When you use a detached secondary storage pool on a different vault it becomes a foreign secondary pool.

If you detach a secondary pool, the safeset status changes to Detached Secondary and the safeset is offline.

A pool system on a detached device is not included in your storage license calculation. However, when you reattach the pool, it is included in your storage license calculation.

To detach a secondary storage pool:

1. From the Storage Management menu, choose Secondary Storage Pools, and then click Detach.
2. Complete the Detach Secondary Pool wizard.
3. Click **Finish**.

## 5.12 Attach secondary storage pools

To access a detached secondary storage pool, you must attach it to its originating vault. When you attach a secondary storage pool, you can use the safesets for restores and exports. You can specify a local secondary pool location, or Universal Naming Convention (UNC) path for the secondary pool location. UNC path names can be a maximum of 128 characters including spaces and slashes (\). You cannot browse to a UNC path; you must enter it manually. To access the secondary pool files of a UNC path, you can use default credentials or provide a valid user name, password and domain name.

If you specify an invalid path or pool, the following message appears: *Configuration file cannot be opened*

To attach a secondary storage pool:

1. From the Storage Management menu, choose Secondary Storage Pools, and then click Attach.
2. Complete the Attach Secondary Pool wizard.
3. Click **Finish**.

## 5.13 Delete storage locations from a secondary storage group

You cannot delete a storage location that has data associated with it. You cannot delete the last location assigned to the last storage group. You must relocate or delete all tasks associated with a location before you can remove it.

To delete a storage location from a secondary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. Click the Secondary Storage Groups tab.
4. Click a secondary storage location.
5. Click Remove Location.
6. Click **Yes**.
7. Click **OK**.

## 5.14 Delete secondary storage groups

To delete a secondary storage group:

1. Select a vault in the left pane of the Director UI.
2. From the Storage Management menu, choose Storage Groups and Locations.
3. Click the Secondary Storage Groups tab.
4. Select a group.
5. Click Delete Group.
6. Click **Yes**.



## 6 Manage archive storage

You can archive safesets offline. Archive storage is used to store data that you access infrequently, but is required to meet the requirements of a corporate retention policy. Archive storage locations can include offline fixed disks, USB drives, and removable media.

Archived data cannot be replicated to another vault.

A safeset's retention policy specifies whether an online safeset is deleted or left online so that it can be archived. Archiving can be enabled when a backup job is run, or enabled in an existing safeset in the vault. See [Enforce retention settings in primary storage](#) and [View and change safeset properties](#).

If you set the archive location to a local disk, the archive job runs as a migration and a migration log is generated. An archive location should provide enough space to store a compressed safeset. Use the Archive command to archive a task and all associated task data to an archive location.

A safeset is not archived if the backup time occurs in the future. To correct this issue, update the timestamp of the safeset on the vault or wait until the backup time for the safeset passes.

### 6.1 Add archive locations

To add an archive location:

1. From the Storage Management menu, choose Storage Groups and Locations.

The Storage Locations dialog box appears.

2. Click the **Archive** tab.
3. Click **Add**.

The Add Location dialog box appears.

4. In the **Path** field, enter or select a UNC or local location for archiving data.
5. Click **OK**.

### 6.2 Assign archive storage locations to tasks

To assign an archive storage location to a task:

1. Right-click a task in the left pane of the Director UI and select **Properties**.
2. Click the **Task Properties** tab.
3. Select an archive storage location in the **Archive area** list.
4. Click **OK**.

## 6.3 Archive safesets from a task

A safeset's retention settings indicate whether the safeset should be archived and stored offline. For more information, see [Enforce retention settings in primary storage](#).

*Note:* Do not move safesets for the same task to both archive storage and secondary storage. Using archive and secondary storage on the same task is not supported, but is not prevented by Director.

An archive log file is created during the archival phase. The log file is stored in the Task Logs folder.

To archive safesets from a task:

1. Select an object in the left pane of the Director UI.
2. From the Storage Management menu, choose Archive Tasks.
3. On the first page of the Archive wizard, click **Next**.
4. On the following pages of the Archive wizard, select the customer, location, computer, and task for which you want to archive safesets. Complete fields in the Archive wizard:
5. On the **When would you like to start the job** page, do one of the following:
  - To run the migration immediately, select **Submit job immediately**, and then click **Next**.
  - To schedule the migration to run, select **Schedule job**, click **Next**, and then specify a schedule for running the migration.
6. Click **Finish**.

## 6.4 Archive safesets to disk

When you select a file system as the archive storage location, the `varchive` command is used to archive data. By default, the `varchive` process is not scheduled. However, you can schedule `varchive` to run automatically when the number of retention days for a safeset is met or exceeded.

Safeset data is transmitted to an archive storage location in Safeset Image (SSI) format. Because each safeset is reconstructed from all of its dependent safesets, this requires a significant amount of storage.

When a task is archived to disk, the computer name and task name are created as directories in order to determine which SSI file belongs to which task. An archived safeset has a different icon and its status is Offline.

All references to pool system data for the archived safeset are removed from the pool index file. An optimization removes all data not referenced by other online safesets, and decreases the pool system size. The only reference remaining is in the safeset list. When a safeset moves to an archive, the physical location path changes to the archive storage location.

## 6.5 Recall safesets

A recall is the retrieval of a safeset from an archive for a restore or export. When you recall a safeset, a copy of the archived safeset is added as a pool file in the main storage area of the task. The original archived safeset is not removed or deleted. The status of the safeset changes from Offline to Recalled.

A retention period is applied to recalled safesets. To expire a safeset that was archived, you must run the `vmigrat offline` command. This command is not scheduled by default.

The Recall log records all changes made to a recalled safeset.

## 6.6 Remove an archive location

To remove an archive location:

1. From the Storage Management menu, choose Storage Groups and Locations.
2. Click the **Archive** tab.
3. Select an archive in the **Location** list.
4. Click **Remove**.
5. In the confirmation message box, click **Yes**.

## 7 Manage customers, locations, accounts and users

To allow an agent to back up data to a vault, a customer, location, account, and user must be created on the vault. Agents use this information to register with the vault.

When an agent is registered with the vault, the vault recognizes the agent and allows it to back up data to and restore data from the vault.

*Note:* You cannot create customers, locations, accounts and users directly on a Satellite vault. When you create customers, locations, accounts and users on a Base vault, they are replicated to Satellite vaults.

### 7.1 Add a customer, location, account and user

Before an agent can back up data to a vault, a customer, location, account, and user must be created on the vault.

You can create a customer, location, account, and user at the same time. You can also create locations, accounts, and users for existing customers. See [Add a location, account and user](#), [Add an account and user](#) and [Add a user](#).

To add a customer, location, account, and user:

1. In the left pane of the Director UI, expand the vault where you want to add a customer, location, account, and user.
2. Right-click Manage Customers/Orgs, Safesets, Tasks, and select Add New Customer.
3. On the Welcome page of the New Organization/Customer wizard, click **Next**.
4. On the General Organization/Customer Information page, type the customer's name and address, and a short name for the customer and click **Next**.

The customer's short name is used to associate their vault data with their site in Portal. Users can then view information from vaults in Portal reports.

5. On the Contact Information page, type the customer's phone number, email address, website, and contact person, and then click **Next**.
6. On the Default Location page, type a default location name and billing code, and then click **Next**.

The billing code (also known as a location code) can be 5-20 characters in length, and can only include alphanumeric characters and dashes. For more information, see [Add a location \(billing\) code](#).

7. On the Account and User Information page, type an account name, user name, and user password, and then click **Next**.

*Notes:*

- The account name must be unique across the entire vault.
- The account and user names cannot include spaces or the following special characters:  
., !, <, >, ' + | \ / ? "
- The maximum password length is 31 characters.

8. If the Base Operating Mode page appears, select the operating mode, and then click **Next**.
9. On the Account Storage Locations page, select storage areas for the account, and then click **Next**.
10. If the Customer Quotas page appears, select each feature (Storage, or a type of agent or plug-in) and click **Set Quota**. In the Organization/Customer Quota dialog box, select **Unlimited** or enter a quota number for the customer in the **Set quota** area.

The Customer Quotas page appears if customer license quotas are used in the vault. For more information, see [Set license quotas for a customer](#).

11. Click **Finish**.

## 7.2 Filter the customer list

Using a filter, you can specify which vault customers appear in the left pane of the Director UI. You can filter the customer list by customer name, or by customer short name. If you specify a customer name and a customer short name, only the customers that match both criteria are displayed. After you apply a filter, the filter parameters appear in angle brackets < > after the vault name.

To filter the customer list:

1. In the left pane of the Director UI, right-click a vault and select **Change filter**.
2. Click Use filter.
3. Do one or both of the following:
  - To filter the list by customer name, enter some or all of a customer name in the **Customer name** field.
  - To filter the list by customer short name, enter some or all of the customer short name in the **Customer short name** field.

Filters are case-sensitive. Spaces entered before or after a name are applied in a filter.

4. Click **OK**.

## 7.3 Add a location, account and user

Before an agent can back up data to a vault, a customer, location, account, and user must be created on the vault.

You can create a location, account, and user for an existing customer at the same time. You can also create accounts and users for existing locations. See [Add an account and user](#) and [Add a user](#).

To add a location, account, and user:

1. In the left pane of the Director UI, expand the vault where you want to add a location, account, and user.
2. Expand Manage Customers/Orgs, Safesets, Tasks.
3. Expand the customer for which you want to add a location.

4. Right-click **Locations** and select **Add New Location**.
5. On the **Welcome** page of the New Location wizard, click **Next**.
6. On the **General Location Information** page, type the location name and address, and then click **Next**.
7. On the **Select a Location Code** page, select a location code or create a new code, and then click **Next**.

The location code (also known as a billing code) can be 5-20 characters in length, and can only include alphanumeric characters and dashes. For more information, see [Add a location \(billing\) code](#).

8. On the **Account Name** page, type an account name, user name, and user password, and then click **Next**.

*Notes:*

- The account name must be unique across the entire vault.
  - The account and user names cannot include spaces or the following special characters:  
.,!,<>'+"|\/?"
  - The maximum password length is 31 characters.
9. On the **Account Storage Locations** page, select storage areas for the account, and then click **Next**.
  10. Click **Finish**.

## 7.4 Add a location (billing) code

When you add a location, you associate the location with a unique alphanumeric code that can be used to calculate storage and usage totals. This location or billing code can be 5-20 characters in length, and can only include alphanumeric characters and dashes.

*Note:* To ensure compatibility with Carbonite Server Backup Reports, the Reports Extractor truncates location codes to five characters.

To add a location (billing) code:

1. From the Reporting menu, choose Location Codes Administration.
2. Click the **Active** tab.
3. Click **New**.
4. Complete the Location Code Creation wizard.
5. Click **Finish**.

## 7.5 Edit a location (billing) code

To edit a location (billing) code:

1. From the Reporting menu, choose Location Codes Administration.
2. Click the **Active** tab.

3. Select a location code in the **Associated Location Codes** pane.
4. Click **Edit**.
5. Edit the location code.
6. Click **Close**.

## 7.6 Delete an unassociated location (billing) code

To delete an unassociated location (billing) code:

1. From the Reporting menu, choose Location Codes Administration.
2. Click the **Delete** tab.
3. Select a location code in the **Unassociated Location Codes** pane.
4. Click **Delete**.
5. Click **OK**.

## 7.7 Add an account and user

Before an agent can back up data to a vault, a customer, location, account, and user must be created on the vault.

You can create an account and user for an existing customer and location in a vault. You can also create a user for an existing account. See [Add a user](#).

*Note:* Each account name must be unique across the entire vault.

If you add an account or a user that already exists, a message states that the account or user could not be added. In addition, one of the following errors might appear in the Admin Service log:

- Violation of UNIQUE KEY constraint 'account name'
- Cannot insert duplicate key in object dbo.VaultAccount.

If these messages appear, select a different account or user name.

To add an account and user:

1. In the left pane of the Director UI, expand the vault where you want to add an account and user.
2. Expand Manage Customers/Orgs, Safesets, Tasks.
3. Expand the customer where you want to add an account.
4. Expand Locations.
5. Right-click the location and select **Add New Account**.
6. On the Welcome page of the New Account wizard, click **Next**.
7. On the Account Name page, type an account name, user name, and user password, and then click **Next**.

*Notes:*

- The account name must be unique across the entire vault.
  - The account and user names cannot include spaces or the following special characters:  
. ! , < > ' + | \ / ? "
  - The maximum password length is 31 characters.
8. On the Account Storage Locations page, select storage areas for the account, and then click **Next**.
  9. Click **Finish**.

## 7.8 Change account properties

To change the properties of an account:

1. In the left pane of the Director UI, expand the vault where you want to change an account's properties.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Right-click an account and select **Properties**.
7. Edit the fields in the **Account Properties** dialog box.

If you change the storage locations for the account, the changes apply to new users. Select **Apply these settings to all the Users in the Account** to apply the changes to existing and future users.

8. Click **OK**.

## 7.9 Add a user

After you create a customer, location, and account, you can add users. Each user name in an account must be unique.

You can also add a user when you create a customer, location, or account. See [Add a location, account and user](#) and [Add an account and user](#).

To add a user:

1. In the left pane of the Director UI, expand the vault where you want to add a user.
2. Expand Manage Customers/Orgs, Safesets, Tasks.
3. Expand the customer.
4. Expand Locations.
5. Expand a location.
6. Right-click an account and select **Add New User**.



7. On the Welcome page of the New User wizard, click **Next**.
8. On the General User Information page, type a user name and password, and then click **Next**.

*Notes:*

- Each user name in an account must be unique.
  - The user name cannot include spaces or the following special characters: . ! , < > ' + | \ / ? "
  - The maximum password length is 31 characters.
9. On the New User Restrictions page, do one of the following, and then click **Next**:
    - To allow vault access for the user from any network address, select **Allow access from any network address**.
    - To allow vault access for the user from a specific IP address, select **Allow access from the following IP address only**, and enter the IP address in the field.
    - To allow vault access for the user from a specific host, select **Allow access from the following host only**, and enter the host name in the field.
  10. On the User Data Locations page, select storage areas for the user, and then click **Next**.
  11. Click **Finish**.


## 7.10 Change user properties

To change the properties of a user:

1. In the left pane of the Director UI, expand the vault where you want to change a user's properties.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Select an account.
7. Right-click a user in the right pane and select **Properties**.
8. Edit the fields in the **User** dialog.
9. Click **OK**.

## 7.11 Disable a user

To temporarily prevent a user from performing a backup or restore, you can disable the user. You can also disable all users in an account.

The following icon appears for a disabled user in the Director UI: 

You can re-enable a disabled user. For more information, see [Enable a user](#).

To disable a user:

1. In the left pane of the Director UI, expand the vault where you want to disable a user.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Select an account.
7. Right-click a user in the right pane and select **Disable**.

To disable all users in an account:

1. In the left pane of the Director UI, expand the vault where you want to disable users.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Right-click an account and select **Disable Users**.

## 7.12 Enable a user

You can re-enable a user that has been disabled, or enable all users in an account.

You can also enable all users and tasks for a customer or location. For more information, see [Enable a disabled or suspect task](#).

To enable a user:

1. In the left pane of the Director UI, expand the vault with the user that you want to enable.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Select an account.
7. Right-click a user in the right pane and select **Enable**.

To enable all users in an account:

1. In the left pane of the Director UI, expand the vault with the users that you want to enable.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.

4. Expand Locations.
5. Expand a location.
6. Right-click an account and select **Enable Users**.

## 7.13 Delete a customer

*Note:* You cannot restore a customer that has been deleted. To temporarily disable tasks and users for a customer rather than deleting the customer permanently, see [Disable a task](#).

To delete a customer:

1. In the left pane of the Director UI, expand the vault where you want to delete a customer.
2. Expand Manage Customers/Orgs, Safesets, Tasks.
3. Right-click a customer and select **Delete**.
4. Enter **YES** and click **OK**.
5. Click **OK** again.

## 7.14 Delete a location

Deleting a location deletes all computers that are assigned to the location and removes all accounts and users.

*Caution:* You cannot restore a location that has been deleted. To temporarily disable all tasks and users for a location rather than permanently deleting the location, see [Disable a task](#).

To delete a location:

1. In the left pane of the Director UI, expand the vault where you want to delete a location.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Right-click the location and select **Delete**.
6. Enter **YES** and click **OK**.
7. Click **OK** again.

## 7.15 Delete an account

*Note:* You can only delete an account if no computers are associated with the account users.

*Note:* You cannot restore a deleted account. To temporarily disable users in an account rather than deleting the account permanently, see [Disable a task](#).

To delete an account:

1. In the left pane of the Director UI, expand the vault where you want to delete an account.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Right-click an account and select **Delete**.
7. Click **Yes**.

## 7.16 Delete a user

*Note:* You can only delete a user that is not associated with a computer.

*Note:* You cannot restore a user that has been deleted. To temporarily disable a user rather than deleting it permanently, see [Disable a task](#).

To delete a user:

1. In the left pane of the Director UI, expand the vault with the user that you want to delete.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Select an account.
7. Right-click a user in the right pane and select **Delete**.
8. Click **Yes**.

## 8 Manage computers, tasks and safesets

In the Director UI, you can view computers with agents that are registered to a vault. You can also view each agent's tasks and safesets.

For most agents, each task in a vault corresponds to a backup job and includes all safesets created by running the job. For a Hyper-V agent, each VM in a backup job is backed up as a separate task.

A safeset is sent to a vault each time a backup is performed successfully. The safeset contains the data protected in that backup session, as well as retention settings and other safeset information.

### 8.1 View registered computers and tasks

In the Director UI, you can view information about computers with agents that are registered to vaults, including their IP addresses, Agent versions, and unique identifiers.

You can also view task information. Each task corresponds to a backup job on an agent, and includes all safesets created by running the job.

To view a registered computer and tasks:

1. In the left pane of the Director UI, expand the vault where you want to view registered computers.
2. Expand the **Manage Customers** list.
3. Expand a customer list.
4. Expand the **Locations** list.
5. Expand a location.
6. Expand the Registered Computers list.
7. To view information about a computer, right-click the computer and choose **Properties** from the menu.

The Computer dialog box shows information about the computer, including its name, domain, IP address, operating system and Agent version.

8. To view information about a task, expand the computer with the task, and then expand **Backup Tasks**. In the Backup Tasks list, right click the task and choose **Properties** from the menu.

The Task dialog box shows information about the backup task, including its pool size and storage groups.

### 8.2 Change a registered computer's Agent type

You can change the Agent type of a registered computer. This may be useful if the vault cannot identify the Agent type during the registration process.

*Note:* The Agent type is considered for licensing purposes.

To change a registered computer's Agent type:

1. In the left pane of the Director UI, expand the vault where you want to change an Agent type.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Expand a **Location**.
6. Expand the Registered Computers list.
7. Right-click a registered computer and select **Properties**.
8. Click Change Agent Type.
9. In the Agent Type dialog box, select an Agent type in the **New Agent Type** list.
10. Click **OK**.

### 8.3 View safesets for a task

You can view safesets for each task in a vault using the Director UI. A safeset is saved in a vault each time a backup is successfully completed. The safeset contains the data protected in that backup session, as well as retention settings and other safeset information.

To view safesets for a task:

1. In the left pane of the Director UI, expand a vault.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Expand a **Location**.
6. Expand the Registered Computers list.
7. Select and expand a computer list.
8. Expand Backup Tasks.
9. Expand a backup task.
10. Click **Backups**.

Safesets for the location appear in the right pane. The following information appears for each safeset:

Column	Description
Backup #	System-generated backup number, starting with 0000001. Successive backups are numbered sequentially. When a backup expires, the retention manager (vmmigrat) deletes it and optimization recovers the disk space as part of regular maintenance.

Column	Description
Backup time	Date and time that the agent started the backup.
Retention #	Retention group to which the safeset belongs. For more information, see <a href="#">Enforce retention settings in primary storage</a> .
Status	Status of the safeset. Each safeset can have one of the following statuses: <ul style="list-style-type: none"> <li>• Online — the safeset is online and can be restored.</li> <li>• Offline — the safeset is archived.</li> <li>• Recalled — the archived safeset can be restored.</li> <li>• Secondary — the safeset has been moved out of primary storage to a secondary storage area.</li> </ul>
Original size	Amount of uncompressed data protected by the safeset. If you fully restore the safeset to a system, this amount of data will be restored.  This value corresponds to the “all stream bytes processed” value in the agent’s backup log.  <i>Note:</i> This value is similar to the “original stream bytes” value in the Director backup log. However, the value in the Director backup log includes additional data required for backups and restores (e.g., blocks used by backup job metadata/catalog files).
Compressed size	Amount of protected data that was new or changed since the last backup, after the data was compressed.  This value corresponds to the “compressed bytes processed” value in the agent’s backup log.
Storage size	Amount of protected data that was new or changed since the last backup, before the data was compressed. For a seed backup, where all data is new, a safeset’s Storage size is the same as its Original size.  This value corresponds to the “deltized bytes processed” value in the agent’s backup log.
Encrypted	Specifies whether the backup data is encrypted.  <i>Note:</i> The password for encrypted data is known only to the user who encrypts the data. The passwords are not stored. If the encryption password is lost, the backup data is inaccessible. If you do not use an encryption password, you can use over-the-wire encryption to secure data going across the network. Over-the-wire encryption does not require any user supplied passwords or encryption methods.

## 8.4 View and change safeset properties

You can view properties for a safeset, including identification numbers and storage location information. You can also change some safeset properties, including retention settings and backup date and time.

To view and change safeset properties:

1. Navigate to the safesets pane. See [View safesets for a task](#).
2. Double-click a safeset in the right pane.

The Safeset Properties dialog box appears.

3. View and edit some or all of the following safeset properties:

Tab	Field	Description
General	Available for restore	If selected, the safeset is online, online-secondary or recalled, and can be restored.
	Backup number	Sequential safeset number.
	Previous backup number	Sequential safeset number for the previous safeset. The number 0 (zero) indicates that the safeset is a seed.
	Serial number	System-generated number for the safeset, used for synchronizing.
Retention	Retention group	Number from 0 to 9 that identifies the Retention Group to which the safeset belongs. A task can belong to different groups to attain different retention values. For more information, see <a href="#">Enforce retention settings in primary storage</a> .  You can change this value, but changing retention settings incorrectly can lead to unexpected data deletion.
	Online safesets	Number of safesets available for immediate recovery and restore. You can change this value.
	Online days	Number of days safesets are kept online. You can change this value.
	Archive days	Number of days safesets are archived. You can change this value.
Location	Physical location	System-generated Globally Unique Identifier (GUID) that identifies the location of the safeset in the vault.
	Location code	Available options are Online (primary or secondary storage) and Offline (archive storage).
Statistics	Backup time	Time (adjusted for GMT offset, to match local time) that the backup started. You can modify this field when there is a discrepancy.



## 8.5 Change the retention settings of multiple safesets

To change the retention settings of multiple safesets:


1. Navigate to the safesets pane. See [View safesets for a task](#).
2. Press **CTRL** and select multiple safesets in the right pane.
3. Right-click a selected safeset and select **Properties**.
4. In the **Safeset Properties** dialog box, edit retention settings for the safesets:


Field	Description
Retention group	Number from 0 to 9 that identifies the Retention Group to which the safeset belongs. A task can belong to different groups to attain different retention values. You can change this value. For more information, see <a href="#">Enforce retention settings in primary storage</a> .
Online safesets	Number of safesets available for immediate recovery and restore. You can change this value.
Online days	Number of days safesets are kept online. You can change this value.
Archive days	Number of days safesets are archived. You can change this value.


5. Click **OK**.

## 8.6 Enabled, disabled and suspect tasks

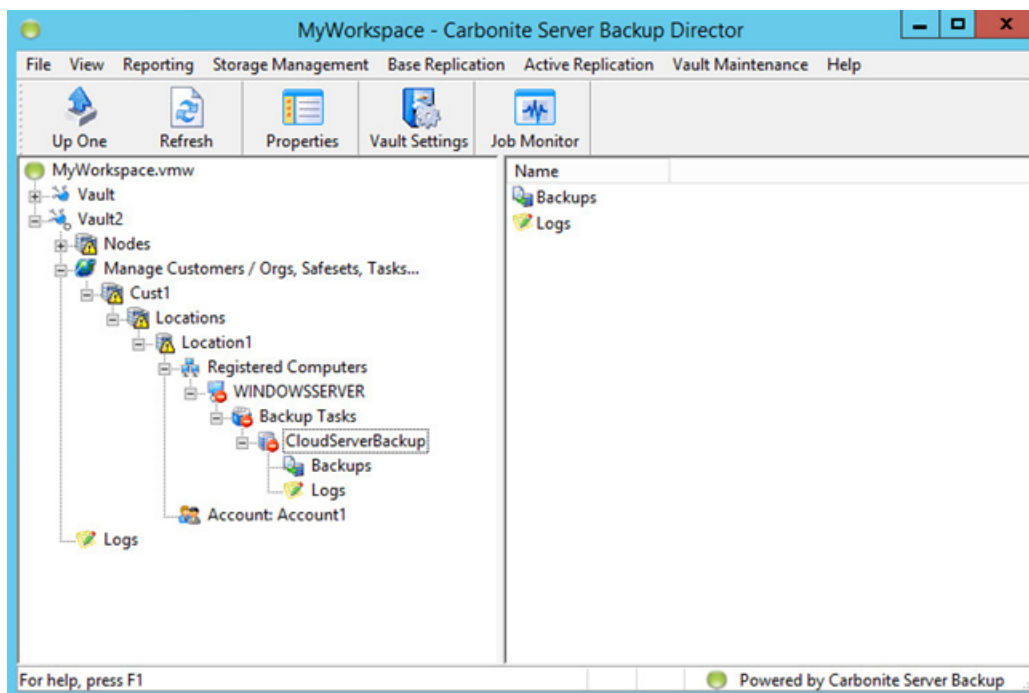
Tasks in the Director can be enabled, disabled or suspect.

Enabled tasks are available for backup, restore, and all other vault processes. The following icon appears in the left pane of the Director UI for each enabled task: 

Disabled tasks are not available for backup or restore, and are skipped by maintenance processes such as optimization, migration and reporting. However, you can run tasks such as "migrate delete" and replication operations for disabled tasks. The following icon appears in the left pane of the Director UI for each disabled task: 

When an item in the hierarchy is disabled, a red Disabled icon appears for the item in the Director UI. For example, the following icon appears for a disabled task: 


When some, but not all, items lower in the hierarchy are disabled for an item, a yellow Warning icon appears for the item. For example, in the following screen, the task named "Backup" is disabled and yellow Warning icons appear for the task's computer, location and customer.



You can manually disable a task to exclude it from maintenance operations. This can be useful if there are problems with the task files or data. See [Disable a task](#). Director can also disable tasks if a clone operation finishes successfully and disables the task on the source. This workflow is part of moving a customer from one vault to another. The tasks are disabled so that backups cannot be accepted during the move.

A task is marked as suspect if the vault detects data integrity issues or possible integrity issues as part of updates on the pool system. Suspect tasks are not available for backup, restore, optimization, migration or reporting. The following restrictions also apply to suspect tasks:

- 1:1 and N:1 Replication — Safesets from a suspect task are not replicated. In addition, the last safeset on the source side is not marked as replicated.
- copy— The suspect status on the source is applied at the target location.

The following icon appears in the left pane of the Director UI for each suspect task: 

You cannot manually change a task to suspect. Director marks a task as suspect if it detects possible data integrity issues. You can only change a task from suspect to enabled or disabled after you identify and correct any data errors in the task pool. If a task is identified as suspect on an Active vault, this status is not replicated on the Passive vault. For more information, see [Enable a disabled or suspect task](#).

## 8.7 Enable a disabled or suspect task

You can re-enable one disabled or suspect task, or all tasks and users for a customer or location.

You can only enable a suspect task if task pool data problems have been resolved. When you try to enable a suspect task, Director checks the physical and logical integrity of the task data and only enables the task if there are no problems with the data. For more information, see [Enabled, disabled and suspect tasks](#).

To enable a disabled or suspect task:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a customer list.
4. Expand Locations.
5. Expand a location.
6. Expand Registered Computers.
7. Expand a computer.
8. Expand Backup Tasks.
9. Right-click a disabled or suspect task and select **Enable**.
10. (If applicable) If you are trying to enable a suspect task, a warning message states that the task pool could be corrupted. Click OK. Director checks the integrity of the task pool data. If there are no data issues, Director enables the task.

To enable all tasks and users for a customer:

1. Select a workspace in the left pane of the Director UI.
2. Expand a vault.
3. Expand the **Manage Customers** list.
4. Right-click a customer and select **Enable all Users and Tasks**.

If a task you are trying to enable is suspect, Director checks the integrity of the suspect task data. If there are no data issues, Director enables the task.

To enable all users and tasks for a location:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Right-click a location and select **Enable all Users and Tasks**.

If a task you are trying to enable is suspect, Director checks the integrity of the suspect task data. If there are no data issues, Director enables the task. This process can take a long time for large tasks.

## 8.8 Disable a task

You can temporarily disable one task, or disable all tasks and users for a customer or location. Disabling a task can be useful if there are problems with the task, such as corrupted files or data. A disabled task is not

available for backup or restore, and is skipped by vault processes such as optimization, migration, reports and storage extractor. For more information, see [Enabled, disabled and suspect tasks](#).

If you disable a task on an Active vault, the task is disabled on the Passive vault.

You can re-enable a disabled task. For more information, see [Enable a disabled or suspect task](#).

To disable a task:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a Customer list.
4. Expand Locations.
5. Expand a location.
6. Expand Registered Computers.
7. Expand a computer.
8. Expand Backup Tasks.
9. Right-click a task and select **Disable**.

To disable all tasks and users for a customer:

1. Select a workspace in the left pane of the Director UI.
2. Expand a vault.
3. Expand the **Manage Customers** list.
4. Right-click a customer and select **Disable All Users and Tasks**.

To disable all users and tasks for a location:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Right-click a location and select **Disable all Users and Tasks**.

## 8.9 Export safesets

You can export safesets from primary or secondary storage as Safeset Image (SSI) files. Agents can restore data from SSI files. You can also import SSI files into a vault. See [Import safesets](#).

You can export safesets to a local disk or a Universal Naming Convention (UNC) device.

An agent requires a corresponding catalog (CAT) file to restore from SSI files. A CAT file provides a visual representation of the directories and files in the backup. You can export the CAT file with SSI files, and add the CAT file to the correct Agent directory to complete the restore.

You can also use the `vvexport` command to export safesets. You can then save the safesets to removable media. For more information, see the [Command Reference](#).

When you export safesets, the original pool system data is unchanged and is available for backups and restores.

To export safesets:

1. Open a command prompt and run the `vvexport` command with the `/simulate` option to determine the size of the backup.
2. If you are using removable media for your backup, connect it to the vault.
3. From the **Storage Management** menu, choose **Export**.
4. Complete fields in the Export wizard:

Field	Description
Export Safesets from this Vault	Exports safesets from the vault.
Export Safesets by specifying the physical path of a Secondary Pool	Exports safesets from the secondary storage pool. When you export safesets from a secondary pool, you can use the secondary pool as an input for the Secondary Restore Server (SRS) utility. The primary storage pool is unchanged, and its safesets are not marked as secondary. In addition, a secondary pool is not created, and the <code>secondarystorage</code> table on the vault is unchanged.
Selected Organization/Customer	Specifies a specific customer. This option is available when you select Export Safesets from this vault.
Selected Location	Specifies a specific customer location. This option is available when you select Export Safesets from this vault.
Selected Computer	Specifies a specific computer. This option is available when you select Export Safesets from this vault.
Selected Task	Specifies a specific task. This option is available when you select Export Safesets from this vault.
Path	The path where the secondary storage pool is located. You can browse to a secondary storage pool on a local disk. If the secondary storage pool is remote, you must enter a UNC path. A maximum of 128 characters are allowed for a UNC path. If the default credentials do not allow you to access the remote path, you must supply custom ones for the specific location or username.
Use Default User Credentials	Applies the default username and password to the location you specify.

Field	Description
Create Custom User Credentials	Creates a new username and password for the location you specify.
User Name	A name that identifies the account user.
Password	The password for the account user.
Test Credentials	Tests the UNC path credentials.
Selected Safeset	Specifies the safesets to export.
Export Safesets (to SSI file)	Exports safesets as SSI files to the specified location.
Export Safesets (to Secondary Pool)	Exports safesets to a secondary storage pool.
Save to	The location to export the safesets.
Deduplicate data. Leave this option selected to reduce secondary pool bloat.	Exporting an already deduplicated pool with the Deduplicate data checkbox selected will allow the pool size to remain the same. Unchecking the Deduplicate data checkbox may cause the exported data to dramatically increase in size.
Safeset (SSI) file size	The size in MB of the exported safeset. The default is 1024MB.
Export Catalogs (to disk or UNC)	Exports catalog (CAT) files to disk or a UNC device.
Save catalogs to a separate location	Saves catalog (CAT) files to a location you specify.
Save catalogs for version 3.2 Agents, and later	Saves catalog (CAT) files for a version 3.2 Agent and later.

5. Click **Finish**.

## 8.10 Import safesets

You can import data in Safeset Image (SSI) file format into the pool system in primary storage. SSI files are typically saved on removable media that you can directly attach to the vault.

Do not run backups before an import completes.

To create SSI files, you can manually back up data on an agent system and choose removable media that is directly attached to the system as the safeset destination. You can also export data from a vault in SSI format. See [Export safesets](#).

To import a safeset:

1. Connect the media containing the safeset to the vault.
2. From the **Storage Management** menu, choose **Import**.
3. On the first page of the Import wizard, click **Next**.
4. On the Organizations/Customer List page, click the customer for the imported safeset, and then click **Next**.
5. On the Locations List page, click the location for the imported safeset, and then click **Next**.
6. On the Computers List page, click the computer for the imported safeset, and then click **Next**.
7. On the Select Task page, click the task for the imported safeset, and then click **Next**.
8. On the Safeset Number page, enter the number of the imported safeset, and then click **Next**.
9. On the Select a Device page, in the **Select Device** field, enter or select the location of the SSI file. If required, specify credentials for connecting to the location. Click **Next**.
10. On the last page of the wizard, click **Finish**.
11. Delete the safeset files on the media.

## 8.11 Delete a registered computer

Deleting a registered computer from a vault removes all information about that computer and its tasks from the Director, and invalidates the vault profile on the agent. Before the computer can connect to the vault again, it must register as a new computer.

To delete a registered computer:

1. In the left pane of the Director UI, expand the vault where you want to delete a registered computer.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand the **Locations** list.
5. Expand a **Location**.
6. Expand the Registered Computers list.
7. Right-click a computer and select **Delete**.
8. In the confirmation message box, click **Yes**.
9. In the Confirm dialog box, enter **YES** and click **OK**.

## 8.12 Delete a task

When you delete a backup task, the entire pool system for the backup is deleted and the task definition on the agent stops working.

*Important:* You cannot recover data from a task that has been deleted from the vault.

To temporarily make a task unavailable for backup or restore, you can disable the task. For more information, see [Disable a task](#).

To delete a task:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Expand Registered Computers.
7. Expand a computer.
8. Expand Backup Tasks.
9. Right-click a task and select **Delete**.
10. Click **Yes**.
11. Click **OK**.

## 8.13 Delete safesets

Deleting a safeset makes it unavailable for restores in the future. Before deleting a safeset, confirm with the safeset owner that you can delete it.

*Caution:* You cannot recover a safeset after you delete it.

When you delete the last safeset, the previous safeset becomes the most current safeset. Because the next safeset is a full, deltized backup, a reseed is not required. However, if you delete the two most recent safesets, a reseed is required for the next backup.

*Note:* In N:1 replication, if you manually delete a safeset from the Base vault, the safeset will be deleted from the Satellite vault when data is synchronized. Similarly, in N:1:1 replication, if you manually delete a safeset from the Active Base vault, the safeset will be deleted from the Satellite vault when data is synchronized.

To delete a safeset:

1. Navigate to the safesets pane. See [View safesets for a task](#).
2. Right-click a safeset in the right pane and select **Delete**.



A confirmation message appears.

3. Click **Yes**.

A message indicates that the Delete command has been submitted to the server.

4. Click **OK**.

To delete multiple safesets:

1. Navigate to the safesets pane. See [View safesets for a task](#).
2. Press **CTRL** and select multiple safesets in the right pane.
3. Right-click a selected safeset and select **Delete**.

A confirmation message appears.

4. Click **Yes**.

A message indicates that the Delete command has been submitted to the server.

5. Click **OK**.

## 9 Manage replication between Satellite and Base vaults

In N:1 replication and in N:1:1 replication, agents send backups to local Satellite vaults, and the data is replicated to a vault in the cloud or in a secondary location in the customer's environment. This ensures that the data is still available for restore if a Satellite vault fails. For more information, see [Many-to-one \(N:1\) replication](#) and [Many-to-one-to-one \(N:1:1\) replication](#).

The same procedures are used to manage replication from Satellite vaults to Base vaults in N:1 replication, and to Active Base vaults in N:1:1 replication. For this reason, the term "Base vault" refers to both Base vaults and Active Base vaults in procedures in this section.

Management tasks are performed on the vault that receives data from a Satellite vault; you cannot perform management tasks directly on a Satellite vault.

When you create customers, locations, accounts and users on a Base vault, they are replicated to Satellite vaults. Customer license quotas limit the amount of storage or number of agents and plug-ins that a specific customer can use.

*Note:* This section describes how to manage replication between Satellite and Base vaults. For information about installing vaults and configuring N:1 and N:1:1 replication, see the *Director Installation Guide*.

**IMPORTANT:** To improve replication security, vault certificates are verified when data is replicated between Director 8.7 vaults. If a vault certificate does not pass verification, replication does not proceed. For more information, see [Certificate verification and pinning for vault-to-vault communications](#).

### 9.1 Schedule replication from Satellite vaults

You can schedule data replication from a Satellite vault to a Base vault or Active Base vault. For example, you can schedule replication during periods of low network activity.

You can also specify whether to replicate data as soon as changes occur, in addition to the scheduled times.

**IMPORTANT:** To ensure that replication operations succeed, be sure that your vault server system clocks are synchronized with a Network Time Protocol (NTP) server.

To schedule replication from a Satellite vault:

1. Select a Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the **Schedule** tab.
5. Do one of the following:
  - To replicate data whenever changes occur as well as according to the schedule, select **Replicate as soon as possible after changes are made and on defined schedule**.
  - To only replicate data when specified by the schedule, select **Replicate on a defined schedule only**.

6. Select the days on which you want replication to occur.
7. In the **Time** field, specify the time at which you want the replication to occur each day.
8. Click **OK**.

## 9.2 Replicate data from Satellite vaults

You can replicate data from a Satellite vault to a Base vault or Active Base vault at any time.

When you run replication, data is not replicated immediately. Instead, replication begins when the Satellite vault next contacts the Base vault or Active Base vault. For more information, see [Set the heartbeat interval for a Satellite vault](#).

To replicate data from a Satellite vault:

1. Select a Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault with data that you want to replicate to the Base vault and click **Replicate Now**.
4. Click **OK**.
5. Click **Close**.

## 9.3 Disable and enable N:1 replication services on a Base vault

You can disable and enable N:1 replication services on a Base vault or Active Base vault.

N:1 replication services are automatically enabled on Base vaults when you add a Replication Many to One license, and on Satellite vaults when the Satellite vaults are installed. When replication services are disabled, neither metadata nor safesets are replicated from Satellite vaults to the Base vault.

To disable or enable N:1 replication services on a vault:

1. Select a Base vault in the left pane of the Director UI.
2. From the Vault Maintenance menu, choose Vault Settings.
3. Click the **Replication** tab and do one of the following:
  - To disable replication services, clear the **Enable N:1 replication services on 'VaultName'** check box.  
*Note:* The actual Windows service (ReplService) still runs when you disable N:1 replication services.
  - To enable replication services, select the **Enable N:1 replication services on 'VaultName'** check box.
4. Click **OK**.

## 9.4 Disable over-the-wire encryption for replication from earlier Satellite vault versions

Communication between a Director 8.7 Satellite vault and a Director 8.7 Base vault or Active Base vault is secured using SSL/TLS.

Communication between a Director 8.6x or earlier Satellite vault and a Director 8.7 Base vault or Active Base vault is encrypted when over-the-wire encryption is enabled. Over-the-wire encryption is enabled by default but can be disabled.

**IMPORTANT:** We do not recommend disabling over-the-wire encryption.

The over-the-wire encryption setting does not affect replication between Director 8.7 vaults.

To disable over-the-wire encryption for replication from earlier vault versions:

1. Select a Base vault or Active Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the **Network Options** tab.
5. Clear Enable over-the-wire encryption.
6. Click **OK**.

## 9.5 Implement bandwidth throttling for replication from Satellite vaults

You can implement bandwidth throttling on a Satellite vault to limit the amount of bandwidth used to send data from the Satellite to the Base vault.

To implement bandwidth throttling for replication from a Satellite vault:

1. Select a Base vault or Active Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the **Network Options** tab.
5. Select Limit bandwidth usage.
6. Enter a bandwidth limit in the **up to** field, and select **MBits/s** or **KBits/s**.
7. Select the days on which you want to restrict bandwidth use.
8. Click **OK**.
9. Click **Close**.

## 9.6 Set the replication policy for a Satellite vault

To set the replication policy for a Satellite vault:

1. Select a Base vault or Active Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the Safeset Management tab.
5. Select **All safesets** to replicate all safesets.

— or —

Select **Only safesets that will be kept on Satellite vault x or more days** to replicate safesets that are stored on the Satellite vault for a specific number of days.

6. Click **OK**.
7. Click **Close**.

## 9.7 Set the retention policy for a Satellite vault

*Note:* The retention period for safesets on a Base vault is usually longer than the retention period on Satellite vaults.

To set the retention policy for a Satellite vault:

1. Select a Base vault or Active Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the Safeset Management tab.
5. Select **All safesets (use the agent's retention policy)** to use the Agent retention policy for safesets.

— or —

Select **Keep at most** to keep a specific number of safesets for a specific period. This setting does not override the Agent retention policy.

6. Click **OK**.
7. Click **Close**.

## 9.8 Set the operating mode for a Satellite vault

You can set the operating mode for each Satellite vault. This operating mode specifies whether or not backups are accepted on the Satellite vault, and replicated to the Base vault or Active Base vault.

*Note:* The operating mode does not affect restores. Regardless of the operating mode, an agent restores data from the Satellite vault if the selected safeset is available on the Satellite vault, and from a Base vault if

the selected safeset is only available on the Base vault. In N:1:1 replication, an agent restores data from the Satellite vault if the selected safeset is available on the Satellite vault, from the Active Base vault if the selected safeset is not available on the Satellite vault, and from the Passive Base vault if the selected safeset is only available on the Passive Base vault.

If the operating mode for a Satellite vault is “Normal operation”, you can set the operating mode for specific customers, locations, computers and tasks on the vault.

To set the operating mode for a Satellite vault:

1. Select a Base vault or Active Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the **Advanced** tab.
5. Select a satellite operating mode:
  - **Normal operation** — Backups are accepted on the Satellite vault, and replicated to the Base vault.
  - **Don't Replicate** — Backups are accepted on the Satellite vault, but data is not replicated to the Base vault. If there are metadata changes on the Base vault, they are applied to the Satellite vault.
  - **Bypass Satellite** — Backups are sent directly from the agent to the Base vault.
  - **Restore Only** — Backups are not allowed and replication is disabled on the Satellite vault.
  - **Customer Only** — Only customer data is replicated from the Satellite vault to the Base vault. This legacy option is not typically used.

*Note:* The operating mode does not affect restores. Regardless of the operating mode, an agent restores data from the Satellite vault if the selected safeset is available on the Satellite vault, and from a Base vault if the selected safeset is only available on the Base vault. In N:1:1 replication, an agent restores data from the Satellite vault if the selected safeset is available on the Satellite vault, from the Active Base vault if the selected safeset is not available on the Satellite vault, and from the Passive Base vault if the selected safeset is only available on the Passive Base vault.

6. Click **OK**.

## 9.9 Set the operating mode for a customer, location or computer on a Satellite vault

You can change replication settings for specific customers, locations and computers on a Satellite vault. The operating mode specifies whether backups are accepted on the Satellite vault and replicated to the Base vault or Active Base vault.

To set the operating mode for a customer, location or computer on a Satellite vault:

1. Expand a Base vault or Active Base vault in the left pane of the Director UI until you see the customer, location, or computer for changing the operating mode.

2. Right-click the customer, location, or computer for setting the operating mode, select **Operating Mode**, and select one of the following:
  - **Use Satellite Operating Mode for all tasks** — Uses the satellite operating mode you selected on the Advanced tab of the Satellite Configuration on Base Vault - *BaseVaultName* dialog box.
  - **Pause Safeset Replication for all tasks** — Temporarily pauses safeset replication for all tasks.  
*Note:* The “Pause Safeset Replication for all tasks” option is only available if the Satellite vault’s operating mode is “Normal operation”. The operating mode for a task is automatically set to “Use Satellite Operating Mode” if the Satellite vault’s operating mode is anything but “Normal Operation”. For more information, see [Set the operating mode for a Satellite vault](#).
3. In the confirmation message box, click **OK**.

## 9.10 Set the operating mode for a task on a Satellite vault

You can change the operating mode for a single task. The operating mode specifies whether or not backups are accepted on the Satellite vault and replicated to the Base vault or Active Base vault.

To set the operating mode for a task on a Satellite vault:

1. Expand a Base vault or Active Base vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a **Customer** list.
4. Expand Locations.
5. Expand a location.
6. Expand Registered Computers.
7. Expand a computer.
8. Expand Backup Tasks.
9. Right-click the task for changing setting the operating mode, click **Operating Mode**, and select one of the following:
  - **Use Satellite Operating Mode** — Uses the satellite operating mode you selected on the Advanced tab of the Satellite Configuration on Base Vault - *BaseVaultName* dialog box.
  - **Pause Safeset Replication** — Temporarily pauses safeset replication for the task.  
*Note:* The “Pause Safeset Replication” option is only available if the Satellite vault’s operating mode is “Normal operation”. The operating mode for a task is automatically set to “Use Satellite Operating Mode” if the Satellite vault’s operating mode is anything but “Normal Operation”. For more information, see [Set the operating mode for a Satellite vault](#).
10. If a confirmation message appears, click **OK**.

## 9.11 Set the heartbeat interval for a Satellite vault

You can set the heartbeat interval for a Satellite vault. The heartbeat interval specifies how frequently the Satellite vault contacts the Base vault or Active Base vault.

*Note:* The term “Base vault” in the remainder of this section refers to either a Base vault or Active Base vault.

During a heartbeat contact, the last heartbeat time is updated on the Base vault, and configuration changes and requests (e.g., replication report requests) from the Base vault are submitted to the Satellite vault.

The Satellite vault gets its licensing from the Base vault. If there is no heartbeat for 30 days, the Satellite vault will act like it has a trial license and expire at the end of the 30 days.

To set the heartbeat interval for a Satellite vault:

1. Select a Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the **Advanced** tab.
5. Enter the number of minutes after which the Satellite vault should contact the Base vault in the **minutes** field.
6. Click **OK**.
7. Click **Close**.

## 9.12 Assign settings control to a Satellite vault

*Note:* The setting described in this procedure was used for a product that is no longer supported: Carbonite Server Backup for Microsoft System Center Data Protection Manager (EDPM)

To assign settings control to a Satellite vault:

1. Select a Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select a Satellite vault and click **Edit**.
4. Click the **Advanced** tab.
5. Select **Allow this Satellite to control its own settings**.
6. Click **OK**.
7. Click **Close**.



## 9.13 View replication activity between Satellite and Base vaults

You can view replication activities between Satellite vaults and Base vaults or Active Base vaults in the Job Monitor. When a replication session is in progress, you can view detailed progress and status information for the task in the Progress Information dialog box.


Each replication process name in the Job Monitor begins with the vault's replication role in the activity: Base or Satellite. Each session name also includes the specific process, such as the type of data being replicated: Configuration (i.e., replication configuration changes), Metadata (e.g., location, computer, account, user and task changes), or Task (i.e., safesets).

Some replication process names also specify the customer, computer or task involved. For example, when database and configuration information is replicated from a Satellite vault to a Base vault, a "Satellite Replication – Metadata – Cust1/Location1/Computer1" message appears in the Job Monitor for the Satellite vault.

When a replication task is in progress, the Job Monitor can also show the process that is currently running. For example, when task data is being replicated from a Satellite vault to a Base vault, a "Base Replication – Task – Cust1/Location1/Computer1/Task (Receiving Data)" message appears in the Job Monitor for the Passive Base vault.

Because multiple replication activities can run at the same time, if a large amount of safeset data is being replicated for one task, Configuration, Metadata and Task replication sessions for other tasks can start and finish while the large task is still running. Configuration replication sessions have priority to run over other sessions, and Metadata replication sessions have priority over Task replication.

To view replication progress between a Satellite and Base vault:

1. In the left pane of the Director UI, select a vault for viewing replication activity.
2. Click Job Monitor: 
3. To show replication activity, do one of the following:
  - If the selected vault is a Base vault, select the **Show Base Replication** option.
  - If the selected vault is a Satellite vault, select the **Show Satellite Replication** option.

The Job Monitor shows the vault's replication activities.

4. (Optional) To view detailed progress information for a replication activity that is running, select the activity and then click **Progress**.

The Progress dialog box shows detailed progress and status information for the replication activity.

5. (Optional) To view replication activity for a different vault, choose the vault from the **Select Vault** list.


## 9.14 Stop Base vault replication processes and services

Using the Job Monitor, you can stop N:1 replication processes and services that are running on a Base vault or Active Base vault.

Each replication process and service listed in the Job Monitor begins with the vault's replication role in the activity: Base or Satellite. For example, when a Base vault receives data from a Satellite vault, the Job Monitor process name begins with "Base".

When you stop a replication service and later want to re-enable it, you must re-enable replication using the Base vault settings.

To stop a Base vault replication process or service:

1. In the left pane of the Director UI, select a Base vault for stopping a replication process or service.
2. Click Job Monitor: 
3. Select the Show Base Replication option.
4. In the list of replication activities, select the "Base" process or service that you want to stop.
5. Click **Stop**.
6. In the confirmation dialog box, click **Yes**.

## 9.15 Run Satellite to Base vault replication reports

To compare data on a Satellite vault and Base vault, or a Satellite vault and Active Base vault, you can run the following reports:

Report	Description
Replication Comparison	Compares objects, including customers, locations, computers and tasks, on Satellite vaults and the Base vault, and indicates whether there are differences. You can specify which objects to compare on each vault, and whether to compare indexes and safesets as well as metadata.
Replication Status	Shows information about each Satellite vault that is registered on the Base vault.

Report	Description
Replication Lag	<p>Shows safesets that have not been replicated from a Satellite vault to the Base vault in a specified number of days. To specify which safesets to include in the report, you can select vaults, customers, locations, computers, or specific tasks.</p> <p>When you run the Replication Lag report, you must enter the number of days after a safeset is created that it should be replicated to the Base vault. For example, if you want a safeset to be replicated to the Base vault two days after a backup, enter 2 as the parameter value.</p> <p>When you run the report, the system compares the date and time of the last safeset created for a task to the date and time of the last replicated safeset for the task. If the time difference is greater than the parameter value, the report lists all online safesets that have not been replicated. An asterisk (*) appears beside any safeset where replication has lagged more than the number of days specified by the parameter value.</p> <p>If the time difference between the last safeset and the last replicated safeset is less than the parameter value, the report does not list any safesets.</p> <p>The report also indicates the percentage of metadata (e.g., catalog files, activity files, index files) that has been replicated to the Base vault for vaults, customers, locations, computers, tasks and safesets in the report.</p> <p><b>IMPORTANT:</b> When you run the Replication Lag report from the UI, the report does not include jobs where replication is paused. To include jobs where replication is paused, run the report from the command line. See <a href="#">replvault</a>.</p>

Reports are saved in the Replication log directory.

You can also run a Replication report from the command line. For more information, see [replvault](#).

To run a Satellite and Base vault replication report:

1. Do one of the following:
  - To run a Replication Comparison or Replication Lag report for one Satellite vault and a Base vault, select the Satellite vault in the left pane of the Director UI. Click **Satellite Replication**, and select **Status Report**.
  - To run a Replication Comparison, Status, or Replication Lag report for one or more Satellite vaults and a Base vault, select the Base vault in the left pane of the Director UI. Click **Base Replication** and select **Status Report**.
2. Complete the fields in the Report wizard. On the **Objects To Include in the Report** page, select the report that you want to run. On the **Report Scope Selection** page, select the data and objects to include.
3. On the last page of the wizard, click **Finish**.

## 9.16 Replace a failed Satellite vault

If a Satellite vault fails, you can generate a new authorization key on the Base vault, and install and register a new Satellite vault to replace the failed vault.

After you replace the vault, customer metadata is transferred from the Base vault to the new Satellite vault so that the new Satellite vault can start receiving backups. If you restore data from a backup that was completed before you replaced the vault, the data will come from the Base vault.

When you run backups after replacing a Satellite vault, existing tasks are reseeded on the Satellite vault. However, existing tasks are not reseeded to the Base vault.

To replace a failed Satellite vault:

1. Select the Base vault in the left pane of the Director UI.
2. Click Base Replication and select Configure Satellites.
3. Select the failed Satellite vault and click **Edit**.
4. Click the **Advanced** tab and select **Bypass Satellite**.
5. Click **OK**.
6. Select the failed Satellite vault and click **Edit**.
7. Click **Reset Key** and record the new authorization key. Click **OK**.
8. Click **Close**.
9. Uninstall the Satellite vault.
10. Install the new Satellite vault. Use the new authorization key and previous IP address. For more information, see the *Director Installation Guide*. Allow replication to finish.
11. Select the Base vault in the left pane of the Director UI.
12. Click Base Replication and select Configure Satellites.
13. Select the Satellite vault and click **Edit**.
14. Click the **Advanced** tab and select **Normal Operation**. Click **OK**.
15. Click **OK** again.
16. Click **Close**.

## 9.17 Set up dual network connections

When you use a single network connection for backup and replication, the Satellite and Base vaults each have one IP address. When you register agents with the Satellite vault, the Satellite vault provides the name of the Base vault. Backups can then be sent to the Base vault if the Satellite vault is unavailable.

*Note:* In this section, “Base vault” refers to a Base vault in N:1 replication, or an Active Base vault in N:1:1 replication. The same information applies in both N:1 and N:1:1 replication.

You can install the Satellite and Base vaults on two separate networks, and use one network connection for backups and one connection for replication. This configuration requires you to install two network cards for the Satellite and Base vaults so that each vault has two IP addresses. With this configuration, backups and replication run faster, and there is no risk of conflicting processes.

Use the **Agent Connection** tab in the **Vault Settings** dialog box to specify the addresses that the agents should use for backups and restores.

## 10 Manage replication between Active and Passive vaults

In 1:1 replication, agents send backups to an Active vault, and the data is replicated to a Passive vault. This ensures that the data is available for restore if the Active vault fails. For more information, see [One-to-one \(1:1\) replication](#).

In N:1:1 replication, data from Satellite vaults is replicated to an Active Base vault, and then to a Passive Base vault, to ensure that the data is always available for restore. For more information, see [Many-to-one \(N:1:1\) replication](#).

The same procedures are used to manage replication between Active and Passive vaults in 1:1 replication, and Active Base vaults and Passive Base vaults in N:1:1 replication. For this reason, the term “Active vault” refers to both Active vaults and Active Base vaults in procedures in this section. The term “Passive vault” refers to both Passive vaults and Passive Base vaults in procedures in this section.

If you disable a task on the Active vault, the task is disabled on the Passive vault. If a task is identified as suspect on the Active vault, this status is not replicated on the Passive vault.

*Note:* This section describes how to manage replication between Active and Passive vaults. For information about installing vaults and configuring 1:1 and N:1:1 replication, see the *Director Installation Guide*.

**IMPORTANT:** To improve replication security, vault certificates are verified when data is replicated between Director 8.7 vaults. If a vault certificate does not pass verification, replication does not proceed. For more information, see [Certificate verification and pinning for vault-to-vault communications](#).

### 10.1 Schedule replication between Active and Passive vaults

You can schedule data replication from an Active vault to a Passive vault, or from an Active Base vault to a Passive Base vault. For example, you can schedule replication during periods of low network activity.

You can also specify whether to replicate data as soon as changes occur, or only according to the defined schedule.

**IMPORTANT:** To ensure that replication operations succeed, be sure that your vault server system clocks are synchronized with a Network Time Protocol (NTP) server.

Safesets are replicated if they are online and not expired. Data stored in archives and secondary storage pools are not replicated.

To schedule replication between an Active and Passive vault:

1. Select an Active vault in the left pane of the Director UI.
2. Click Active Replication and select Configure.
3. Click the **Schedule** tab.
4. Do one of the following:
  - To replicate data whenever changes occur as well as according to the schedule, select **Replicate as soon as possible after changes are made and on defined schedule**.

- To only replicate data when specified by the schedule, select **Replicate on a defined schedule only**.
5. Select the days you want the replication to occur.
  6. Select a time for the replication in the **Time** field.
  7. Click **OK**.

## 10.2 Replicate data from an Active vault to a Passive vault

You can replicate data at any time from an Active vault to a Passive vault, or from an Active Base vault to a Passive Base vault.

To run replication from the command line, see [replvault](#).

To replicate data from an Active vault to a Passive vault:

1. In the left pane of the Director UI, select the Active vault with data that you want to replicate to a Passive vault.
2. Click Active Replication and select Replicate Now.
3. Click **OK**.

## 10.3 Allow or pause replication from Active to Passive vaults for customers, locations or computers

You can allow or pause safeset replication for specific customers, locations and computers from an Active vault to a Passive vault, or from an Active Base vault to a Passive Base vault.

When safeset replication is paused, metadata changes (e.g., customer, location, and computer changes) are replicated to the Passive vault, but safesets are not replicated.

To allow or pause replication from an Active vault to a Passive vault for a customer, location or computer:

1. Expand an Active vault in the left pane of the Director UI until you see the customer, location, or computer for which you want to allow or pause safeset replication.
2. Right-click the customer, location, or computer for allowing or pausing safeset replication, select **Operating Mode**, and select one of the following:
  - **Allow Safeset Replication for all tasks**
  - **Pause Safeset Replication for all tasks**
3. In the confirmation dialog box, click **OK**.

## 10.4 Allow or pause replication from Active to Passive vaults for a task

You can allow or pause safeset replication from an Active vault to a Passive vault, or from an Active Base vault to a Passive Base vault for a single task.

When safeset replication is paused, metadata changes (e.g., task changes) are replicated to the Passive vault, but safesets are not replicated.

To allow or pause replication from an Active to a Passive vault for a task:

1. Expand an Active vault in the left pane of the Director UI until you see the task for which you want to allow or pause safeset replication.
2. Right-click the task for allowing or pausing safeset replication, select **Operating Mode**, and select one of the following:
  - Allow Safeset Replication
  - Pause Safeset Replication
3. If a confirmation dialog box appears, click **OK**.

## 10.5 Disable and enable replication services on an Active or Passive vault

1:1 Replication services are automatically enabled on an Active or Passive vault, or an Active Base vault or Passive Base vault when you add a Replication One to One license. However, you can disable and re-enable the replication services.

When replication services are disabled, neither metadata nor safesets are replicated from the Active vault to the Passive vault.

To disable or enable replication services on an Active or Passive vault:

1. Select an Active or Passive vault in the left pane of the Director UI.
2. From the Vault Maintenance menu, choose Vault Settings.
3. Click the **Replication** tab, and do one of the following:
  - To disable replication services, clear the **Enable 1:1 replication services on 'vaultName'** check box.
  - To enable replication services, select (check) the **Enable 1:1 replication services on 'vaultName'** check box.
4. Click **OK**.



## 10.6 Disable over-the-wire encryption for replication from earlier Active vault versions

Communication between a Director 8.7 Active or Active Base vault and a Director 8.7 Passive or Passive Base vault is secured using SSL/TLS.

Communication between a Director 8.6x or earlier Active or Active Base vault and a Director 8.7 Passive vault or Passive Base vault is encrypted when over-the-wire encryption is enabled. Over-the-wire encryption is enabled by default, but can be disabled.

The over-the-wire encryption setting does not affect replication between Director 8.7 vaults.

To disable over-the-wire encryption for replication between Active and Passive vaults:

1. In the left pane of the Director UI, select an Active vault.
2. Click Active Replication and select Configure.
3. Click the **Network Options** tab.
4. Clear Enable over-the-wire encryption.
5. Click **OK**.

## 10.7 Implement bandwidth throttling for replication from an Active vault

You can implement bandwidth throttling to limit the amount of bandwidth used to send data from an Active vault to a Passive vault, or from an Active Base vault to a Passive Base vault.

To implement bandwidth throttling for 1:1 replication:

1. Select an Active vault in the left pane of the Director UI.
2. Click Active Replication and select Configure.
3. Click the **Network Options** tab.
4. Select Limit bandwidth usage.
5. Enter a bandwidth limit in the **up to** field and select **MBits/s** or **KBits/s**.
6. Select the days that you want to restrict bandwidth use.
7. Click **OK**.

## 10.8 View replication activity between Active and Passive vaults

You can view replication activities between Active vaults and Passive vaults, or between Active Base vaults and Passive Base vaults in the Job Monitor. When a replication session is in progress, you can view detailed replication progress and status information in the Progress Information dialog box.


Each replication session name in the Job Monitor begins with the vault's replication role in the activity: Active or Passive. Each session name also includes the specific process, such as the type of data being replicated: Configuration (i.e., replication configuration changes), Metadata (e.g., location, computer, account, user and task changes), or Task (i.e., safesets).

Some replication session names include the name of the customer, computer or task involved. For example, when task data is replicated from an Active vault to a Passive vault, a "Passive Replication – Task – Cust1/Location1/Computer1/Backup" session appears in the Job Monitor for the Passive vault.

When a replication task is in progress, the Job Monitor message can also show the process that is currently running. For example, when metadata changes are replicated from an Active vault to a Passive vault, an "Active Replication – Metadata – Vault (Sending Data)" message appears in the Job Monitor for the Active vault.

Because multiple replication activities can run at the same time, if a large amount of safeset data is being replicated for one task, Configuration, Metadata and Task replication sessions for other tasks can start and finish while the large task is still running. Configuration replication sessions have priority to run over other sessions, and Metadata replication sessions have priority over Task replication.

To view replication activity between Active and Passive vaults:

1. In the left pane of the Director UI, select a vault for viewing replication activity.
2. Click Job Monitor: 
3. To show replication activity, do one of the following:
  - If the selected vault is an Active vault, select the **Show Active Replication** option.
  - If the selected vault is a Passive vault, select the **Show Passive Replication** option.

The Job Monitor shows the vault's replication activities.

4. (Optional) To view detailed progress information for a replication activity that is running, click the activity and then click **Progress**.

The Progress dialog box shows detailed progress and status information for the replication activity.

5. (Optional) To view replication activity for a different vault, choose the vault from the **Select Vault** list.

## 10.9 Run Active to Passive vault replication reports

To compare data on an Active vault and Passive vault, or an Active Base vault and Passive Base vault, you can run the following reports:

Report	Description
Replication Comparison	Compares objects, including customers, locations, computers and tasks, on the Active vault and the Passive vault, and indicates whether there are differences. You can specify which objects to compare on each vault, and whether to compare indexes and safesets as well as metadata.

Report	Description
Replication Lag	<p>Shows safesets that have not been replicated from the Active vault to the Passive vault in a specified number of days. To specify which safesets to include in the report, you can select vaults, customers, locations, computers, or specific tasks.</p> <p>When you run the Replication Lag report, you must enter the number of days after a safeset is created that it should be replicated to the Passive vault. For example, if you want a safeset to be replicated to the Passive vault two days after a backup, enter 2 as the parameter value.</p> <p>When you run the report, the system compares the date and time of the last safeset created for a task to the date and time of the last replicated safeset for the task. If the time difference is greater than the parameter value, the report lists all online safesets that have not been replicated. An asterisk (*) appears beside any safeset where replication has lagged more than the number of days specified by the parameter value.</p> <p>If the time difference between the last safeset and the last replicated safeset is less than the parameter value, the report does not list any safesets.</p> <p>The report also indicates the percentage of metadata (e.g., catalog files, activity files, index files) that has been replicated for vaults, customers, locations, computers, tasks and safesets in the report.</p>

Reports are saved to the Replication log directory.

You can also run a replication report from the command line. See [replvault](#).

To run an Active to Passive vault replication report:

1. Select an Active vault in the left pane of the Director UI. Click **Active Replication**, and select **Status Report**.
2. Complete the fields in the Report wizard. On the **Objects To Include in the Report** page, select the report you want to run. On the **Report Scope Selection** page, select the data and objects to compare.
3. On the last page of the Report wizard, click **Finish**.

## 10.10 Fail over from an Active vault to a Passive vault

If an Active vault fails or becomes unavailable for some reason, you can “fail over” to the Passive vault so that it becomes the new Active vault. Agents then send data directly to the new Active vault instead of to the unavailable vault. Similarly, if an Active Base vault is unavailable, you can fail over to the Passive Base vault so that it becomes the new Active Base vault.

### WARNINGS:

- If you fail over to a Passive vault before data is fully synchronized with the Active vault, unreplicated data might be lost. Safesets that are not replicated from the Active vault to the Passive vault before failover are marked for deletion on the formerly Active vault and remain only until they are deleted by maintenance processes.
- Before failing over from an Active Base vault to a Passive Base vault, ensure that the vaults are properly configured and licensed. Failover to a vault that is licensed only as a Passive vault (not as a Passive Base vault) is not supported in N:1:1 replication.
- Do not register an agent directly to a Passive vault. If an agent exists on the Passive vault but not on the Active vault, the agent will be removed from the Passive vault when the vaults are synchronized. To restore data from a Passive vault, fail over to the Passive vault.

When the unavailable vault becomes available again, you must replicate data from the current Active vault to the current Passive vault to synchronize the data. You can then “fail back” and change the current Passive vault into the Active vault again. See [Fail back to a formerly Active vault](#).

*Note:* Vault certificates are verified when a Director 8.7 source vault tries to connect to a Director 8.7 target vault. After a failover in an N:1:1 replication configuration, data will replicate successfully from Satellite vaults to the new Active Base vault. However, if replication from Satellite vaults to the new Active Base vault fails, check the Replication Service log to determine whether there is a certificate issue. For more information, see [vaultop certificate subcommands](#).

To fail over from an Active vault to a Passive vault:

1. Ensure that:
  - The Listener service is not running on the Active vault. To check whether the Listener service is running, see [View, start and stop vault services](#). If the Listener service is running, ramp down the vault and wait for it to go offline so that backups cannot run. See [Ramp down a vault](#). If you fail over from an Active vault to a Passive vault while a backup is running, your backup data could be corrupted.
  - No replication sessions are running. Use the Job Monitor to check whether replication sessions are running. If replication sessions are running, wait for them to finish or stop the replication sessions. See [View and stop vault processes in the Job Monitor](#).  
If you stop a replication session before a safeset is completely replicated, the next replication session begins where the last session ended. Data that has already been replicated does not have to be replicated again.
2. Select the Passive vault in the left pane of the Director UI.

3. Click **Passive Replication** and select **Failover**.
4. Do one of the following:
  - If a dialog box states that there are no pending or running sessions, type “YES”, and click **OK**.  
*Note:* You must type the word “YES” in all capital letters.
  - If a dialog box states that there are pending or running replication events, read the warning and information carefully. If you still want to fail over from the Active to the Passive Base vault, type “I AGREE” in capital letters, and click **Force**.
5. In the confirmation message box, click **OK**.

## 10.11 Fail back to a formerly Active vault

As described in [Fail over from an Active to a Passive vault](#), if an Active vault is unavailable, you can “fail over” to the Passive vault so that it becomes the new Active vault. Similarly, if an Active Base vault is unavailable, you can fail over to the Passive Base vault so that it becomes the new Active Base vault.

When the unavailable vault is available again, you must replicate data from the current Active vault to the current Passive vault to synchronize the data. You can then “fail back” and change the current Passive vault into the Active vault again.

*Warning:* If you fail back to a replacement vault that has the same hostname or IP address as the original vault but a different self-signed certificate, replication will fail from source vaults where the original vault’s certificate is pinned. To resolve this issue, you can import a CA-signed certificate for the replacement vault or delete the pinned certificate from the replication source vaults. For more information, see [vaultop certificate subcommands](#).

*Warning:* If you fail back to a formerly Active vault before data is fully synchronized with the current Active vault, unreplicated data might be lost. Safesets that are not replicated from the current Active vault before failback are marked for deletion on the formerly Active vault and remain only until they are deleted by maintenance processes.

To fail back to a formerly Active vault:

1. Replicate data from the current Active vault to the formerly Active (current Passive) vault to ensure that the data is synchronized. See [Replicate data from an Active vault to a Passive vault](#).
2. Fail back to the previously Active vault. The failback procedure is the same as the initial failover procedure. See [Fail over from an Active vault to a Passive vault](#).

*Note:* If the **Passive Replication** menu does not appear when you select the current Passive vault in the left pane of the Director UI, click **Refresh** to update the Director UI.

## 11 Configure and run reports

You can configure and run the following Director reports using the Report wizard:

- **Storage.** The Storage report shows the number of safesets and amount of data in a vault for a single task, or for each task in a vault, organization/customer, location, or computer. See [Create a Storage report](#).
- **Storage Location.** The Storage Location report lists all primary, secondary and archive storage locations for one task, or for each task in a vault, organization/customer, location, or computer. See [Create a Storage Location report](#).
- **Vault Storage.** The Vault Storage report shows the number of customers, number of servers and the total amount of storage in a vault. The amount of storage includes both primary and secondary storage. See [Create a Vault Storage report](#).
- **Last Backup Status.** The Last Backup Status report shows information about the last backup for one task, or for each task in a vault, organization/customer, location, or computer. See [Create a Last Backup Status report](#).
- **Late Server Status.** The Late Server Status report lists tasks that were not backed up in a specified amount of time before the current time, and shows the last backup date and time for each task. See [Create a Late Server Status report](#).
- **Missed Backups.** The Missed Backups report lists backup tasks that were scheduled or expected to run, but for which a safeset does not exist. See [Create a Missed Backups report](#).
- **Storage Pool Summary.** The Storage Pool Summary report shows the amount of disk space allocated to one task, or to each task in a vault. See [Create a Storage Pool Summary report](#).

You can also run a Storage, Storage Location, or Last Backup Status report quickly from the left pane of the Director UI. For more information, see [Run a Storage, Storage Location or Last Backup Status report](#).

When you configure a report, you can run the report immediately, or schedule it to run on specific days in a week or month at a specific time. For more information, see [Create schedules for running reports](#).

### 11.1 Run a Storage, Storage Location or Last Backup Status report

You can run a Storage, Storage Location, or Last Backup Status report from the left pane of the Director UI. When you run a report from the left pane of the Director UI, it is saved as a log file. Storage and Storage Location reports are created in organized column format, and Last Backup Status reports are created in CSV format.

If you want to create a schedule for running a report, e-mail a report to a recipient, or create a Storage report in CSV format, you must configure the report using the Report wizard. For more information, see [Create a Storage report](#), [Create a Storage Location report](#), and [Create a Last Backup Status report](#).

To run a Storage, Storage Location or Last Backup Status report:

1. In the left pane of the Director UI, right-click the vault, organization/customer, location, computer, or task for the report.
2. From the shortcut menu, choose the type of report to run: Storage Report, Storage Location Report, or Last Backup Status.

## 11.2 Create a Storage report

The Storage report shows the number of safesets and amount of data in a vault for a single task, or for each task in a vault, organization/customer, location, or computer. A report sample appears at the end of the procedure that follows.

For each task, and in totals for each organization/computer, location, customer and vault, the report provides information in the following fields:

Field	Description
Media	Location of the task's safesets. Possible values are: <ul style="list-style-type: none"> <li>• PR — primary storage</li> <li>• SE — secondary storage</li> <li>• OF — offline (i.e., detached) storage</li> <li>• N/A — not applicable. There are no safesets for the task.</li> </ul>
Gen	Number of safesets for the task. The number of safesets is shown separately for each storage location (i.e., primary, secondary and offline).
Original	Total amount of original data protected by the task's safesets. This value is the sum of the safesets' Original size values. The Original size values are provided by the agent. The amount of data is shown separately for each storage location (i.e., primary, secondary and offline).
Compres'd	Total amount of compressed data in the safesets. This value is the sum of the safesets' Compressed size values. The Compressed size values are provided by the agent. The amount of data is shown separately for each storage location (i.e., primary, secondary and offline).
Poolsize	Size of the task's pool. The amount of data is shown separately for each storage location (i.e., primary, secondary and offline).

You can use the following procedure to configure and run or schedule a Storage report. You can also run a Storage report quickly from the left pane of the Director UI. For more information, see [Run a Storage, Storage Location or Last Backup Status report](#).

To create a Storage Report:

1. In the left pane of the Director UI, select the vault for running a Storage History report. From the **Reporting** menu, choose **Storage**.

The Report wizard appears.

2. On the Welcome screen, click **Next**.
3. On the Report Media screen, specify the report destination. For more information, see [Select the destination for a report](#).
4. Select one of the following Output Style Format options:
  - To format the report in columns, select **Organized column format**.
  - To create the report in comma-separated values (CSV) format, select **CSV format (comma separated)**.
5. Click **Next**.
6. On the Select Scope screen and any subsequent selection screens, choose data to include in the report. For more information, see [Select the scope for a report](#).
7. On the Command Execution Time screen, do one of the following:
  - To run the report, click **Submit job immediately**, and then click **Next**.
  - To create a schedule for running the report, click **Schedule job**, and then click **Next**. On the **Select the Execution cycle** screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).
8. On the Report Wizard completion page, review the report information, and then click **Finish**.

### 11.2.1 Storage report sample

```

+-----+
| CARBONITE STORAGE HISTORY REPORT 3-JAN-2023 09:27:20.46 -0500 |
| |
| Server: VAULT Customer: <all> |
| Location: <all> |
| Start Date: <none> Computer: <all> |
| End Date: <none> Task: <all> |
+-----+

```

Customer	Location	Computer	Task	Media	Gen	Original	Compres'd	~Poolsize
CUSTOMER1	LOCA1	COMPUTER1	TASK1	PR	1	1.9MB	0bytes	1.1MB
			TASK2	PR	1	1.9MB	0bytes	878.9KB
			TASK3	PR	7	5.8GB	736.9MB	1.1GB
			TASK4	PR	15	12.4GB	1.5GB	1.5GB
			TASK5	PR	3	3.1GB	315.7MB	1.1GB
		Computer Total		PR	27	21.2GB	2.6GB	3.8GB
	Location Total			PR	27	21.2GB	2.6GB	3.8GB
Customer Total				PR	27	21.2GB	2.6GB	3.8GB
Vault Total				PR	27	21.2GB	2.6GB	3.8GB



## 11.3 Create a Storage Location report

A Storage Location report lists all primary, secondary and archive storage locations for one task, or for each task in a vault, organization/customer, location, or computer. A report sample appears at the end of the procedure that follows.

You can use the following procedure to configure and run or schedule a Storage Location report. You can also run a Storage Location report quickly from the left pane of the Director UI. For more information, see [Run a Storage, Storage Location or Last Backup Status report](#).

To create a Storage Location report:

1. In the left pane of the Director UI, select the vault for running a Storage Location report.
2. From the Reporting menu, choose Storage Locations.

The Report wizard appears.

3. On the Welcome screen, click **Next**.
4. On the Report Media screen, specify the report destination, and then click **Next**. For more information, see [Select the destination for a report](#).
5. On the Select Scope screen and any subsequent selection screens, choose data to include in the report. For more information, see [Select the scope for a report](#).
6. On the Command Execution Time screen, do one of the following:
  - To run the report, click **Submit job immediately**, and then click **Next**.
  - To create a schedule for running the report, click **Schedule job**, and then click **Next**. On the **Select the Execution cycle** screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).
7. On the Report Wizard completion page, review the report information, and then click **Finish**.

### 11.3.1 Storage Location report sample

```
+-----+
| CARBONITE STORAGE LOCATION REPORT 3-JAN-2022 09:30:27.77 -0500 |
| |
| For Server: VAULT |
+-----+
Customer Location Computer Task Storage Locations
-----
CUSTOMER1 LOCATION1 COMPUTER1 TASK1 Primary: C:\Vault\W2008\pool1\
Secondary: C:\SECONDARY\ {SEC}
Archive: C:\ARCHIVE\
TASK2 Primary: C:\Vault\W2008\pool2\
Secondary: C:\SECONDARY\ {SEC}
Archive: C:\ARCHIVE\
TASK3 Primary: C:\Vault\W2008\pool3\
Secondary: C:\SECONDARY\ {SEC}
Archive: C:\ARCHIVE\
TASK4 Primary: C:\Vault\W2008\pool4\
```

```
Secondary: C:\SECONDARY\ {SEC}
Archive: C:\ARCHIVE\
TASK5 Primary: C:\Vault\W2008\pool5\
Secondary: C:\SECONDARY\ {SEC}
Archive: C:\ARCHIVE\
```

## 11.4 Create a Vault Storage report

The Vault Storage report shows the number of customers, number of servers and the total amount of storage in a vault. A report sample appears at the end of the procedure that follows.

The amount of storage includes both primary and secondary storage. Secondary storage is included in the total amount of storage regardless of whether it is attached or not.

To create a Vault Storage report:

1. In the left pane of the Director UI, select the vault for running a Vault Storage Report.
2. From the **Reporting** menu, choose **Vault Storage**.

The Report wizard appears.

3. On the Welcome screen, click **Next**.
4. On the Report Media screen, in the **Send generated report by email to** field, enter one or more email addresses where you want to send the report. Email addresses must be separated by commas.

*Note:* The Vault Storage report cannot be saved as a log file. Before you can run the Vault Storage report, an SMTP server must be specified for the vault. For more information, see [Configure email notifications](#).

5. Click **Next**.
6. On the Command Execution Time screen, do one of the following:
  - To run the report, click **Submit job immediately**, and then click **Next**.
  - To create a schedule for running the report, click **Schedule job**, and then click **Next**. On the **Select the Execution cycle** screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).
7. On the Report Wizard completion page, review the report information, and then click **Finish**.

### 11.4.1 Vault Storage report sample

```
+-----+
| VAULT STORAGE REPORT 3-JAN-2022 09:32:12.30 -0500 |
| |
| Vault Name: Vault1 Vault Domain: |
| Vault ID: 2136085941 |
+-----+
Total Customers      Total Servers      Total Storage
-----
1                    1                    3.81 GB
```

## 11.5 Create a Last Backup Status report

The Last Backup Status report shows information about the last backup for one task, or for each task in a vault, organization/customer, location, or computer. The report is available only in comma-separated values (CSV) format. A report sample appears at the end of the procedure that follows.

For each task, the report includes information in the following fields:

Field	Description
LastSynch#	Highest safeset number for the task.
Date	Date of the last safeset for the task
Original	Amount of original data protected by the safeset. The Original size value is provided by the agent.
Compres'd	Amount of compressed data in the safeset. The Compressed size value is provided by the agent.
Delta	Difference in amount of data protected by the last safeset and the previous safeset.

You can use the following procedure to configure and run or schedule a Last Backup Status report. You can also run a Last Backup Status report quickly from the left pane of the Director UI. For more information, see [Run a Storage, Storage Location or Last Backup Status report](#).

To create a Last Backup Status report:

1. In the left pane of the Director UI, select the vault for running a Last Backup Status report.
2. From the Reporting menu, choose Last Backup Status.  
The Report wizard appears.
3. On the Welcome screen, click **Next**.
4. On the Report Media screen, specify the report destination, and then click **Next**. For more information, see [Select the destination for a report](#).
5. On the Select Scope screen and any subsequent selection screens, choose data to include in the report. For more information, see [Select the scope for a report](#).
6. On the Command Execution Time screen, do one of the following:
  - To run the report, click **Submit job immediately**, and then click **Next**.
  - To create a schedule for running the report, click **Schedule job**, and then click **Next**. On the **Select the Execution cycle** screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).
7. On the Report Wizard completion page, review the report information, and then click **Finish**.

### 11.5.1 Last Backup Status report sample

```
"Customer","Location","Computer","Task","LastSynch#","Date","Original","Compres'd","Delta"
"CUSTOMER1","LOCATION1","COMPUTER1","TASK1","2054","18-OCT-2021 21:20","1,997,776","0","0"
"CUSTOMER1","LOCATION1","COMPUTER1","TASK2","2068","03-NOV-2021 21:20","1,997,776","0","0"
"CUSTOMER1","LOCATION1","COMPUTER1","TASK3","11","27-DEC-2021
14:25","1,217,595,364","110,195,704","110,690,444"
"CUSTOMER1","LOCATION1","COMPUTER1","TASK4","15","27-DEC-2021
14:47","1,660,357,140","110,563,064","110,690,444"
"CUSTOMER1","LOCATION1","COMPUTER1","TASK5","11","27-DEC-2021
12:46","1,217,595,364","110,236,504","110,690,444"
```

## 11.6 Create a Late Server Status report

A Late Server Status report lists tasks that do not have safesets from the specified amount of time before the current time, and shows the last backup date and time for each task. A report sample appears at the end of the procedure that follows.

To create a Late Server Status report:

1. In the left pane of the Director UI, select the vault for running a Late Server Status report.
2. From the Reporting menu, choose Late Server Status.

The Report wizard appears.

3. On the Welcome screen, click **Next**.
4. On the Report Media screen, specify the report destination, and then click **Next**. For more information, see [Select the destination for a report](#).
5. On the Report Settings screen, in the **Report all tasks that have not backed up in the past x hours** field, enter the number of hours before the current time for including tasks.

Tasks appear in the report if they have not been backed up in the specified amount of time before the current time. By default, the report lists tasks that have not been backed up in the past 48 hours.

6. On the Select Scope screen and any subsequent selection screens, choose data to include in the report. For more information, see [Select the scope for a report](#).
7. On the Command Execution Time screen, do one of the following:
  - To run the report, click **Submit job immediately**, and then click **Next**.
  - To create a schedule for running the report, click **Schedule job**, and then click **Next**. On the **Select the Execution cycle** screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).
8. On the Report Wizard completion page, review the report information, and then click **Finish**.

## 11.6.1 Late Server Status report sample

Late Server Status Report

Report Scope: Entire Vault;

Report Time Frame: 1-JAN-2022 09:34:04.05 -0500 To 3-JAN-2022 09:34:04.05 -0500

The following tasks have not been backed up in the past 48 hours

```
"Customer","CustomerEmail","Location","LocationEmail","Computer","Task","LastBackupDate"
"CUSTOMER1","customer1@address.com","LOCATION1","customer1@address.com","COMPUTER1","TASK
1","18-OCT-2021 21:20:09.70 -0500"
"CUSTOMER1","customer1@address.com","LOCATION1","customer1@address.com","COMPUTER1","TASK
2"," 3-NOV-2021 21:20:26.91 -0500"
"CUSTOMER1","customer1@address.com","LOCATION1","customer1@address.com","COMPUTER1","TASK
3","27-DEC-2021 14:25:20.69 -0500"
"CUSTOMER1","customer1@address.com","LOCATION1","customer1@address.com","COMPUTER1","TASK
4","27-DEC-2021 14:47:48.33 -0500"
"CUSTOMER1","customer1@address.com","LOCATION1","customer1@address.com","COMPUTER1","TASK
5","27-DEC-2021 12:46:57.63 -0500"
```

## 11.7 Create a Missed Backups report

The Missed Backups report lists backup tasks that were scheduled or expected to run, but for which a safeset does not exist. A task appears in the report even if the scheduled or expected backup ran successfully but the committed safeset was later deleted. The report lists tasks with missed backups, and shows the expected backup time and detected backup time (time of the last existing safeset) of each task. A report sample appears at the end of the procedure that follows.

A backup task appears in the Missed Backups report if one of these is true:

- The task is scheduled, but there is no safeset from the last scheduled backup. A task appears only once in the report even if there are no safesets for multiple scheduled backups.
- The task is not scheduled, a safeset exists for the task from before the report timeframe, and no safeset exists for the task from during the timeframe.

The report timeframe is based on the amount of time entered when you create the Missed Backups report, and is measured back from the current local time on the vault when the report runs. For example, if you specify 48 hours for the report and the report runs on November 21, 2021 at 10:00, the report includes unscheduled backup tasks with safesets dated before November 19, 2021 at 10:00, and no safesets dated between November 19, 2021 at 10:00 and November 21, 2021 at 10:00.

*Note:* The timeframe does not have any effect on tasks that are not scheduled.

For a missed backup that is scheduled, the expected backup time in the report is the last date when the report was scheduled to run, and the time of the last existing safeset (or 00:00 GMT if there are no safesets). For example, if a backup task named Job1 was scheduled to run on November 20, 2021, but the last safeset is from November 16, 2021 at 13:01, the expected backup time is November 20, 2021 at 13:01, as shown in the following example:

```
Computer TaskName ExpectedBackupTime DetectedBackupTime DetectedFailure
```

```
-----
COMPUTER1 JOB1 20-NOV-2021 13:01 16-NOV-2021 13:01 00-000-0000 00:00
```

For a missed backup that is not scheduled, the expected backup time in the report is the start of the timeframe. For example, if the report runs on November 21, 2021 at 10:00 with a timeframe of 48 hours before the current vault time, and the last safeset for a task named Job2 that is not scheduled is dated before the timeframe, the expected backup time is November 19, 2021 at 10:00, as shown in the following example:

```
Computer TaskName ExpectedBackupTime DetectedBackupTime DetectedFailure
-----
```

```
COMPUTER2 JOB2      19-NOV-2021 10:00   16-NOV-2021 13:01   00-000-0000 00:00
```

**Note:** Backup tasks which occur in the future according to the vault time do not appear in the Missed Backups report.

To create a Missed Backups report:

1. In the left pane of the Director UI, select the vault for running a Missed Backups report.
2. From the **Reporting** menu, choose **Missed Backups**.
3. In the left pane of the Missed Backups Report Configuration screen, select the customers, locations, computers, and tasks for the report.
4. In the right pane of the screen, enter an amount of time in the **Show Reports for Missed or Failed Backups Before the Previous x Hours** field.
 

**Note:** The amount of time only affects whether tasks that are not scheduled appear in the report. The time does not affect whether scheduled tasks appear in the Missing Backups report.
5. In the Send Report To pane, specify recipients for the report.
6. (Optional) To format report data as a CSV file, select **Format Report Data in CSV (Comma Separated Values)**.
7. Do one of the following:
  - To run the report immediately, click **Submit Now**.
  - To schedule the report to run weekly or monthly on a specific day at a specific time, click **Schedule**. On the Change Schedule Entry Wizard screen, click **Next**. On the Select the Execution Cycle screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).

### 11.7.1 Missed Backups report sample

```
+-----+
| 22-NOV-2021 11:06:47.33 -0500 |
| |
| Server: SERVER1 Customer: CUSTOMER1 |
| Location: <all> |
| Start Date: <none> Computer: <all> |
| End Date: <none> Task: <all> |
```

+-----+  
+-----+

\*\*\*\*\*

Customer name: CUSTOMER1

\*\*\*\*\*

Computer	TaskName	ExpectedBackupTime	DetectedBackupTime	DetectedFailure
COMPUTER1	TASK1	22-NOV-2021 09:25	16-NOV-2021 09:25	00-000-0000 00:00
COMPUTER1	TASK2	15-NOV-2021 11:51	14-NOV-2021 11:51	00-000-0000 00:00
COMPUTER1	TASK3	22-NOV-2021 09:30	21-NOV-2021 09:30	00-000-0000 00:00
COMPUTER1	TASK4	22-NOV-2021 09:30	21-NOV-2021 09:30	00-000-0000 00:00
COMPUTER1	TASK5	22-NOV-2021 09:30	21-NOV-2021 09:30	00-000-0000 00:00

## 11.8 Create a Storage Pool Summary report

The Storage Pool Summary report provides information about the amount of disk space allocated to one task, or to each task in a vault. For each task, the report shows the total physical size of the pool, the total amount of actual data in the pool, and an estimated amount of storage space that could be reclaimed through optimization. A report sample appears at the end of the procedure that follows.

To create a Storage Pool Summary report:

1. In the left pane of the Director UI, select the vault for running a Storage Pool Summary report.
2. From the Reporting menu, choose Storage Pool Summary.  
The Report wizard appears.
3. On the Welcome screen, click **Next**.
4. On the Scope screen, do one of the following:
  - To include data for a specific task, click **By Computer**.
  - To include task data for the entire vault selected in the Director UI, select **Entire Vault**.
5. (Optional) To include secondary storage information in the report, select the **Include secondary storage in this report** option.
6. Click **Next**.
7. (If applicable) If you are running the report for a specific computer, the Select Organization/ Customer screen appears. Select the organization or customer for the report, and then click **Next**. On the Select Location screen, select the location for the report, and then click **Next**. On the Select Computer screen, select the computer for the report, and then click **Next**. On the Select Task screen, select the task for the report, and then click **Next**.
8. On the Command Execution Time screen, do one of the following:
  - To run the report, click **Submit job immediately**, and then click **Next**.

- To create a schedule for running the report, click **Schedule job**, and then click **Next**. On the Select the Execution cycle screen and subsequent screens, create a schedule for running the report. For more information, see [Create schedules for running reports](#).

9. On the Report Wizard completion page, review the report information, and then click **Finish**.

### 11.8.1 Storage Pool Summary report sample

Vault: VAULT1

Carbonite Server Backup | Director Pool Report

task details: summaryall task \*

task started at 3-JAN-2022 09:37:30.35 -0500

Analyzing task: 'CUSTOMER1/LOCATION1/COMPUTER1/TASK1'

UPGV5\_CD529D22-C209-47FE-8762-5A1143C224CE Online Pool

total physical size of the pool: 1,149,660 (1.1 MB)

total actual data in the pool: 855,841 (835.8 KB)

total space to reclaim, if optimized: 293,819 (286.9 KB)

Analyzing task: 'CUSTOMER1/LOCATION1/COMPUTER1/TASK2'

UPGV5\_F35AF390-C9C7-4274-9AD3-CA96037CEF8C Online Pool

900,036 (878.9 KB)

total actual data in the pool: 848,916 (829.0 KB)

POOL-I-0000 total space to reclaim, if optimized: 51,120 (49.9 KB)

Analyzing task: 'CUSTOMER1/LOCATION1/COMPUTER1/TASK3'

UPGV5\_10EC0212-116C-4400-A8A2-6A618EA18F74 Online Pool

total physical size of the pool: 1,217,602,629 (1.1 GB)

total actual data in the pool: 1,217,601,413 (1.1 GB)

total space to reclaim, if optimized: 1,216 (1.2 KB)

Analyzing task: 'CUSTOMER1/LOCATION1/COMPUTER1/TASK4'

UPGV5\_C7CB069D-FD04-4241-8A34-03AEFB433EAD Online Pool

total physical size of the pool: 1,658,906,301 (1.5 GB)

total actual data in the pool: 1,658,906,301 (1.5 GB)

total space to reclaim, if optimized: 0 (0 bytes)

Jan03 09:37:34.118 [3732] POOL-I-0001

Analyzing task: 'CUSTOMER1/LOCATION1/COMPUTER1/TASK5'

UPGV5\_C522A150-0209-40B9-8D3D-A1B33331ACB3 Online Pool

total physical size of the pool: 1,217,262,658 (1.1 GB)

total actual data in the pool: 1,217,260,154 (1.1 GB)

total space to reclaim, if optimized: 2,504 (2.4 KB)

Totals:

total physical size of analyzed pools: 4,095,821,284 (3.8 GB)

total actual data in analyzed pools: 4,095,472,625 (3.8 GB)

total space to reclaim, if optimized: 348,659 (340.5 KB)



POOL-I-0000 total # of errors: 0

```
=====
Total # of tasks processed 5
# of successful tasks 5
# of disabled tasks 0
# of tasks with error(s) 0
# of errors in failed tasks 0
task completed at 3-JAN-2022 09:37:34.18 -0500
=====
```

## 11.9 Select the destination for a report

When creating a report, you can specify whether to save a report as a log file or e-mail the report to a recipient. You can select the report destination for a Storage, Storage Location, Vault Storage, Last Backup Status, or Late Server Status report.

To select the destination for a report:

1. When creating a report using the Report wizard, navigate to the Report Media screen. For more information, see [Create a Storage report](#), [Create a Storage Location report](#), [Create a Vault Storage report](#), [Create a Last Backup Status report](#) or [Create a Late Server Status report](#).
2. On the Report Media screen of the Report wizard, select one of the following Output Location options:
  - To save the report as a log file, select **Save report in a vault log file**.  
*Note:* This option is not available for the Vault Storage report.
  - To email the report to one or more recipients, select **Send generated report by Email to**. In the field below this option, enter one or more email addresses to send the report to. Multiple email addresses must be separated by commas.  
*Note:* Before reports can be sent by e-mail, an SMTP server must be specified for the vault. For more information, see [Configure email notifications](#).  
*Note:* When creating a Storage report, you can also specify whether data appears in columns or in CSV format on the Report Media screen. For more information, see [Create a Storage report](#).
3. Continue creating the report.

### 11.9.1 Select the scope for a report

When creating a report, you can specify whether to run the report for an entire vault, or restrict the report to a specific organization/customer, location, computer or task.

You can select the report scope for a Storage, Storage Location, Last Backup Status, or Late Server Status report. You can also select the report scope for a Missed Backups or Storage Pool Summary report, using different procedures than described here. For more information, see [Create a Missed Backups report](#) and [Create a Storage Pool Summary report](#).

To select the scope for a report:

1. When creating a report using the Report wizard, navigate to the Select Scope screen.
2. Do one of the following:
  - To include data for the entire vault selected in the Director UI, click **Vault**, and then click **Next**.
  - To restrict the report to a specific organization or customer, click **Organization/Customer**, and then click **Next**. On the Select Organization/Customer screen, select the organization or customer for the report, and then click **Next**.
  - To restrict the report to a specific location, click **Location**, and then click **Next**. On the Select Organization/Customer screen, select the organization or customer for the report, and then click **Next**. On the Select Location screen, select the location for the report, and then click **Next**.
  - To restrict the report to a specific computer, click **Computer**, and then click **Next**. On the Select Organization/Customer screen, select the organization or customer for the report, and then click **Next**. On the Select Location screen, select the location for the report, and then click **Next**. On the Select Computer screen, select the computer for the report, and then click **Next**.
  - To restrict the report to a specific task, click **Task**, and then click **Next**. On the Select Organization/Customer screen, select the organization or customer for the report, and then click **Next**. On the Select Location screen, select the location for the report, and then click **Next**. On the Select Computer screen, select the computer for the report, and then click **Next**. On the Select Task screen, select the task for the report, and then click **Next**.
3. Continue creating the report.

### 11.9.2 Create schedules for running reports

After creating a report, you can run the report immediately, or create a schedule for running the report. For example, you can create a schedule for running a report monthly or weekly on a specific day and time.

To create a schedule for running a report:

1. When creating a report using the report wizard, navigate to the **Select the Execution Cycle** screen.
2. Do one of the following:
  - To run the report on one or more days each week, click **Weekly**, and then click **Next**. On the **Weekly Cycle** screen, click each day to run the report each week. In the **Time** field, enter the time when you want to run the report each day. Click **Next**.
  - To run the report on one or more days each month, click **Monthly**, and then click **Next**. On the **Monthly Cycle** screen, click each day to run the report each month. In the **Time** field, enter the time when you want the run the report each date. Click **Next**.
3. On the completion screen, review the report information, and then click **Finish**.

### 11.9.3 View reports

After running a report, you can find the report in the Director UI and view the report. For information about running reports, see [Configure and run reports](#).

*Note:* The following procedure is used to open and view most reports in the Director UI. To view a Storage Pool Summary report for a single task, see [View a Storage Pool Summary report for a task](#).

To view a report:

1. In the left pane of the Director UI, select the report's vault.
2. In the right pane of the Director UI, double-click **Logs**.

A list of report and log files appears in the right pane of the Director UI. Report files have the following names:

Report	File Name in the Director UI
Storage	StorageUsageReport
Storage Locations	StorageLocationReport
Vault Storage	VaultStorageReport
Last Backup Status	LastBackupStatusReport
Late Server Status Report	LateServerStatusReport
Missed Backups	MissedBackupsReport
Storage Pool Summary	Pool Summary

3. (As applicable) If the report you want to view does not appear in the right pane of the Director UI, click **Refresh** to update the report and log file list.
4. In the Log name column, double-click the report you want to view.  
The report file appears in a word processor or text editor window.

#### 11.9.4 View a Storage Pool Summary report for a task

After you run a Storage Pool Summary report for one task, the report is available in the task's Logs folder in the Director UI.

To view a Storage Pool Summary report for a task:

1. In the left pane of the Director UI, expand the vault, customer, location, computer and backup tasks, and task for the report.
2. Click the Logs folder for the task.
3. In the right pane of the Director UI, double-click the Pool Summary report log.

## 12 Automate maintenance processes

Automated maintenance enforces safeset retention settings, optimizes and deduplicates data, and checks pools for physical corruption. Automating maintenance can improve the efficiency of jobs such as backups and restores.

Beginning in Director 8.60, when automated maintenance is enabled on a vault, some maintenance operations are triggered to run for a task after another operation finishes. See [Triggered maintenance and replication](#).

Other maintenance operations are scheduled to run. See [Scheduled maintenance operations](#).

### 12.1 Scheduled maintenance operations

The following scheduled maintenance operations are available:

- Purge expired logs
- Pool system optimization. By default, this operation (vvpoolop dedup) runs every Saturday.
- CheckCRC for the whole vault. By default, this operation runs once per month.
- Report maintenance jobs older than 14 days
- Delete expired secondary safesets
- Optimize secondary pools
- Replication satellite vault

*Note:* Online area migration for the whole vault is also available but, beginning in Director 8.60, cannot be scheduled to run. Instead, migration runs as part of the 'vvpoolop optimize' process that is triggered after backups and replications. Schedules for both 'vvmigrat online' and 'vvpoolop optimize' will be ignored.

Scheduled maintenance operations run as VVQmanager processes during periods of low activity (disk I/O and CPU). To avoid duplicate maintenance operations, the Queue Manager verifies if newly submitted jobs are the same as the jobs currently running. If the jobs are identical, the Queue Manager cancels the running jobs and starts the new jobs. To increase efficiency and prevent the locking of files and resources for extended periods, the system breaks large operations into smaller operations.

Tasks are sometimes locked to prevent multiple operations from processing the same data. When a locked task is identified, it is added to a list of tasks that are not available for scheduled maintenance. The list of unavailable tasks expires in 5 minutes and the maintenance operation is run again.

Each day, the system cancels scheduled maintenance operations older than 14 days that have not yet started or finished. Cancelled operations are noted in the maintenance log and in e-mails to the system administrator, if an email address is provided for the vault. For more information, see [Configure email notifications](#). If scheduled maintenance operations do not run for 14 days, the vault might be under-powered or under-sized.

## 12.2 Enable or disable automated maintenance

You can enable or disable automated maintenance on a vault.

If you disable automated maintenance:

- Scheduled maintenance operations will not run. Pending operations are removed from the maintenance queue.
- Triggered maintenance operations, described in [Triggered maintenance and replication](#), will not run.

To enable or disable automated maintenance:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Vault Settings**.
3. Click the **Maintenance** tab.
4. Do one of the following:
  - To enable automated maintenance on the vault, select the **Enable maintenance host on <Vault name>** checkbox.
  - To disable automated maintenance on the vault, clear the **Enable maintenance host on <Vault name>** checkbox.

*Note:* If automated maintenance is disabled, neither scheduled nor triggered maintenance operations will run.
5. Click **OK**.

## 12.3 Verify that automated maintenance is running

To verify that automated maintenance is running:


1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Services**.
3. Verify that the Queue Manager state is **Running**.
4. Click **Close**.
5. Click **Job Monitor** and select **Show Internal Jobs**.
6. Verify that the Maintenance Host status is **Running**.
7. Click **Close**.

## 12.4 Stop automated maintenance

Stopping the Maintenance Host stops scheduled maintenance operations that are in progress and removes pending operations from the maintenance queue. When the Maintenance Host restarts, operations that were already in progress restart automatically.

Stopping the Maintenance Host does not stop triggered maintenance operations that are in progress, but no new maintenance operations will be triggered.

To stop automated maintenance:

1. Click Job Monitor: 
2. Select Maintenance Host and click Stop.
3. Click **Yes**.

## 12.5 Run a maintenance operation on demand

To run a maintenance operation on demand:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Select a maintenance operation in the **Description** list.
4. Click **Run Now**.
5. In the confirmation dialog box, click **Yes**.
6. Click **OK**.

## 12.6 Change the time of a scheduled maintenance operation

To change the time of a scheduled maintenance operation:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Select a maintenance operation in the **Description** list.
4. Click **Edit**.
5. Complete the Change Schedule Entry wizard.
6. Click **Finish**.

## 12.7 Enable automated secondary storage maintenance

To automate secondary storage pool maintenance, you can enable the following scheduled maintenance operations which are disabled by default:

- Delete expired secondary safesets
- Optimize secondary pools


To enable automated scheduled storage maintenance:

1. Select a vault in the left pane of the Director UI.

2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Select Delete expired secondary safesets in the Description list.
4. Click **Enable**.
5. Select Optimize secondary pools in the Description list.
6. Click **Enable**.
7. Click **OK**.

## 12.8 Verify maintenance performance

To verify maintenance performance:

1. Click Job Monitor: 
2. Select **Maintenance Host**.
3. Click **Progress**.

The Progress Information dialog box shows the number of maintenance operations with the following statuses:

- **Running** – Scheduled maintenance operations that are currently running and using computer resources. A fixed number of operations can run at a time, even though the other queues may add up to more than the limit. A running operation can have multiple waiting and pending objects.
- **Waiting** – Submitted maintenance operations that are idle and waiting for an event, such as another spawned task to finish.
- **Pending** – Queued maintenance operations that are waiting for resources.

*Note:* The running, waiting and pending counts do not include triggered maintenance operations.

If the number of waiting and pending maintenance operations is continually increasing, contact Support for assistance.

4. Click **Close**.
5. Click **Close** again.

## 12.9 View maintenance logs

Maintenance logs that are generated during a failed backup remain in the global log location because of storage policy restrictions. If your storage policy limit for disk storage is exceeded, logs cannot be written to storage. Normally, backup logs are moved to the task log location.

Log names for maintenance operations include the characters “Mc”. Log names for maintenance operations for all tasks include the characters “McG”, where G means Global scope.

To view maintenance logs:

1. Select a connection in the left pane of the Director UI.

2. Double-click **Logs** in the right pane.
3. Double-click a maintenance log file in the right pane.

## 12.10 Create a custom scheduled operation

The command or batch file must reside in the <...>\Director\prog directory. You cannot edit a custom scheduled entry after you create it. To edit a custom scheduled entry, you must delete it and then recreate it. To view log entries for custom scheduled entries, view the spawn logs.

To create a custom scheduled operation:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Click **Custom**.
4. Complete the Custom Command Scheduling wizard. For supported commands, see the [Command Reference](#).
5. Click **Finish**.

## 12.11 Disable a scheduled operation

To disable a scheduled operation:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Select an enabled scheduled operation in the **Description** list.
4. Click **Disable**.
5. Click **OK**.

## 12.12 Enable a scheduled operation

To enable a scheduled operation:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Select a disabled scheduled operation in the **Description** list.
4. Click **Enable**.
5. Click **OK**.



## 12.13 Remove a scheduled operation

To remove a scheduled operation:

1. Select a vault in the left pane of the Director UI.
2. From the **Vault Maintenance** menu, choose **Schedule Entries**.
3. Select a scheduled operation in the **Description** list.
4. Click **Remove**.
5. Click **OK**.

## 13 Monitor and manage vaults

Director includes tools for monitoring and managing vault processes, including the following:

- Notifications. You can set up email notifications when jobs fail or when other events occur in a vault. See [Configure email notifications](#) and [Set primary storage thresholds](#).
- Job Monitor. See [View and stop vault processes in the Job Monitor](#).
- Vault Services dialog box. See [View, start and stop vault services](#).
- Logs. See [View log files](#).

To ensure that a Director vault database is protected and can be recovered in case of failure, see [Back up a vault database](#).

### 13.1 View and stop vault processes in the Job Monitor


Using the Job Monitor, you can view current and recent vault processes, including:

- Backups
- Restores
- Replication
- Internal jobs
- Completed jobs. Completed jobs appear in the Job Monitor for one hour. You cannot edit this setting.
- Processes waiting to run (in a queue)

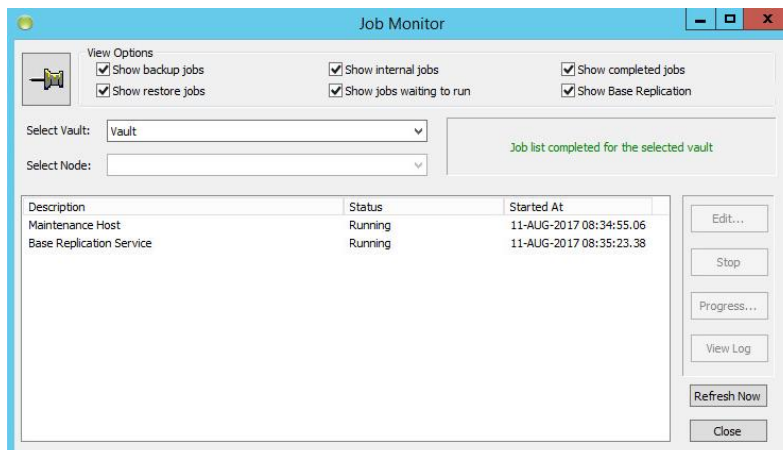
You can also stop processes that are running.


During periods of high activity, the Job Monitor might be slow to respond to commands. In addition, a large number of jobs that are running concurrently can affect the responsiveness of Job Monitor commands.

To view vault and stop processes in the Job Monitor:

1. In the Director UI, click **Job Monitor**: 

The Job monitor shows processes running on the vault.



2. In the **Select Vault** list, click the vault with activities that you want to view.
3. Select the check box for each process type that you want to view. Available process types include:
  - Show backup jobs
  - Show restore jobs
  - Show internal jobs
  - Show jobs waiting to run
  - Show completed jobs
  - Show replication (if configured)
4. To view progress information for a vault process, click the process, and then click **Progress**.
5. To view the log for a process, click the process, and then click **View Log**.
6. To update information in the Job Monitor, click **Refresh Now**.
7. To stop a vault process, click the process, and then click **Stop**. If a confirmation dialog box appears, read the information. Click **Yes** to stop the process.
8. To permanently display the Job Monitor on your desktop, click the Pin icon: 

## 13.2 View, start and stop vault services

When you install Director, the following services are automatically initialized and started:

- Admin Service — authorizes requests from the Director UI to the vault.
- Listener — waits for customer connection requests and then spawns a Director process.
- Replication Service — controls and runs vault replication for 1:1, N:1 and N:1:1 replication configurations. This service starts automatically for vaults that are licensed for replication.
- Queue Manager — monitors and manages Director specific job queues.

*Note:* Stopping the Queue Manager service stops any jobs that have started. When you start the Queue Manager again, the jobs start again. As a result, you may see two operations running and two logs generated, one of which will have errors, and one that should be successful.
- Scheduler — monitors and spawns regular scheduled Director functions.
- Node Health Monitor — monitors the health status of a vault node

You can view the status of these vault services, and start, restart and stop the services using the Director UI.

**IMPORTANT:** Stopping or restarting vault services using this method could result in data corruption. Instead, we recommend “ramping down” the vault. See [Ramp down a vault](#).

When you install Director, you can also install the Reporting service. When registered to Carbonite Server Backup API, the Reporting service sends data to and receives messages from the API. The Reporting service does not appear in the Director UI and cannot be managed from the UI. Instead, view the Reporting service using Microsoft Management Console.

*Note:* You can open Reporting service logs from the Director UI. See the procedure for viewing global logs in [View log files](#).

To view, start and stop vault services:

1. Select a vault in the left pane of the Director UI.
2. On the **Vault Maintenance** menu, click **Services**.

The Vault Services dialog box lists Director services and indicates whether they are running.

3. Click a service in the list. To select multiple services, hold down **CTRL** and click each service.
4. Do one of the following:
  - To start the selected services, click **Start**.
  - To restart the selected services, click **Restart**.
  - To stop the selected services, click **Stop**.
5. Click **OK**.

### 13.2.1 Ramp down a vault


To shut down the vault gracefully after current processes are finished, you can “ramp down” the vault using this procedure. Depending on the operations that are in progress, it can take some time to ramp down a vault.

You can also ramp down a vault using the vaultop set\_node\_state rampdown command. See [vaultop](#).

To ramp down a vault:

1. In the left pane of the Director UI, click the plus sign beside the vault that you want to ramp down, and then click the plus sign beside **Nodes**.

The vault server name appears.



2. Right-click the vault server name and click **Rampdown** in the context menu.

## 13.3 View log files

You can open log files from the right pane of the Director UI. By default, two log folder types are available in the Director UI:

- Task logs. A task log provides information about a backup, migration, restore or synchronization for a specific task.
- Global logs. A global log provides information about scheduled and ad hoc vault processes that are not associated with a specific task.

You can also choose to view the following log file types:

- Spawn logs. Spawn logs provide detailed information about system tasks and custom commands.

- License logs. License logs provide vault licensing information.
- Server logs. Server logs show ServerHost logs created by VVServer, a vault service that controls backups and restores.
- Replication logs. For more information, see [View replication logs](#).

To show these log file types in the Director UI, see [Select log file types to display](#). You can also access log files through the Job Monitor. See [View and stop vault processes in the Job Monitor](#).

When you view a log file that is larger than 256 KB, the log file can be abbreviated. If part of a log file is not shown, it is replaced by diagonal lines and the following text: This file was abbreviated for viewing.

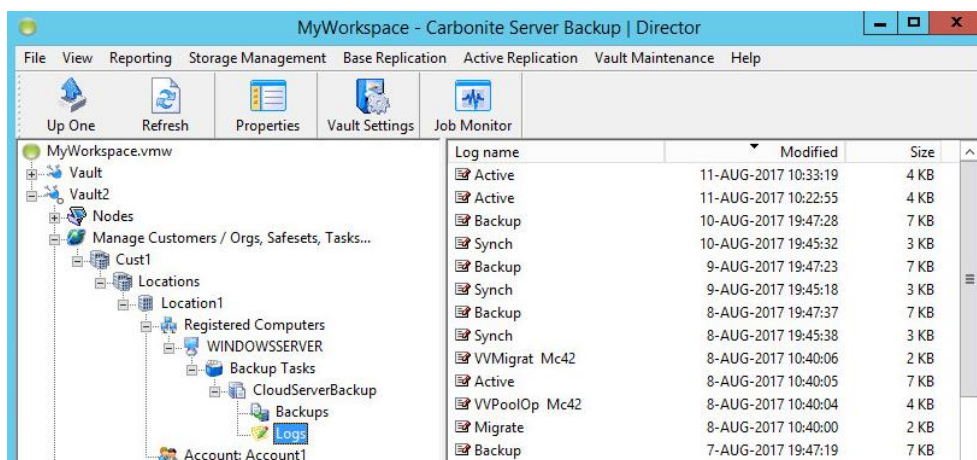
If you use policies and quotas, the following error message might appear:

VVLT-W-0002 Log file will be kept in the global logs folder. Error(2): Vault is out of storage

If this occurs, you can view the log file in the Job Monitor initially. To locate it later, go to the global logs folder.

To view log files:

1. In the Director UI, do one of the following:
  - To view a global log, click a vault in the left pane, and then double-click **Logs** in the right pane.
  - To view a task log, expand the vault, customer, location, computer, and task in the left pane, and then click the **Logs** folder for the task.



- To view a spawn log, click a vault in the left pane of the Director UI, and then double-click **Spawn Logs** in the right pane.
- To view a license log, click a vault in the left pane of the Director UI, and then double-click **License Logs** in the right pane.
- To view a server log, click a vault in the left pane of the Director UI, and then double-click **Server Logs** in the right pane.

If the Spawn Logs, License Logs, Server Logs option is not available, see [Select log file types to display](#).

2. Double-click a log file in the right pane.

3. If a message asks whether you want to abbreviate the log file, click **Yes** to hide parts of the log file. Click **No** to show the entire log file.

## 13.4 Select log file types to display

To select log file types to display:

1. Select a vault in the left pane of the Director UI.
2. From the **View** menu, choose **Options**.
3. Select one or more of the following in the **Logs** area:
  - Show Spawn Logs and License Logs
  - Show Replication Logs
  - Show Server Logs
4. Click **OK**.

## 13.5 View replication logs

When replication is used, a log is generated for each replication service and session.

Configuration and metadata replication logs are available in a vault's Replication Logs folder. Each replication log name includes the vault's replication role in the activity (Active, Passive, Base or Satellite) and the type of information being replicated (Config or Meta).

Task data replication logs are available in a task's Logs folder. Each task replication log name is the vault's replication role in the activity (Active, Passive, Base or Satellite).

To view a configuration or metadata replication log:

1. Select a vault in the left pane of the Director UI.
2. Double-click **Replication Logs** in the right pane. If replication logs are not available, see [Select log file types to display](#).
3. Double-click a replication log in the right pane.

To view a task replication log:

1. Expand a vault in the left pane of the Director UI.
2. Expand the **Manage Customers** list.
3. Expand a Customer list.
4. Expand Locations.
5. Expand a location.
6. Expand Registered Computers.
7. Expand a computer.
8. Expand Backup Tasks.

9. Select a task.
10. Double-click **Logs** in the right pane of the Director UI.
11. Double-click a replication log in the right pane.

## 13.6 Log message codes

Each log message includes the following information:

- Date and time
- Four-letter code that describes the task. See the list of task codes below.
- Single letter code that describes the severity of the issue. See the list of severity codes below.
- Four-digit number for Support personnel

In the following sample log message, REPL is the task code and I is the severity code:

```
Jan15 15:05:56.109 [1696] REPL-I-0001 Collecting children for item: CBTd907bf01-a5a1-41cf-a52a-de8288611deaServer Agent
```

*Note:* Some messages, such as messages with the REPL task code, might display unprintable characters as empty boxes. These boxes represent special characters that are necessary for the item name.

### Log Message Task Codes

The following task codes appear in log messages:

Code	Description
VVLT	A general message
DISK	Disk I/O
DEVC	Device
POOL	Pool system
COPY	Copying
VVCP	vvcopy
MIGR	Migrate
SSET	Safeset
MERG	Merge
PMGR	Pool Manager
DAEM	Daemon
EVLT	Log header

Code	Description
B RTP	Backup Restore Transfer Protocol
SRVC	Service
MOVE	Move
ALYZ	Analyze
UTIL	Activation utility
UPDT	Verify license key
REPL	Vault Replication

### Log Message Severity Codes

The following severity codes appear in log messages:

Code	Description
I	Informational
W	Warning
E	Error
F	Failure
S	Success

## 13.7 Purge activity records and logs

You can purge, or delete, activity records and logs. It is recommended that you purge log files weekly to save space on the vault.

Unlike log files, activity records are not deleted automatically. Activity records list the time of completion, file size, and success messages for backup, restore, and migrate tasks. Activity report information is used to generate reports.

To purge activity records and logs:

1. Select a vault in the left pane of the Director UI.
2. From the Vault Maintenance menu, choose Purge Activity/Logs.  
The Purge wizard appears.
3. On the Welcome page, click **Next**.
4. On the Select Purge Type page, do one of the following:
  - To delete log files, select **Purge log files**, and then click **Next**.



- To delete activity records, select **Purge activity records**. In the **Purge activities older than x days** field, enter number of days of activity records that you want to keep. Click **Next**.
5. On the Command Execution time page, do one of the following:
    - To delete files immediately, select **Submit job immediately**, and then click **Next**.
    - To delete files later, select **Schedule job**, click **Next**, and then specify a weekly or monthly schedule for deleting the files.
  6. Click **Finish**.

## 13.8 Schedule a log file purge

To schedule a log file purge:

1. Select a connection in the left pane of the Director UI.
2. From the Vault Maintenance menu, choose Vault Settings.
3. Click the **Logs** tab.
4. Enter or select the number of days old that a log file must be to be deleted in the **Purge logs after** field.  
  
For example, if the value is 180, logs are only deleted
5. Enter or select the minimum number of log files to keep in the **Minimal number of log files** field.
6. Click **OK**.

## 13.9 Configure email notifications

Director can send automated email notifications for events such as job failures and replication failures.

You must specify an SMTP server for sending email notifications. The SMTP server is also used for sending emailed reports. See [Select the destination for a report](#).

Beginning in Director 8.70, you can specify an SMTP server that requires authentication and uses a secure email protocol (StartTLS or SSL/TLS). You can specify a private SMTP server or a cloud email service provider (e.g., Amazon Simple Email Service (SES), Google Gmail email service or Microsoft 365). If you specify a cloud email service provider, please see the provider's documentation for requirements.

You can configure email notifications using the Director UI, as described in this procedure. You can also configure email settings using the vaultop command. See [vaultop email subcommands](#).

You can set a maximum size for Director emails. This size limit applies to email notifications and emailed reports. See [Change the maximum email size](#).

To configure email notifications:

1. From the **Vault Maintenance** menu, choose **Vault Settings**.
2. In the Vault Settings dialog box, click the **Email Notifications** tab.

3. In the Notification area, select the events for sending email notifications:
  - **On successful job completion** — If selected, notifications are sent when jobs finish successfully.
  - **On job failure** — If selected, notifications are sent when jobs fail.
  - **On device notification** — If selected, notifications are sent when Universal Naming Convention (UNC) connections are inaccessible.
  - **On replication failure** — If selected, notifications are sent when replication sessions fail.
4. In the SMTP Configuration area, enter SMTP server settings in the following fields:
  - **Server** — Enter the fully-qualified domain name (FQDN) of the SMTP server for sending email notifications.
  - **From** — Enter the email address that will appear as the From address in notification emails. To include a display name, specify the display name followed by the email address in angle brackets (e.g., Director Alerts <alerts@company.com>).

*Note:* Emails from some service providers might not include the entire From address. For example, emails might only include the display name.
  - **Security** — Select the email encryption type: None, SSL/TLS or StartTLS.

We do not recommend “None” for the encryption type. Unencrypted emails can be read if they are intercepted.

The **Port** field is automatically set to the appropriate port number for the selected encryption type. You can edit this value, if required.
5. In the **Specify address(es) of recipients** field, enter each email address that will receive email notifications. Use a comma as a separator between email addresses (e.g., user1@company.com,user2@company.com). Display names are not supported in this field.
6. If the specified SMTP server requires authentication, select the **Require Authentication** check box. In the **Username** and **Password** fields, enter credentials for authenticating with the SMTP server.
7. (Optional) To test whether email notifications can be sent using the specified settings, click **Send Test Email**.

*Note:* You can test the specified settings before saving them.
8. Click **OK**.

## 13.10 Set primary storage thresholds

You can set warning and critical thresholds for primary storage in vault. One or more users receive email notifications when the primary storage group reaches a specified threshold:

- **Warning** — All primary storage locations are running out of disk storage. Disk usage is above warning level.
- **Error** — All primary storage locations are running out of disk storage. Disk usage is above critical level. Backups may start failing until more disk storage is available.
- **Maintenance has been disabled** — After making more disk storage available, please enable Maintenance.

- Normal — Disk usage is OK.

To set primary storage thresholds:

1. Ensure that email notifications are configured for the vault. See [Configure email notifications](#).
2. Select a vault in the left pane of the Director UI.
3. From the Storage Management menu, choose Storage Monitor Configuration.
4. Enter a value in the **Warning Threshold** field. You can enter values from 15 to 95. The default value is 80.
5. Enter a value in the **Critical Threshold for Backups** field. This value must be greater than the value in the **Warning Threshold** field. You can enter values from 20 to 98. The default value is 98.
6. Click **OK**.

## 13.11 Modify advanced vault settings

To modify advanced vault settings:

1. From the Vault Maintenance menu, choose Vault Settings.
2. Click the **Advanced** tab.
3. Edit fields on the **Advanced** tab:

Field	Description
<b>Communication</b>	
Server ports	The server ports configured for the vault. The default port is 2546. Separate ports with a comma. Multiple ports are attempted in order from left to right.  If you change ports, you must restart the Listener service or the computer.
Socket timeout	The time in seconds that the socket remains open as the Listener waits to complete an agent request. The default is 30 seconds.  Lowering this value can cause communication errors. Raising this value can consume system resources and cause delays.
Maximum socket connections	The maximum number of allowed socket connections. The default is 250.
<b>Optimization</b>	
Open pool files limit	The number of pool files that run concurrently. The default is 100.
Maximum number of jobs	The number of jobs that can run concurrently. The default is 25

DBA password	
New password	Resets the current Database Administrator password. The new password must adhere to the Windows password policy.  If you change the DBA password, you must restart the vault services.
Confirm password	Confirms the new Database Administrator password.

4. Click **OK**.

## 13.12 View, add and remove a vault's alternate vaults

When you set up replication between vaults, vaults that receive replicated data are automatically configured as “alternate vaults” for the source vaults. For example, in 1:1 replication, Passive vaults are configured as alternate vaults for Active vaults. When an agent backs up data to or restores data from a vault, the agent also tries to connect to each of the vault's alternate vaults.

You can view a vault's alternate vaults on the Agent Connection tab of the Vault Settings dialog box. You can also manually add and remove alternate vaults using this tab. However, adding alternate vaults for a vault can increase the amount of time required for a backup or restore.

To view, add, or remove a vault's alternate vaults:

1. In the left pane of the Director UI, select the vault for viewing, adding, or removing alternate vaults.
2. From the Vault Maintenance menu, choose Vault Settings.
3. Click the **Agent Connection** tab.

The tab lists the IP address of the vault's alternate vault, if any.

4. (Optional) Do one of the following:
  - To add an alternate vault for the selected vault, enter the alternate vault's IP address in the **New address** field, and then click **Add**.
  - To remove an alternate vault from the selected vault, select the alternate vault's IP address, and then click **Remove**.
5. Click **OK**.

## 13.13 Back up a vault database

A Director vault database includes important configuration, node and account information, and should be backed up regularly.

Use the dbbackup utility to back up the database. The dbbackup utility is installed with a vault.

By default, dbBackup runs automatically at midnight every day. The scheduled dbBackup automatically extracts the database data from the <database> directory, to the data.bin file. When dbBackup is running, the vault database is read only.

If a folder is not specified during the backup, the C:\DBBACKUP folder is created. When a folder is not specified during a restore, this default folder is used. The default user name and password are used. The default username and password are stored in a VaultDB.cfg.

To view information about a backup, open the DBBackup log file in the Director\logs directory. You can also use the Job Monitor to view backup information.

To view the syntax and options available for the dbBackup command, see [dbbackup](#) or open a command prompt and enter the following command:

```
dbbackup /?
```

## 14 Command Reference

This command reference provides information about the following Director commands:

- [agentinfoports](#)
- [dbbackup](#)
- [migratesecondary](#)
- [replvault](#)
- [secondaryop](#)
- [vaultop](#)
- [Vault Settings](#)
- [vanalyz](#)
- [varchive](#)
- [vvexport](#)
- [vimport](#)
- [vmigrat](#)
- [vmove](#)
- [vpoolop](#)
- [vvpurge](#)
- [vrecall](#)

*Note:* Only the listed commands are supported.

The command documentation uses the following conventions:

- Variables appear in *italics*.
- Optional parameters appear in square brackets [ ].
- Sets of options appear in curly brackets { }. Options are separated by vertical lines |.

If you must enter a password for a command, the password appears as clear text.

Commands that you run appear in the log files.

You can use the Custom Command option to run a group of Director commands in a batch file. Add the batch file to the prog folder in the Director installation folder. Specify the FORCEDOS option.

If you want the Scheduler to run a program, you must copy the program to the prog folder in the Director installation folder.

## 14.1 agentinforeports

The agentinforeports command writes information about the agents registered on the vault into comma-separated value (CSV) files in the Director\bill directory.

Unless the command is used with the /check subcommand described below, the agentinforeports command creates LicenseUsageReport\_YYYYMMDD.csv and DatatypeReport\_YYYYMMDD.csv files, where YYYYMMDD is the year, month and day when the report was generated. The LicenseUsageReport CSV file includes version, type, and license information for each agent that is registered to the vault. The DatatypeReport CSV file includes the version and type of each agent that is registered to the vault, and provides pool footprint information for each task.

### Syntax

```
agentinforeports [subcommand]
```

### Subcommands

Subcommand	Description
/id	Writes customer IDs instead of customer names in the CSV files.
/check <i>filename</i>	Verifies that the file has not been modified since it was generated. A <i>filename</i> must be specified.
/defaultquotas	Sets all customer quotas to equal current usage for licensing of Agent features.
/simdefaultquotas	Simulates setting all customer quotas to equal current usage for licensing of Agent features.

## 14.2 dbbackup

The dbbackup command backs up and restores the database for a vault. The command also reduces the size of the database transaction log file. For more information, see [Back up a vault database](#).

The command produces a DBBackup log file in the Director\logs directory.

### Syntax

```
dbbackup subcommand
```

### Subcommands

Subcommand	Description
backup <i>directory</i>	Creates a backup of a vault database. <i>directory</i> is the physical path to the location where the database backup will be saved.
restore <i>directory</i>	Restores a previously backed up vault database. <i>directory</i> is the physical path to the location of the database backup to restore.

Subcommand	Description
shrink	Shrinks a transaction log file.

## 14.3 migratesecondary

Moves safesets from primary storage to the assigned secondary storage group.

### Syntax

To migrate specific safesets to secondary storage, use the following command:

```
migratesecondary taskid=taskid safeset=x,y,...z
[last_online] [nointerrupt] [dedup={0|1}] [path=directory
[{user=userName password=password|useDefaultCredentials}]]
[preview={1|2}] [verify_every=value] [verify_last]
```

To migrate safesets to secondary storage based on retention settings, use the following command:

```
migratesecondary taskid=taskid use_retention [last_online]
[nointerrupt] [dedup={0|1}]
```

To specify migratesecondary command parameters in a text file, use the following command:

```
migratesecondary param=parameterFilename
```

### Parameters

Parameter	Description
<i>taskid=taskid</i>	Specifies the task with safesets to move to secondary storage. To move safesets from all tasks, use an asterisk (*) as the <i>taskid</i> .
<i>safeset=x,y,...z</i>	Specifies safesets to move to secondary storage, where <i>x</i> , <i>y</i> and <i>z</i> represent the safeset numbers to move. Separate safeset numbers with commas.
<i>use_retention</i>	Migrates safesets based on retention settings
<i>last_online</i>	Specifies that the last online safeset should be migrated to secondary storage. If you do not include this parameter, the last online safeset will never be migrated.
<i>nointerrupt</i>	Specifies that an incoming backup or replication would not cancel the migration.
<i>dedup</i> ={0 1}	Specifies whether deduplication should be enabled. To enable deduplication, specify <i>dedup</i> =1. To disable deduplication, specify <i>dedup</i> =0 (zero). The default value is 1 (deduplication enabled).



Parameter	Description
path= <i>directory</i> [ <i>{user=userName</i> <i>password=password </i> <i>useDefaultCredentials}</i> ]	Copies safesets to the specified location without modifying the primary storage.  The user/password or useDefaultCredentials parameters are used when the <i>directory</i> is a UNC storage location.
preview={1 2}	Specifies that you want to preview which safesets are going to be migrated. To preview the safesets which will move to secondary storage, specify preview=1. To preview the safesets which will move to secondary storage and the amount of storage required in secondary storage, specify preview=2.
verify_every= <i>value</i>	Specifies which safesets to verify, if any. To verify every safeset, specify 1 as the value. To verify every second safeset, specify 2. To verify every third safeset, specify 3. To not verify safesets, specify 0 (zero).
verify_last	Verifies the last safeset.
param= <i>parameterFilename</i>	Specifies all parameters in a text file. The parameter file format is shown below.

### Parameter File Format

The parameter file has the following format:

=====

```

CONFIGVERSION=3
CONTENT_VERSION=1
# Global flags:
ALLTASKS=TRUE           # set to TRUE or FALSE
MARK_AS_OFFLINE=FALSE   # set to TRUE or FALSE
MIGRATE_LAST_ONLINE=FALSE # set to TRUE or FALSE
PREVIEW=2               # 1 for safesets, 2 for safesets and storage sizes
VERIFY_EVERY=<value>    # verify every so many safeset (0 for no verification)
VERIFY_LAST=FALSE       # always verify the last safeset
DEDUP=TRUE              # enable or disable deduplication

# You MUST specify one and ONLY one of the following methods for migration:
# 1. Using a Business Rule or 2. Explicit Specification of Safesets

# METHOD 1: Using a Business Rule
# Specify ONLY one of the following 3 rules:
MAX_DAYS_ONLINE=90      # Specify max days online
# Or
MAX_SAFESSETS_ONLINE=2  # Specify max online safesets
# Or
USE_RETENTION=TRUE      # Use Archive retentions (TRUE or FALSE)

```

```
# Scope selection: The specified objects are processed based on the business rule.
CUSTOMER=1,4,5 # Process these customers
LOCATION=1,4,5 # Process these customer locations
COMPUTER=10,11 # Process these computers
TASK_1=25 # Process this task
TASK_2=28 # Process this task
# ... More tasks can go here ...

# METHOD 2: Explicit Specification of Safesets.
TASK_1=23:2,4,7,8 # Explicit migration. Process task 23, safesets 2,4,7,8
# ... More tasks with safesets can go here ...
```

## Examples

The following command migrates safesets 3, 4, 5 and 6 from task 23:

```
migratesecondary taskid=23 safeset=3,4,5,6
```

The following command copies online safesets 3, 4, 5 and 6 from task 23 to C:\secondary. This effectively creates a new detached pool system.

```
migratesecondary taskid=23 safeset=3,4,5,6 path=C:\secondary
```

The following command migrates expired safesets from task 23:

```
migratesecondary taskid=23 use_retention
```

The following command migrates expired safesets from all tasks:

```
migratesecondary taskid=* use_retention
```

The following command migrates expired safesets for all tasks and includes the last safest in the migration:

```
migratesecondary taskid=* use_retention last_online
```

The following command migrates safeset 3 from task 2, with deduplication enabled:

```
migratesecondary.exe taskid=2 safeset=3 dedup=1
```

The following parameter file could be used to migrate safesets 1 and 2 from task 2, with deduplication enabled and verification of every safeset:

```
CONFIGVERSION=3
CONTENT_VERSION=1
DEDUP=TRUE
TASK_1=1:1,2
VERIFY_EVERY=1
```

## 14.4 replvault

Controls and runs vault replication for 1:1 and N:1 configurations.

### Syntax

```
replvault subcommand [parameters]
```

### Subcommands

Subcommand	Description
-allowpassive	If this parameter is included in the command, the command allows a non-empty vault to become Passive. If this parameter is not included in the command, a non-empty vault cannot become Passive.
cfgfile fileName	Takes all parameters from configuration file.
DIFF <i>scope</i> [ <i>object-id</i> ] [-lastsafeset] [-diffmethod={IDX CRC DATA}]	<p>Compares replication objects and creates a difference report. Where:</p> <ul style="list-style-type: none"> <li><i>scope</i> — VAULT, CUSTOMER, LOCATION, COMPUTER, TASK, ACCOUNT, or USER</li> <li><i>object-id</i> — the object's database ID, not required for VAULT.</li> <li>-lastsafeset — Compares hash values for last safesets only.</li> </ul> <p>-diffmethod — Controls how safeset data is compared. Available -diffmethod options are:</p> <ul style="list-style-type: none"> <li>IDX — Compares pool indexes only (fast but least reliable)</li> <li>CRC — Uses file CRCs to compare data (fast and reliable but will not detect data block level pool file corruption)</li> <li>DATA — Reads and compares the entire backup data (very slow but will detect data block level corruption)</li> </ul> <p>If a -diffmethod option is not specified, replvault compares tasks based on their synch entries (safesets information) in the synch.dat file. This comparison method is faster than the -diffmethod=IDX and -diffmethod=DATA options.</p>
DISABLE	Disables processing of replication events.
ENABLE	Enables processing of replication events.
GENSATREPORT <i>scope</i> [ <i>object-id</i> ] [-lastsafeset] [-diffmethod=IDX CRC DATA]	Compares replication objects and creates a difference report for Satellite vaults.

Subcommand	Description
REGSAT basevault= <i>hostname</i> port= <i>port</i> otrk= <i>authorization_code</i> [altvault= <i>address:port</i> ] [-force]	Registers the Satellite vault on a Base vault. Where: <ul style="list-style-type: none"> <li>• basevault — The IP or hostname of the Base vault.</li> <li>• port — The port where the Base vault is listening.</li> <li>• otrk — The authorization code for the Satellite vault.</li> <li>• altvault — Base vault address:port for the agent connection (if different from &lt;basevault&gt; and &lt;port&gt;)</li> <li>• -force — Forces registration even if data exists on this Satellite vault which could be overwritten</li> </ul>
REPSAT [mailto= <i>email_address</i> ]	Generates a status report for all of the satellites registered on the Base vault. Email_address is the email address where the report is sent.
RESET [-scriptingmode] [-scriptingmode -allowpassive]	Resets 1:1 replication configuration, if it exists.  If the -scriptingmode parameter is not included, the command prompts for user input.  If the -scriptingmode parameter is included, the command runs in silent mode and does not require user input.  If the -scriptingmode -allowpassive parameters are included, the command runs in silent mode and allows a non-empty vault to become a Passive vault.  <i>Warning:</i> If you run the replvault reset command on a Passive vault, agents that were registered to the associated Active vault could start sending backups directly to the Passive vault. To prevent a Passive vault from receiving backups, ensure that the vault is offline before running the replvault reset command on it.
SETACTIVE [-force]	Changes the current vault replication role to active.  -force forces the replication role on the vault to active.

Subcommand	Description
<p>SETSATREPORT <i>satelliteID</i> <i>requestID=paramsFile</i></p> <p><b>OR</b></p> <p>SETSATREPORT <i>satelliteID scope</i> [<i>object-id</i>] [-lastsafeset] [-diffmethod={IDX CRC DATA}] [mailto=<i>email_address</i>]</p> <p><b>OR</b></p> <p>SETSATREPORT <i>satelliteID scope</i> [<i>object-id</i>] [-days=<i>x</i>] [-handlepaused] [mailto=<i>email_address</i>] -LATE</p>	<p>Compares replication objects and creates a difference report.</p> <ul style="list-style-type: none"> <li>• <i>satelliteID</i> – ID of the satellite vault for the report.</li> <li>• <i>paramsFile</i> – A file that contains parameters for the diff report.</li> <li>• <i>scope</i> [<i>object-id</i>] – <i>scope</i> is CUSTOMER, LOCATION, COMPUTER, TASK. <i>object-id</i> is the object's database ID.</li> <li>• -lastsafeset – Compares hash values for last safesets only.</li> <li>• -diffmethod – Controls how the safeset data gets compared: <ul style="list-style-type: none"> <li>• IDX – Compares pool indexes only (fast but least reliable)</li> <li>• CRC – Use file CRCs to compare data.</li> <li>• DATA – Read and compare the entire backup data.</li> </ul> </li> <li>• Emailto – The email address to which to send the report.</li> </ul> <p>To run the Replication Lag Report for N:1 replication, run the REPLVAULT SETSATREPORT -LATE command on a Base vault.</p> <ul style="list-style-type: none"> <li>• -days=<i>X</i> – The number of days behind in replication. The default is 5 days.</li> <li>• -handlePaused – Includes tasks which are paused from replication</li> </ul>
STATSAT	Uploads the statistics from the Satellite vault to the Base vault.
SYNCH <i>scope</i> [ <i>object-id</i> ] [replmode=ACTIVE BAV]	<p>Submits a replication event for the specified object.</p> <ul style="list-style-type: none"> <li>• <i>scope</i> – VAULT, CUSTOMER, LOCATION, COMPUTER, TASK, ACCOUNT, or USER</li> <li>• <i>object-id</i> – The object's database ID, not required for VAULT</li> <li>• <i>replmode</i> – ACTIVE or BAV. The default is both.</li> </ul>
SYNCHSAT <i>satelliteID</i>	Submits a replication event for the specified Satellite vault, where <i>satelliteID</i> is the ID of the Satellite vault to synchronize.
CLEARREPLBIT <i>taskID safesets</i>	<p>Removes "replicated" bit from safesets so the next satellite replication replicates this safeset. Where:</p> <ul style="list-style-type: none"> <li>• <i>taskID</i> – a DB identifier for a particular task or "*" for all tasks.</li> <li>• <i>safesets</i> – a comma-separated list of safesets to process or "*" for all safesets.</li> </ul>

Subcommand	Description
LATE scope [object-id] [-days=X] [-handlePaused] [emailto=]	<p>Produces a report for tasks that have fallen a specified number of days behind in replication. Where:</p> <ul style="list-style-type: none"> <li>• <i>scope</i> – VAULT, CUSTOMER, LOCATION, COMPUTER, TASK</li> <li>• <i>object-id</i> – The object's database ID, not required for VAULT.</li> <li>• <i>-days=X</i> – The number of days behind in replication. The default is 5 days.</li> <li>• <i>-handlePaused</i> – Includes tasks which are paused from replication</li> </ul> <p>To run the Replication Lag Report for 1:1 replication, run the REPLVAULT LATE command.</p>

### Examples

The following command, run on a Base vault, produces a Replication Lag Report for task 3 on a Satellite vault and includes tasks where replication is paused:

```
replvault setsatreport 8d057ebc-3644-4fba-9457-1815c9ec0ce7 task 3
-handlepaused -late
```

The following command produces a late report for an entire vault, includes tasks that are 10 days behind in replication, including paused tasks, and sends the report to one email address:

```
replvault late vault -days=10 -handlepaused emailto=admin@carbonite.com
```

## 14.5 secondaryop

The secondaryop command performs operations in secondary storage pools.

The secondaryop command can:

- check physical and logical data integrity
- delete safesets
- deduplicate data
- optimize pools
- attach, close and detach pools

### Syntax

```
secondaryop subcommand scopeParameter
```

## Subcommands

Subcommand	Description
checkcrc	Validates the physical integrity of files in secondary storage.
close [offline] [new_group= storageGroupName]	<p>Closes a secondary storage pool so that it is read-only and cannot be changed or updated.</p> <p>To mark safesets as offline after closing the pool, include the offline parameter.</p> <p>To specify a secondary storage group to use for a task after closing the previous one, include the new_group=storageGroupName parameter. You can only include the new_group=storageGroupName parameter if the command also includes the taskid=id scope parameter.</p> <p><i>Important:</i> If you close a secondary storage pool, the affected tasks start with a new pool the next time that migration occurs. This results in a full seed of data in the new secondary storage location, and could result in a significant increase in the amount of disk space required for secondary storage.</p>
dedup	Deduplicates data in a secondary storage pool.
delete safeset={id last *}	Deletes one or more safesets from a secondary storage pool: the safeset with the specified id, the last committed safeset, or all safesets.
expire	<p>Expires safesets in secondary storage according to retention settings. This command is similar to the vvmigrat online command for safesets in primary storage, but applies retention settings across safesets in both primary and secondary storage and only expires safesets from secondary storage.</p> <p><i>Warning:</i> Running secondaryop expire followed by secondaryop optimize can result in the deletion of all safesets in secondary storage. For example, if a task has one backup in primary storage and nine in secondary storage, and the retention settings are Online safesets=1, Online days=0 and Archive days=0, running the secondary op expire command marks all safesets in secondary storage for deletion.</p>
mount storage_location=path [reattach_deleted]	<p>Attaches a secondary storage pool. A secondary storage pool can only be attached for its associated task, and is then read-only.</p> <p>Use the storage_location=path parameter to specify the location of the secondary storage pool to attach. You cannot include other scope parameters with this subcommand.</p> <p>To re-attach a deleted safeset to the primary pool, include the reattach_deleted parameter.</p>

Subcommand	Description
mountall storage_location= <i>path</i>	Attaches all secondary storage pools under a specified directory.  Use the storage_location= <i>path</i> scope parameter to specify the directory with secondary storage pool locations to attach. You cannot include other scope parameters with this subcommand.
Optimize	Optimizes a secondary storage pool.
optimize dedup	Optimizes and deduplicates data in a secondary storage pool.
ssiverone safeset={ <i>id</i>  last} [/local]	Checks a safeset in a secondary storage pool for logical corruption: either the safeset with the specified <i>id</i> , or the last committed safeset. This verification process is more reliable but slower than the verifyidxone process.  To get information from the database for all tasks and skip detached pools, include the /local parameter. You can only include the /local parameter if the command also includes the storage_location= <i>path</i> scope parameter.  <i>Note:</i> The system assigns an identification number ( <i>id</i> ) to each safeset. To determine the <i>id</i> for a safeset, you can view its properties in the Director UI.
ssiverall [/local]	Checks all safesets in a secondary storage pool for logical corruption. This verification process is more reliable but slower than the verifyidxall process.  To get information from the database for all tasks and skip detached pools, include the /local parameter. You can only include the /local parameter if the command also includes the storage_location= <i>path</i> scope parameter.
Syncfix	Checks for and fixes discrepancies between main and pool synchronization information in a secondary storage pool.
Umount	Detaches a secondary storage pool.  Use the storage_location= <i>path</i> parameter to specify the location of the secondary storage pool to detach. You cannot include other scope parameters with this subcommand.
unmountall storage_location= <i>path</i>	Detaches all secondary storage pools under the specified directory.  Use the storage_location= <i>path</i> scope parameter to specify the directory with secondary storage pools to detach. You cannot include other scope parameters with this subcommand.



Subcommand	Description
verifyidxone safeset={ <i>id</i>  last}	<p>Checks a safeset in a secondary storage pool for logical corruption: either the safeset with the specified <i>id</i>, or the last committed safeset. This verification process is less reliable but faster than the <code>ssiverone</code> process. Unlike <code>ssiverone</code>, this process does not read the pool files.</p> <p>To get information from the database for all tasks and skip detached pools, include the <code>/local</code> parameter. You can only include the <code>/local</code> parameter if the command also includes the <code>storage_location=</code><i>path</i> scope parameter.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each safeset. To determine the <i>id</i> for a safeset, you can view its properties in the Director UI.</p>
verifyidxall [/local]	<p>Checks all online safesets in a secondary storage pool for logical corruption. This verification process is less reliable but faster than the <code>ssiverall</code> process. Unlike <code>ssiverall</code>, this process does not need pool files.</p> <p>To get information from the database for all tasks and skip detached pools, include the <code>/local</code> parameter. You can only include the <code>/local</code> parameter if the command also includes the <code>storage_location=</code><i>path</i> scope parameter.</p>

## Scope Parameters

Scope Parameter	Description
taskid= <i>id</i>	<p>Performs the operation for the secondary storage pool for the task with the specified <i>id</i>. To perform the operation in pools for all tasks, enter <code>*</code> as the <i>id</i>. This parameter can only be used for secondary storage pools with tasks still on the local vault.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each task. To determine the <i>id</i> for a task, you can run a report for the task or view its properties in the Director UI.</p> <p><i>Note:</i> The <code>taskid</code> parameter is not used with the <code>mount</code>, <code>mountall</code>, <code>umount</code> or <code>umountall</code> subcommands.</p>
storage_location= <i>path</i> [user=[ <i>domain</i> ]\ <i>username</i> password= <i>password</i>   /useDefaultCredentials]	<p>Performs the operation for the secondary storage pool in the specified <i>path</i>. The <i>path</i> can be a local or a UNC path.</p> <p>If the location is a network share, <code>user=[<i>domain</i>]\<i>username</i></code> <code>password=<i>password</i></code> specify the domain, username and password for connecting to the network share. Alternatively, to connect to the network share using default credentials, include the <code>/useDefaultCredentials</code> parameter.</p>

Scope Parameter	Description
<code>storage_group=storageGroupName</code>	For use with the <code>checkcrc</code> , <code>close</code> , <code>ssiverone</code> , <code>ssiverall</code> , <code>verifyidxone</code> and <code>verifyidxall</code> subcommands only.  Performs the operation for the secondary storage group with the specified <code>storageGroupName</code> .

## Examples

The following command performs a CRC in the "secondary" storage group:

```
secondaryop checkcrc storage_group=secondary
```

The following command closes secondary storage for task 2 and marks safesets in the pool as offline:

```
secondaryop close taskid=2 offline
```

The following command deduplicates data in the `\\computer1\c$\secondary` storage pool location:

```
secondaryop dedup storage_location=\\computer1\c$\secondary
```

The following command deletes the last safeset from the `e:\secondary` storage pool location:

```
secondaryop delete safeset=last storage_location=e:\secondary
```

The following command deletes expired safesets for task 3 from secondary storage:

```
secondaryop expire taskid=3
```

The following command attaches the secondary storage pool in `c:\secondary3`:

```
secondaryop mount storage_location=c:\secondary3
```

The following command attaches all secondary storage pools under the `e:\secondarystorage` directory:

```
secondaryop mountall storage_location=e:\secondarystorage
```

The following command optimizes data in secondary storage for task 4:

```
secondaryop optimize taskid=4
```

The following command optimizes and deduplicates data in the secondary storage pool in `f:\secondary`:

```
secondaryop optimize dedup storage_location=f:\secondary
```

The following command verifies the last safeset in the "secondary2" storage group:

```
secondaryop ssiverone safeset=last storage_group=secondary2
```

The following command verifies all safesets in the `\\computer1\c$\secondary` storage pool location:

```
secondaryop ssiverall storage_location=\\computer1\c$\secondary
```

The following command synchronizes main and pool synchronization information for task 1 in secondary storage:

```
secondaryop syncfix taskid=1
```

The following command detaches the secondary storage pool in the `f:\secondary` storage directory:

```
secondaryop umount storage_location=f:\secondary
```

The following command detaches all secondary storage pools under the d:\secondarystorage directory:

```
secondaryop unmountall storage_location=d:\secondarystorage
```

The following command verifies safeset 6 for task 5 in secondary storage:

```
secondaryop verifyidxone safeset=6 taskid=5
```

The following command verifies safesets in the f:\secondary storage pool location:

```
secondaryop verifyidxall storage_location=f:\secondary
```

## 14.6 vaultop

The vaultop command executable is available in the Utils directory where Director is installed.

The vaultop command includes subcommands for:

- Remotely starting or stopping vault services, or shutting down vault services and spawned vault processes. See [vaultop service and node subcommands](#).

*Note:* You cannot manage the Reporting service using the vaultop command.

- Managing vault certificates and enabling the agent-vault certificate pinning feature. See [vaultop certificate subcommands](#).

- Configuring email notifications and settings. See [vaultop email subcommands](#).

*Note:* You can also configure email settings using the Director UI. See [Configure email notifications](#).

## 14.6.1 vaultop service and node subcommands

### Syntax

```
vaultop subcommand [service1...servicen] [timeout]
```

### Subcommands

Subcommand	Description
start_services	Starts one or more specified vault services.
stop_services <i>timeout</i>	Stops one or more specified vault services. <i>timeout</i> is the number of seconds for each vault service to stop gracefully. The default value for <i>timeout</i> is 30 seconds.  IMPORTANT: Stopping vault services using this command could result in data corruption. Instead, use the vaultop set_node_state rampdown command to ramp down a vault.
restart_services	Stops then starts one or more specified vault services.  IMPORTANT: Restarting services using this command could result in data corruption. Instead, use the vaultop set_node_state rampdown command to ramp down a vault. You can then start the services.
shutdown_vault <i>timeout</i>	Stops all vault services and all spawned vault processes. <i>timeout</i> is the time for each vault service and process to stop gracefully. After the vault services and processes are stopped, all active vault processes are terminated. The default value for <i>timeout</i> is 30 seconds.  <i>Note:</i> You cannot specify <i>services</i> with the shutdown_vault subcommand.  IMPORTANT: Stopping vault services using this command could result in data corruption. Instead, use the vaultop set_node_state rampdown command to ramp down a vault.

Subcommand	Description
get_node_state	<p>Gets the operational state of the vault. A vault can have any of the following states:</p> <ul style="list-style-type: none"> <li>• Online. The vault is available to perform Director operations.</li> <li>• Rampdown. The vault is finishing the current operations before it goes offline. No new operations will start on the node before it goes offline.</li> <li>• Offline. The vault is not available to perform Director operations and will not accept requests.</li> </ul>
kill_node_processes	Stops all vault processes on the vault node.
set_node_state {ONLINE RAMPDOWN OFFLINE} [/t <i>timeout</i> ]	<p>Sets the operational state of the vault to one of the following:</p> <ul style="list-style-type: none"> <li>• ONLINE - The vault is available to perform Director operations.</li> <li>• RAMPDOWN - The vault finishes the current operations and then goes offline. Note that this can take some time. Some operations (replication, in particular) can take hours to finish under some circumstances.</li> <li>• OFFLINE - The vault is not available to perform Director operations and will not accept requests.</li> </ul> <p>For all state transitions, the <i>/t timeout</i> parameter specifies the number of seconds that the command waits for the system to transition to the requested state. If the system has not transitioned to the desired state after the specified number of seconds, the command returns a timeout failure. The transition to the desired state proceeds after the command times out, until the desired state is reached. The default value for <i>timeout</i> is 30 seconds.</p> <p><i>Note:</i> For RAMPDOWN, the <i>timeout</i> does not specify the amount of time after which processes are terminated.</p>

Subcommand	Description
update_node_in_cluster [externalAddress:<externalWorkerIPaddress>] [internalAddress:<internalWorkerIPaddress>]	<p>Updates the external and/or internal network address of a vault.</p> <p>The <i>externalWorkerIPaddress</i> is a public IP address for connections from outside the firewall (e.g., agents, Satellite vaults, source vaults for replication). Specify an externally-available IP address or fully-qualified domain name (FQDN).</p> <p>The <i>internalWorkerIPaddress</i> is a private IP address for internal vault connections (e.g., with the Director UI).</p> <p><i>Note:</i> When the vault is installed on a virtual machine in Microsoft Azure, an externally-accessible IP address or FQDN must be specified for both the external and internal addresses. See the <i>Director Installation Guide</i>.</p>

## Services

*service1...servicen* specify one or more vault services to start or stop. The following table lists possible values.

Service	Description
* (asterisk)	All vault services except the Reporting service. This is the default value.
AdminService	Interface service for the Director UI. This service translates between commands made in the UI and the command line. It provides feedback to the UI on jobs that are running. This service produces Admin Service logs in the Director\logs directory.
ReplService	Service that handles replication on the vault.
VVListen	Service that listens to agent-initiated traffic and spawns a server process. It validates the traffic based on the signature of the packets it receives. It also listens on the ports specified for the vault. If you add, change, or delete a port, you must restart the VVListen service to allow the service to listen to the new port. It produces a SPAWN log for each server process it starts.
VVQmanager	Service that manages the launch sequence of non-backup, restore, and synch processes on the vault. It enforces the Maximum Number of Jobs setting. It produces QManager logs in the Director\logs directory.

Service	Description
VVSvrDae	Scheduler service that enforces the Director schedule settings. It produces VVSvrDae logs in the Director\logs directory.
VVSvrMonitor	Service that monitors the health of other vault services and manages node state transitions (e.g., from Online to Rampdown to Offline).

## Timeout

For the `stop_services` and `shutdown_vault` subcommands, you can specify a timeout value. The default value is 30 seconds.

For the `stop_services` subcommand, *timeout* specifies the number of seconds given to each vault service to stop gracefully.

For the `shutdown_vault` subcommand, *timeout* specifies the number of seconds given to all vault services and process to stop gracefully. After that, all active vault processes are terminated.

## Examples

The following command starts the ReplService and VVListen services:

```
vaultop start_services ReplService VVListen
```

The following command stops all vault services with a 120-second timeout:

```
vaultop stop_services 120 *
```

The following command restarts the ReplService and VVListen services:

```
vaultop restart_services ReplService VVListen
```

The following command shuts down the vault in 120 seconds:

```
vaultop shutdown_vault 120
```

### 14.6.2 vaultop certificate subcommands

When you install a Director 8.70 vault or upgrade a vault from a Director 8.56 or earlier, a self-signed TLS certificate that is valid for 10 years is generated for the vault.

The self-signed certificate that was generated for Director 8.6x vaults was valid for one year. If you upgrade a Director 8.6x vault that has its original generated self-signed certificate, the certificate is regenerated with a 10-year validity period. When the certificate is regenerated, the certificate's public and private keys are preserved so that, if agent-certificate pinning is enabled on the vault, backups continue to function without interruption.

The vault certificate is stored in the Local Computer certificate store in `\Carbonite Server Backup\Certificates`.

Using `vaultop certificate` subcommands, you can:

- Import a new vault certificate. You can replace a vault certificate with a self-signed certificate or a certificate from an enterprise or commercial Certificate Authority (CA).

- Export a vault certificate. Exporting a vault certificate can be useful if you want to import the certificate into another vault.
- Enable agent-vault certificate pinning. When this security feature is enabled in a vault, when an agent that supports certificate pinning tries to connect to the vault (e.g., to run a backup or restore), it checks whether the public key of the vault's TLS certificate is the same as when the agent previously connected to the vault. If the public key of the vault certificate is different, the agent reports a certificate failure and will not connect to the vault until the certificate failure is resolved.

If agent-vault certificate pinning was enabled in a Director 8.6x vault, the feature remains enabled after the vault is upgraded to version 8.7.

- List or delete pinned and recorded certificates on a vault. When a Director 8.7 source vault first connects to a Director 8.7 target vault to replicate data, copy or clone data, or run a replication report, the source vault pins or records the target vault's certificate. When the source vault tries to connect to the target vault again, it checks whether the target vault certificate has changed. If the target vault certificate has changed to a self-signed certificate, the source vault will not connect to the target vault and "certificate verify fail" messages appear in the Replication Service log. The source vault will not connect to the target vault again unless you delete the pinned or recorded certificate from the source vault. For more information, see [Certificate verification and pinning for vault-to-vault communications](#).

**IMPORTANT:** Do not delete a pinned or recorded vault certificate unless you are sure that there is no security risk. Please contact your IT security staff or service provider to determine whether the certificate change was expected or whether further investigation is required.

## Syntax

`vaultop subcommand`



## Subcommands

Subcommand	Description
<p>import_certificate certPathAndFilename certPassword</p>	<p>Imports a vault certificate. The certificate can be a self-signed certificate or a certificate from an enterprise or commercial Certificate Authority (CA) and can be a wildcard certificate. The certificate must be in .pfx format.</p> <p><i>certPathAndFilename</i> is the full path and file name of the certificate that you are importing. Enclose the certificate path and filename in quotation marks if it includes spaces. <i>certPassword</i> is the password of the .pfx file that you are importing. The certificate is imported into the Local computer certificate store in \Carbonite Server Backup\Certificates.</p> <p>IMPORTANT: If you import a new vault certificate, additional actions might be required to allow backups and vault-to-vault communications (e.g., replication) to continue:</p> <ul style="list-style-type: none"> <li>• If you import a vault certificate with a different public key than the previous vault certificate, and agent-vault certificate pinning is enabled in the vault, a Portal user must resolve agent certificate failures before backups can continue. For more information, see the <a href="#">vaultop certificate pinning subcommand</a>.</li> <li>• If you replace a Director 8.7 target vault certificate with a new self-signed certificate, Director 8.7 source vaults (e.g., Satellite vaults or Active vaults) will not connect to the target vault again unless you delete the pinned or recorded certificate from the source vault. For more information, see the <a href="#">vaultop delete pinned certificate subcommand</a>.</li> <li>• If you replace a Director 8.7 target vault certificate with a CA-signed certificate and the signing authority is a Microsoft trusted certificate authority, the signing certificate is automatically downloaded to the Trusted Root Certification Store on the Director 8.7 source vault server. In rare cases, you might have to manually download and import the root signing certificate on the source vault server.</li> </ul> <p>Beginning in Director 8.70, when you import a new vault certificate using the vaultop import_certificate command, the certificate being imported is verified. If the root certificate is not found on the vault server, the import fails. If another issue is found with the certificate (e.g., the certificate has expired), a warning message is logged but the certificate is imported.</p> <p><i>Note:</i> Director services must be restarted after any change to the vault certificate or vault certificate chain (e.g., a new signing certificate for a CA-signed certificate). The vaultop import_certificate command automatically restarts Director services.</p>

Subcommand	Description
export_certificate <i>certPathAndFilename</i> <i>certPassword</i>	Exports the vault certificate in .pfx format. <i>certPathAndFilename</i> is the full path and file name of the exported certificate. <i>certPassword</i> is the password for the exported .pfx file.
certificate_pinning	<p>Enables agent-vault certificate pinning. When this feature is enabled in a vault, when an agent that supports certificate pinning tries to connect to the vault (e.g., to run a backup or restore), it checks whether the public key of the vault's TLS certificate is the same as when the agent previously connected to the vault. If the public key of the vault certificate is different, the agent reports a certificate failure and will not connect to the vault. For more information, see the <i>Director Installation Guide</i>.</p> <p>If agent-vault certificate pinning was enabled in a Director 8.6x vault, the feature remains enabled after the vault is upgraded to version 8.7.</p> <p><b>IMPORTANT:</b> Do not enable agent-vault certificate pinning until the intended TLS certificate for the vault is installed. If you enable this feature and then import a certificate with a different public key, agents that support certificate pinning will not connect to the vault until a Portal user resolves the certificate failures.</p> <p><b>IMPORTANT:</b> Agent-vault certificate pinning cannot be turned off in a vault after it is enabled.</p> <p><i>Note:</i> You do not have to manually enable certificate verification and pinning for vault-to-vault communications. Certificates are always verified when the source and target vaults are Director version 8.7.</p>
list_certificates [verbose]	<p>Lists pinned and recorded certificates on a vault. A Director 8.7 source vault pins or records the target vault certificate when it first connects to a Director 8.7 target vault to replicate data, copy or clone data, or run a replication report. When the source vault tries to connect to the target vault again, it checks whether the target vault certificate is the same as the pinned or recorded certificate. For more information, see <a href="#">Certificate verification and pinning for vault-to-vault communications</a>.</p> <p>For each pinned self-signed certificate, this command returns the target vault address, certificate thumbprint and certificate subject name. To also return each self-signed certificate in PEM format, include the verbose parameter.</p> <p>For each recorded CA-signed certificate, the command returns the target vault address and indicates that the vault certificate is not self-signed.</p>

Subcommand	Description
<code>delete_pinned_certificate</code> <i>serverAddress</i>	<p>Deletes a pinned or recorded vault certificate from a source vault. <i>serverAddress</i> is the IP address or hostname of the target vault with the certificate that was pinned or recorded on the source vault. If a target vault certificate changes to a self-signed certificate, the source vault will not connect to the target vault unless you delete the pinned certificate using this command. For more information, see <a href="#">Certificate verification and pinning for vault-to-vault communications</a>.</p> <p><b>IMPORTANT:</b> Do not delete a pinned or recorded vault certificate unless you are sure that the certificate change was expected and there is no security risk. Please contact your IT security staff or service provider to determine whether a certificate change was expected or whether further investigation is required.</p>

## Examples

The following command imports a C:\Certificate\cert.pfx file with the password “password1” into a vault:

```
vaultop import_certificate C:\Certificate\cert.pfx password1
```

The following command exports the current vault certificate to a cert.pfx file in C:\Exported Cert with the password “password2”:

```
vaultop export_certificate "C:\Exported Cert\cert.pfx" password2
```

The following command enables the certificate pinning feature:

```
vaultop certificate_pinning
```

**IMPORTANT:** Certificate pinning cannot be turned off in a vault after it is enabled.

The following command lists pinned and recorded certificates on a source vault, including each self-signed certificate in PEM format:

```
vaultop list_certificates verbose
```

The following command deletes the pinned certificate on a source vault for a target vault with IP address 198.51.100.23:

```
vaultop delete_pinned_certificate 198.51.100.23
```

### 14.6.3 vaultop email subcommands

Using vaultop email subcommands, you can:

- Configure email notifications, including the SMTP server, email addresses and events for sending notifications. See [vaultop configure\\_send\\_email subcommand](#).
- Send a test email from Director. See [vaultop send\\_test\\_email subcommand](#).

You can also configure and test email settings using the Director UI. For more information, see [Configure email notifications](#).

### 14.6.3.1 vaultop configure\_send\_email subcommand

The vaultop configure\_send\_email subcommand configures email notifications, including the SMTP server, email addresses and events for sending notifications. You can configure initial notification settings or change existing settings using this command.

#### Syntax

```

vaultop configure_send_email server=server [port=port] [transport-
security={None|StartTLS|SSL/TLS}] [authentication=authentication] [user=username
password=password] [test-timeout=testTimeout] [use-admin-service={true|false}]
from="fromAddress" [recipients="emailAddresses"] [notify-on-success={0|1}] [notify-
on-failure={0|1}] [notify-on-device={0|1}] [notify-on-replication-
failure={0|1}] [test-connection={true|false}] [timeout=timeout]

```

#### Parameters

Parameter	Description
server= <i>server</i>	Specifies the SMTP server to use for sending email notifications. <i>server</i> is the fully-qualified domain name (FQDN) or IP address of the SMTP server. The SMTP server can be a private email server or a cloud email service provider (e.g., Amazon SES, Google Gmail email service or Microsoft 365). If you specify a cloud email service provider, please see the provider's documentation for requirements.
port= <i>port</i>	Specifies the SMTP server port number. If you do not include this parameter or specify 0 (zero) as the value, the port is set to the appropriate port number for the specified encryption type: port 25 for no encryption, port 587 for StartTLS, or port 465 for SSL/TLS.
transport-security={None StartTLS SSL/TLS}	Specifies the email encryption type: None, StartTLS or SSL/TLS. If you do not include this parameter, the email encryption type is SSL/TLS. We do not recommend "None" for the encryption type. Unencrypted emails can be read if they are intercepted.
authentication= <i>authentication</i>	Specifies the SMTP authentication method: Auto, Anonymous, CRAM-MD5, CRAM-SHA1, LOGIN, PLAIN or NTLM. If you do not include this parameter, the authentication method is Auto.
user= <i>username</i> password= <i>password</i>	Specifies the <i>username</i> and <i>password</i> for authenticating with the SMTP server, if required. If you specify a <i>username</i> and <i>password</i> but the authentication method is Anonymous, the credentials are stored but are not used for SMTP authentication.

Parameter	Description
<code>test-timeout=<i>testTimeout</i></code>	<p>Specifies the number of seconds to wait when testing the SMTP server connection before running the command. If a connection is not established during this time, the command fails. If you do not include this parameter, the <i>testTimeout</i> value is 30 seconds.</p> <p><i>Note:</i> If the test-connection parameter value is false, the SMTP server connection is not tested before running the command.</p>
<code>use-admin-service={true   false}</code>	<p>Specifies whether the command is processed by the Director Admin service. To run the command outside of the Admin service (e.g., if the Admin service is down or does not trust the email server certificate), specify <code>use-admin-service=false</code>. If you do not include this parameter, the command is processed by the Admin service.</p>
<code>from=<i>fromAddress</i></code>	<p>Specifies the email address that will appear as the From address in notification emails.</p> <p>You cannot specify a display name for the From address in this command, but you can specify a display name using the Director UI. For more information, see <a href="#">Configure email notifications</a>.</p>
<code>recipients="<i>emailAddresses</i>"</code>	<p>Specifies the email address or addresses that will receive notifications. Enclose multiple addresses in quotation marks and use a comma as a separator between email addresses. If you do not include this parameter, the email notification recipient list does not change.</p>
<code>notify-on-success={0   1}</code>	<p>Specifies whether notifications are sent when jobs finish successfully. To send notifications when jobs finish successfully, specify 1. To not send notifications when jobs finish successfully, specify 0. If you do not include this parameter, the notify-on-success setting does not change.</p>
<code>notify-on-failure={0   1}</code>	<p>Specifies whether notifications are sent when jobs fail. To send notifications when jobs fail, specify 1. To not send notifications when jobs fail, specify 0. If you do not include this parameter, the notify-on-failure setting does not change.</p>
<code>notify-on-device={0   1}</code>	<p>Specifies whether notifications are sent when UNC connections are inaccessible. To send notifications when UNC connections are inaccessible, specify 1. To not send notifications when UNC connections are inaccessible, specify 0. If you do not include this parameter, the notify-on-device setting does not change.</p>
<code>notify-on-replication-failure={0   1}</code>	<p>Specifies whether notifications are sent when replication sessions fail. To send notifications when replication sessions fail, specify 1. To not send notifications when replication sessions fail, specify 0. If you do not include this parameter, the notify-on-replication setting does not change.</p>

Parameter	Description
<code>test-connection={true false}</code>	Specifies whether to test the connection to the SMTP server before running the command. To test the connection to the SMTP server, specify true. To not test the connection to the SMTP server, specify false. If you do not include this parameter, the test-connection value is true.
<code>timeout=<i>timeout</i></code>	Specifies the number of seconds for Director to wait for a connection to the SMTP server when sending a notification. If a connection is not established during this time, the email notification fails. If you do not include this parameter, the <i>timeout</i> value is 600 seconds.

### Example

The following command specifies an SMTP server for sending email notifications, a From address (`notification@company.com`) for the emails, and two email addresses (`user1@company.com` and `user2@company.com`) that will receive notifications when jobs or replication sessions fail:

```
vaultop configure_send_email server=smtp.company.com transport-security=SSL/TLS
user=user@company.com password=StrongPassword from=notification@company.com
recipients="user1@company.com,user2@company.com" notify-on-failure=1 notify-on-
replication-failure=1
```

#### 14.6.3.2 vaultop send\_test\_email subcommand

The `vaultop send_test_email` subcommand sends a test email from Director using a specified SMTP server.

### Syntax

```
vaultop send_test_email server=server [port=port] [transport-
security={None|StartTLS|SSL/TLS}] [authentication=authentication] [user=username
password=password] [test-timeout=seconds] [use-admin-service={true|false}]
from="fromAddress" recipients="emailAddresses" "subject" "body"
```

Parameter	Description
<code>server=<i>server</i></code>	Specifies the SMTP server to use for sending the test email. <i>server</i> is the fully-qualified domain name (FQDN) or IP address of the SMTP server. The SMTP server can be a private email server or a cloud email service provider (e.g., Amazon SES, Google Gmail email service or Microsoft 365). If you specify a cloud email service provider, please see the provider's documentation for requirements.
<code>port=<i>port</i></code>	Specifies the SMTP server port number. If you do not include this parameter or specify 0 (zero) as the value, the port is set to the appropriate port number for the specified encryption type: port 25 for no encryption, port 587 for StartTLS and port 465 for SSL/TLS.

Parameter	Description
transport-security={None StartTLS SSL/TLS}	Specifies the email encryption type: None, StartTLS or SSL/TLS. If you do not include this parameter, the email encryption type is SSL/TLS.  We do not recommend “None” for the encryption type. Unencrypted emails can be read if they are intercepted.
authentication= <i>authentication</i>	Specifies the SMTP authentication method: Auto, Anonymous, CRAM-MD5, CRAM-SHA1, LOGIN, PLAIN or NTLM. If you do not include this parameter, the authentication method is Auto.
user= <i>username</i> password= <i>password</i>	Specifies the <i>username</i> and <i>password</i> for authenticating with the SMTP server, if required.
test-timeout= <i>testTimeout</i>	Specifies the number of seconds to wait when testing the SMTP server connection before sending a test email. If a connection is not established during this time, the command fails. If you do not include this parameter, the <i>testTimeout</i> value is 30 seconds.
use-admin-service={true false}	Specifies whether the command is processed by the Director Admin service. To run the command outside of the Admin service (e.g., if the service is down or does not trust the email server certificate), specify use-admin-service=false. If you do not include this parameter, the command is processed by the Admin service.
from= <i>fromAddress</i>	Specifies the email address that will appear as the From address in the test email.  You cannot specify a display name for the From address in this command, but you can specify a display name using the Director UI. For more information, see <a href="#">Configure email notifications</a> .
recipients=" <i>emailAddresses</i> "	Specifies the email address or addresses that will receive the test email. Enclose multiple addresses in quotation marks and use a comma as a separator between email addresses.
<i>Subject</i>	Specifies a subject line for the test email. Enclose the subject line in quotation marks.
<i>Body</i>	Specifies the message body for the test email. Enclose the message body in quotation marks.

## Example

The following command specifies an SMTP server for sending a test email, a From address (notification@company.com) for the email, and an email address (user1@company.com) that will receive the test email:

```
vaultop send_test_email server=smtp.company.com transport-security=SSL/TLS  
user=user@company.com password=StrongPassword from=notification@company.com  
recipients=user1@company.com
```

## 14.7 Vault Settings

The VaultSettings script gets and populates vault settings in the Director database.

This PowerShell script is available in the Scripts directory where Director is installed.

### Syntax

```
.\VaultSettings.ps1 [ get | set ] vaultSetting [value]
```

Where:

- *vaultSetting* is the vault setting that you want to get or set.
- *value* is value for the vault setting. This parameter is only used with the “set” command.

The following sections describe vault settings that you can change using the VaultSettings script:

- [Enable the Rapid VM Restore feature on a vault](#)
- [Change the date when an upgraded vault stops accepting replication from vaults that are susceptible to a security issue](#)
- [Change the maximum email size](#)
- [Enable or disable block size cache files](#)

For more information about the VaultSettings script, please contact Support.

### 14.7.1 Enable the Rapid VM Restore feature on a vault

Beginning with vSphere Recovery Agent 8.80 and Hyper-V Agent 9.00, you can restore a VM within minutes using the Rapid VM Restore feature. The backup must be saved in a local version 8.50 or later vault that has the Rapid VM Restore feature enabled. For additional requirements, see the [Server Backup online help](#), *vSphere Recovery Agent User Guide* or *Hyper-V Agent User Guide*.

The Rapid VM Restore feature is enabled by default on Satellite vaults (e.g., on appliances). On Base vaults that are installed locally, you must enable the Rapid VM Restore feature using the following procedure.

To enable the Rapid VM Restore feature on a vault:

1. On the server where the vault is installed, open a Powershell window as administrator, and navigate to the Scripts subfolder in the Director installation directory.



2. Run the following command:

```
.\VaultSettings.ps1 set IsRVMRAAllowed 1
```

### 14.7.2 Change the date when an upgraded vault stops accepting replication from vaults that are susceptible to a security issue

Because Director 8.60, 8.51, 8.50, 8.4x and 8.30 vaults were susceptible to a critical security issue, replication from these vault versions was disabled in Director 8.62 or 8.61 vaults on May 31, 2022. However, the cutoff date could be changed so that replication from these vault versions continued after May 31, 2022.

If you upgrade a Director 8.62 or 8.61 vault that receives replicated data from a Director 8.60, 8.51, 8.50, 8.4x or 8.30 vault, you can change the cutoff date again so that vaults with the security issue continue to replicate data to the target vault.

To change the date when an upgraded vault stops accepting replication from vaults that are susceptible to a security issue:

1. On the target vault server, open a PowerShell window as administrator. In the PowerShell window, navigate to the Scripts subfolder in the Director installation directory.
2. Run the following command:

```
.\VaultSettings.ps1 set LegacyReplicationCutoffDate yyyymmdd
```

Where *yyyymmdd* is the date when replication from vaults that are susceptible to the security issue will be automatically disabled in the vault. You cannot specify a date in the past.

When you run the command, the date will be applied but the following message will appear: *To protect your environment, it is recommended that this value is set to an earlier date.*

For example, to change the date for disabling replication from vaults that are susceptible to the security issue to September 1, 2023, run the following command:

```
.\VaultSettings.ps1 set LegacyReplicationCutoffDate 20230901
```

### 14.7.3 Change the maximum email size

Beginning in Director 8.62, you can change the maximum size for Director emails. This size limit applies to email notifications and emailed reports.

The default maximum email size is 20 MB. In Director 8.61 and earlier versions, the default maximum email size was 2 MB and could not be changed.

If a Director email is larger than the maximum size, a “THIS MESSAGE WAS TRUNCATED at *maximumSize* MB” message appears at the end of the email and the remaining lines are removed. If an emailed report is truncated, you can view the complete report in the log file. You can also increase the maximum email size so that reports will not be truncated in the future.

If you set the maximum email size to your SMTP server’s message size limit, Director emails will be delivered instead of bouncing back to the sender.

To change the maximum email size:

1. On the vault server, open a PowerShell window as administrator. In the PowerShell window, navigate to the Scripts subfolder in the Director installation directory.
2. Run the following command:

```
.\VaultSettings.ps1 set MaxEmailSizeInMBytes maximumSizeInMB
```

Where *maximumSizeInMB* is the maximum email size in megabytes. The maximum email size can be set from 1 MB to 3072 MB (i.e., 3 GB). If you specify 0 (zero) or a value larger than 3072 MB as the *maximumSizeInMB*, the maximum size will be set to 3 GB.

For example, to change the maximum email size to 50 MB, run the following command:

```
.\VaultSettings.ps1 set MaxEmailSizeInMBytes 50
```

#### 14.7.4 Enable or disable block size cache files

During optimization, migration and replication, pool size calculations are performed on a regular basis. Depending on various factors, these calculations can take a long time.

Beginning in Director 8.70, block size data that is used in pool size calculations can be held in cache files at the task level. Block size cache files (file extension: .pfblocksizocache) can significantly reduce the time required for pool size calculations, especially on vaults that have storage with poor read performance or high latency.

IMPORTANT:

- Extra storage is used by block size cache files.
- We recommend enabling block size cache files on one vault and monitoring the vault performance before enabling these files on other vaults. If the files use too much storage, you can disable block size cache files on the vault.
- If you disable block size cache files, cache files will no longer be created but existing cache files will not be deleted automatically. You must manually delete files with the .pfblocksizocache extension.

To enable block size cache files:

1. On the server where the vault is installed, open a PowerShell window as administrator, and navigate to the Scripts subfolder in the Director installation directory.
2. Run the following command:

```
.\VaultSettings.ps1 set UsePfBlockSizeCacheFile 1
```

To disable block size cache files:

1. On the server where the vault is installed, open a PowerShell window as administrator, and navigate to the Scripts subfolder in the Director installation directory.
2. Run the following command:

```
.\VaultSettings.ps1 set UsePfBlockSizeCacheFile 0
```

## 14.8 vvanalyze

The vvanalyze command produces reports on tasks, pools and vaults. You must run the vvanalyze command locally on the vault you want to report on.

### Syntax

```
vvanalyze subcommand [parameters]
```

### Subcommands

Subcommand	Description
/lateststatus [/timeframe= <i>hours</i> ] [ <i>scopeParameter</i> ] [ <i>destination</i> ]	<p>Runs a Late Server Status report. The report lists tasks that were not backed up in a specified number of hours before the current time. For more information, see <a href="#">Create a Late Server Status Report</a>.</p> <p>To specify the number of hours before the current time for finding tasks that were not backed up, include the <code>/timeframe=<i>hours</i></code> parameter, where <i>hours</i> is the number of hours before the current time. If you do not include the <code>/timeframe=<i>hours</i></code> parameter, the report lists tasks that were not backed up in the past 48 hours.</p> <p>For <i>scopeParameter</i> and <i>destination</i> descriptions, see the <a href="#">Scope Parameters</a> and <a href="#">Destinations</a> tables.</p>
/originalsizeupdate	<p>Updates the Original size value in the database for one or more safesets. The Original size value is the amount of original data protected by a safeset, and is provided by the Agent.</p> <p><i>Note:</i> This subcommand is only used for safesets created with Director 6.31 or earlier. The Original size is automatically updated in the database for safesets created with Director 7.0 or later.</p>

Subcommand	Description
/missedbackups [ <i>missedBackupParameters</i> ] [ <i>scopeParameter</i> ][ <i>destination</i> ]	<p>Runs a Missed Backups report. The report lists backup tasks that were scheduled or expected to run, but for which a safeset does not exist. A task appears in the report even if the scheduled or expected backup ran successfully but the committed safeset was later deleted. For more information, see <a href="#">Create a Missed Backups Report</a>.</p> <p>For a list of <i>missedBackupParameters</i>, see the <a href="#">vvanalyze /missedbackups parameters</a> table. For <i>scopeParameter</i> and <i>destination</i> descriptions, see the <a href="#">Scope Parameters</a> and <a href="#">Destinations</a> tables.</p> <p><i>Note:</i> If the vault's schedule.cfg file has been modified manually, the Missed Backups report does not report tasks correctly.</p>
/storage [/bare] [ <i>scopeParameter</i> ][ <i>destination</i> ]	<p>Runs a Storage report. The report shows the number of safesets and amount of data in a vault for a single task, or for all tasks in a vault, organization/customer, location, or computer. For more information, see <a href="#">Create a Storage Report</a>.</p> <p>To produce a report in CSV format, include the /bare parameter. If you do not include the /bare parameter, report data appears in columns.</p> <p>For <i>scopeParameter</i> and <i>destination</i> descriptions, see the <a href="#">Scope Parameters</a> and <a href="#">Destinations</a> tables.</p>
/storageLocation [ <i>scopeParameter</i> ][ <i>destination</i> ]	<p>Runs a Storage Location report. The report lists all primary, secondary and archive storage locations for one task, or for each task in a vault, organization/customer, location, or computer. For more information, see <a href="#">Create a Storage Location report</a>.</p> <p>For <i>scopeParameter</i> and <i>destination</i> descriptions, see the <a href="#">Scope Parameters</a> and <a href="#">Destinations</a> tables.</p>

Subcommand	Description
<p>/storageextract /poolsize [<i>storageExtractParameters</i>] [<i>destination</i>]</p> <p>OR</p> <p>/storageextract /original [<i>storageExtractParameters</i>] [<i>destination</i>]</p>	<p>Reports the amount of data for each task, computer, location and customer in a vault, and provides the total for the vault as a whole. The report is saved in CSV and log files.</p> <p>If you include the /poolsize parameter, the report shows the pool size for each task, computer, location, customer and the vault as a whole.</p> <p>If you include the /original parameter, the report shows the original data size, or uncompressed protected data, for each task, computer, location, customer and the vault as a whole.</p> <p>The resulting CSV files are saved in the Director\bill directory, and are named <i>yyyymmdd_ONLINE.csv</i> and <i>yyyymmdd_OFFLINE.csv</i>, where <i>yyyymmdd</i> is the year, month and day when the report was generated. The CSV files include the customer name, location code, computer name, and total gigabytes (GB) for each computer. You can import the contents of the output file into a spreadsheet or other application to produce reports.</p> <p>The log file is saved in a StorageExtract log file in the Director\logs directory. When report information appears in the log file, it contains more detail than the CSV files, and includes pool size information for each task, with totals at the computer, location, customer, and vault levels.</p> <p>Numbers in the CSV files are in gigabytes, and numbers in the log file are in bytes, kilobytes, megabytes and gigabytes. Discrepancies between the CSV and log files can occur because of conversion and rounding.</p> <p>For a list of <i>storageExtractParameters</i>, see the <a href="#">vanalyze /storageExtract Parameters</a> table. For <i>destination</i> descriptions, see the <a href="#">Destinations</a> table.</p>
<p>/syncscan [<i>scopeParameter</i>] [<i>destination</i>]</p>	<p>Runs a Last Backup Status report. The report shows information about the last backup for one task, or for each task in a vault, organization/customer, location, or computer. For more information, see <a href="#">Create a Last Backup Status Report</a>.</p> <p>For <i>scopeParameter</i> and <i>destination</i> descriptions, see the <a href="#">Scope Parameters</a> and <a href="#">Destinations</a> tables.</p>

Subcommand	Description
/vaultstorage [ <i>destination</i> ]	<p>Runs a Vault Storage report. The report shows the amount of vault utilization. For more information, see <a href="#">Create a Vault Storage Report</a>.</p> <p>For <i>destination</i> descriptions, see the <a href="#">Destinations</a> table.</p>

## Scope Parameters

*Note:* If you do not specify a *scopeParameter*, the vvanalyze subcommand runs for the entire vault.

ScopeParameter	Description
/Customer= <i>id</i>	<p>Runs the report for the customer with the specified <i>id</i>.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each customer. To determine the <i>id</i> for a customer, you can run a report for the customer from the Director UI and view the customer <i>id</i> in the command line in the resulting report log file.</p>
/Location= <i>id</i>	<p>Runs the report for the location with the specified <i>id</i>.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each location. To determine the <i>id</i> for a location, you can run a report for the location from the Director UI. The location <i>id</i> appears in the command line in the resulting report log file.</p>
/Computer= <i>id</i>	<p>Runs the report for the computer with the specified <i>id</i>.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each computer. To determine the <i>id</i> for a computer, you can run a report for the computer from the Director UI. The computer <i>id</i> appears in the command line in the resulting report log file.</p>
/Task= <i>id</i>	<p>Runs the report for the task with the specified <i>id</i>. To run the report for all tasks, enter * as the <i>id</i>.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each task. To determine the <i>id</i> for a task, you can view its properties or run a report for the task in the Director UI.</p>

## Destinations

*Note:* If you do not specify a destination for a report, the report is saved as a log file.

Destination	Description
/logfile	Saves the report as a log file.

Destination	Description
/email= <i>user@address</i>	Emails the report to the specified <i>user@address</i> and saves the report as a log file.  <i>Note:</i> Before reports can be sent by email, an SMTP server must be specified for the vault. For more information, see <a href="#">Configure email notifications</a> .
/output= <i>filename</i>	Saves the report as a text file with the specified <i>filename</i> , where <i>filename</i> is the path and filename for the output file.

### vvanalyz /missedbackups Parameters

In addition to the parameters listed in the ScopeParameters and Destinations tables, the vvanalyz /missedbackups subcommand can include the following parameters:

vvanalyz /missedbackups Parameter	Description
/bare	Creates the Missed Backups report in CSV format. If you do not include the /bare parameter, report data appears in columns.
/timeframe= <i>hours</i>	Specifies the number of hours for the Missed Backups report timeframe. A backup task that is not scheduled in Portal appears in the Missed Backups report if a safeset exists for the task from before the report timeframe, and no safeset exists for the task during the timeframe.  <i>Note:</i> The timeframe does not affect whether scheduled tasks appear in the Missed Backups report.  The report timeframe is measured back from the current local time on the vault when the report runs. For example, if you specify 24 hours for the report timeframe and the report runs on November 21, 2021 at 10:00, the report includes unscheduled backup tasks with safesets dated before November 20, 2021 at 10:00 and no safesets from between November 20, 2021 at 10:00 and November 21, 2021 at 10:00.
/sendtoadmin	Sends the Missed Backups report to the administrator.
/detailed	Specifies that the Missed Backups report should show errors for failed backups.

<b>vvanalyz /missedbackups Parameter</b>	<b>Description</b>
<code>/sendtocustomer</code> <code>[/header=<i>headerFile</i>]</code> <code>[/footer=<i>footerFile</i>]</code>	Sends the Missed Backups report to the customer email address. To specify a header file for the report, include the <code>/header=<i>headerFile</i></code> parameter, where <i>headerFile</i> is the path and filename for the header file. To specify a footer file for the report, include the <code>/footer=<i>footerFile</i></code> parameter, where <i>footerFile</i> is the path and filename for the footer file.  <i>Note:</i> Before you can email the report, an email address must be specified for the customer and an SMTP server must be specified for the vault. For more information, see <a href="#">Configure email notifications</a> .

### vvanalyz /storageextract Parameters

The vvanalyz /storageextract subcommand can include the following parameters:

<b>vvanalyz /storageextract Parameter</b>	<b>Description</b>
<code>/billing</code>	Creates CSV files in the \Director\bill directory, but does not include report data in the Director\logs directory.
<code>/analyzedisabledtask</code>	Includes disabled tasks in the report.
<code>/url</code>	Includes Customer URLs in the CSV files.
<code>/nodata</code>	Includes computers with no tasks in the pool information
<code>/header</code>	Includes field names at the top of the CSV files.
<code>/native</code>	Includes native storage amounts in the CSV files. The native storage amount is the amount of original client data that is protected by a safeset.
<code>/vaultname</code>	Includes the vault name in the CSV files.

### Examples

The following command runs a Late Server Status report that lists all tasks that were not backed up in the 24 hours before the current time:

```
vvanalyz /lateststatus /timeframe=24
```

The following command runs a Missed Backups report with a 72-hour timeframe in CSV format, and emails the report to user@address.com:

```
vvanalyz /missedbackups /timeframe=72 /bare /sendtoadmin  
/email=user@email.com
```



The following command runs a Missed Backups report with a 48-hour timeframe, and sends the report, with the specified header, to the customer e-mail address:

```
vvanalyz /missedbackups /sendtocustomer /header=c:\headerfile.txt
```

The following command runs a Storage report for Location 1, and emails the report in CSV format to user@address.com:

```
vvanalyz /storage /location=1 /email=user@address.com /bare
```

The following command runs a Storage report for Customer 3, and saves the report in CSV format in a log file:

```
vvanalyz /storage /customer=3 /bare
```

The following command runs a Storage Location report for Computer 1:

```
vvanalyz /storageLocation /computer=1
```

The following command runs a Last Backup Status report for Task 1, and saves the report in a log file:

```
vvanalyz /syncscan /task=1
```

The following command runs a Last Backup Status report for Location 3, and emails the report to user@address.com:

```
vvanalyz /syncscan /location=3 /email=user@address.com
```

The following command runs a Vault Storage report, and emails the report to user@address.com:

```
vvanalyz /vaultstorage /email=user@address.com
```

The following command reports pool type and pool data file sizes for a vault in CSV and log files, and includes Customer URLs, field names, and native storage amounts in the CSV files:

```
vvanalyz /storageextract /poolsize /url /header /native
```

The following command reports pool type and pool data file sizes for a vault in CSV files, but does not include pool information in the log file:

```
vvanalyz /storageextract /poolsize /billing
```

## 14.9 vvarchive

Archives safesets from the primary pool.

### Syntax

```
vvarchive [ taskid | -disk ]
```

If you do not specify a taskid, all safesets are archived.

If you do not specify a taskid, but include -disk, only the tasks that are set to archive to disk are archived.

### Example

The following command archives all safesets from the primary pool:

vvarchive

## 14.10 vvexport

Exports safesets from a vault to a disk media. The media can then be transported to another system, to be restored using an Agent or imported into another vault.

### Syntax

```
vvexport taskid safeset(s) {devicename | /nobackupdata}
[/maxfilesize=sizeInMegabytes] [/catalog={4|5} [/catalogpath=location]]
vvexport path safeset(s) {devicename | /nobackupdata}
[/maxfilesize=sizeInMegabytes] [/catalog={4|5}
[/catalogpath=location]] [/useDefaultForSrcUser
| [/srcUser=<domain/username> /srcPassword=<encryptedPassword>]]
[/useDefaultForDstUser | [/destUser=<domain/username>
/destPassword=<encryptedPassword>]] [/catUser=<domain/username>
/catPassword=<encryptedPassword>] [/simulate]
```

### Parameters

Parameter	Description
taskid	Unique internal task number. Found in the logs for that task's backup.
safeset(s)	Number in online list of safeset(s).
devicename	Local disk path, or UNC path to remote disk for safesets (i.e., \\<Server>\<Share>\<path>).  <i>Note:</i> You can export safesets to a local disk or a Universal Naming Convention (UNC) device.
nobackupdata	Exports catalogs and not safesets.
maxfilesize	Specifies the maximum SSI file size, in megabytes. If the SSI file is larger in total than this size, it is split into multiple SSI files. You can have multiple SSI files on a CD or DVD, or have a single SSI file that spans multiple CD or DVD disks.
catalog	Catalog file version. The value is 5 for Agent versions 3.2 and later.
catalogpath	Path on disk to which catalogs are exported.
simulate	Allows you to view what would be generated, and the size required for the output storage. The simulate option writes the output to the VVExport log file.

### Examples

```
vvexport 1 4 c: /maxfilesize=1024 /catalog=5 /catalogpath=c:\Documents
and Settings
```

## 14.11 vvimport

Imports a safeset from disk into the pool system in primary storage. Use for seeding or re-seeding. The safeset is stored in compressed SSI format, and named after the safeset being stored. You can use a UNC path to specify non-local disk.

### Syntax

```
vvimport taskid synchnum sourcepath [/useDefaultCredentials |
  [/user=<username> /password=<password>]]
```

### Parameters

Parameter	Description
Tasked	Unique internal task number. Refer to the task for this safeset.
Synchnum	Safeset number.
Sourcepath	Refers to local or remote (via UNC path) disk.

## 14.12 vvmigrat

Migrates safesets in accordance with retention policies.

### Syntax

```
vvmigrat {online|offline|genoptidxv5|initsync} {taskID|last}
  [noninterruptible]
```

Subcommand	Description
Online	Determines which safesets in primary storage should be kept based on the safeset retention settings and retention groups. Safesets that are not kept are marked for deletion. For more information, see <a href="#">Enforce retention settings in primary storage</a> .
Offline	Determines which safesets in archive storage, recalled secondary storage, or detached secondary storage should be kept based on the safeset retention settings and retention groups. Safesets that are not kept are marked for deletion.
genoptidxv5	Generates index files to optimize searches within the pool system.
Initsync	Recreates synch.dat or synch.mir depending on which is missing.
taskID	vvmigrat will be run on the task with this task identification.
Last	Works in conjunction with genoptidxv5. Generates optimize index for the last safeset.

Noninterruptible	Sets the vvmigrat command to have the highest priority.
------------------	---

or

```
vvmigrat delete {taskID} {safeset#} [{safeset#},,,]
```

Subcommand	Description
taskID	Task identification on which to delete.
safeset#	Safeset number to delete.
{safeset#},,,	Safeset numbers to delete. Use a comma separator for each number.

or

```
vvmigrat param=<parameters_file>
```

Subcommand	Description
param=<parameters_file>	vvmigrat param=<parameters_file> is only used for deleting objects.

**Important:** Use the vvmigrat param=<parameters\_file> command with caution. Running the vvmigrat command with the parameters file will completely delete the specified customer from the vault even only one job is specified in the parameters file.

## Remarks

The parameters file has the following format:

```
=====
```

```
CONFIG_VERSION=3
CONTENT_VERSION=1
# Command can be one of: DELETE
COMMAND=DELETE
# Scope selection:
CUSTOMER=1,2,3 # Process these customers
LOCATION=4,3 # Process these customer locations
COMPUTER=49,60 # Process these computers
ACCOUNT=5,6 # Process these accounts
USER=11,12 # Process these users
BILLINGCODE=250,139,3 # Process these billing codes
# Tasks can be specified without safesets as follows:
TASK=25,26,27,28 # Process these tasks
# OR you can specify the tasks with/without safesets as follows:
TASK_1=25 # Process this task
TASK_2=28 # Process this task
TASK_3=23:2,4,7,8 # Explicitly process safesets 2,4,7,8 only in task 23
# ... More tasks with/without safesets go here ...
```

## 14.13 vvmove

The vvmove command moves files between storage locations in the primary storage group. For example, you can use the vvmove command to move data from a storage location that you want to retire, or move index files to a faster location. Unless you specify a destination for files, the system moves files to the storage location with the most available space.

**Warning:** Do not use the vvmove command to move pool files to the same physical storage location with a new UNC path, or your pool system will be corrupted.

**Note:** In Director versions earlier than 8.4x, a vvmove operation failed if it tried to move a file to a storage location where there was a file with the same name for the same task. A vvmove operation now continues if a duplicate file name is detected, and does one of the following:

- If the files have matching CRC signatures, vvmove does not replace the existing file. Messages such as the following appear in the vvmove log:

```
19-Jul-2018 11:27:50.759 -04:00 [15872] POOL-I-0001 Copying file:
C:\Vault438151955\COMPUTER-2012_4\testVVMove\poolfiles.dat
19-Jul-2018 11:27:50.775 -04:00 [15872] POOL-W-0002 Will not copy file because it
already exists: D:\Vault438151955\COMPUTER-2012_4\testVVMove\poolfiles.dat
19-Jul-2018 11:27:50.775 -04:00 [15872] POOL-I-0001 new location
D:\Vault438151955\COMPUTER-2012_4\testVVMove\poolfiles.dat
```

- If the files have different CRC signatures, vvmove chooses the file that is referenced by the pool system and saves the other file version with the “.duplicate” extension. Messages such as the following appear in the vvmove log:

```
19-Jul-2018 11:28:06.562 -04:00 [15872] POOL-W-0002 Duplicate file
{C:\Vault438151955\COMPUTER-2012_4\testVVMove\differentFile} exists, renaming it to
{C:\Vault438151955\COMPUTER-2012_4\testVVMove\differentFile.2018-07-18 15-44-
52.duplicate}
19-Jul-2018 11:28:06.578 -04:00 [15872] POOL-I-0001 Copying file:
C:\Vault438151955\COMPUTER-2012_4\testVVMove\differentFile.2018-07-18 15-44-
52.duplicate
19-Jul-2018 11:28:06.593 -04:00 [15872] POOL-I-0001 new location
D:\Vault438151955\COMPUTER-2012_4\testVVMove\differentFile.2018-07-18 15-44-
52.duplicate
```

If duplicate files are saved after a vvmove operation, we recommend running a task verification. Files with the .duplicate extension are removed by subsequent cleanup operations.

### Syntax

```
vvmove subcommand [filter1... filtern]
```

### Subcommands

Subcommand	Description
<code>/removeraidpath=<i>physicalPath</i></code>	Moves files from the specified storage location to other locations in the primary storage group. <i>physicalPath</i> is the path of the location from which you want to move files.

Subcommand	Description
<code>/consolidatetopath=<i>physicalPath</i></code>	Moves files to the specified storage location from other locations in the primary storage group. <i>physicalPath</i> is the path of the location to which you want to move files.
<code>/applystoragepolicy</code> <code>[/filetype=<i>filetype</i>]</code>	Moves files to preferred storage locations, as specified by the location storage policies. <i>filetype</i> can have the following values: <ul style="list-style-type: none"> <li>• <code>idx</code> — Moves index files to their preferred locations.</li> <li>• <code>log</code> — Moves log files to their preferred locations.</li> <li>• <code>pool</code> — Moves pool data files to their preferred locations.</li> </ul> You can include one or more <code>/filetype</code> parameters with the <code>/applystoragepolicy</code> subcommand. If you do not include a <code>/filetype</code> parameter, files with all types are moved to the preferred locations.
<code>/retiregroup=<i>storageGroupID</i></code> <code>/username=<i>username</i></code>	Moves files from the specified storage group to locations in another storage group, and marks all storage locations in the storage group as read-only. <i>storageGroupID</i> is the name of the primary storage group. <i>username</i> is an account that can connect to the vault. The account can be a domain or local account that can access to the vault machine.

## Parameters

Parameter	Description
<code><i>filter1... filtern</i></code>	Specifies which files to move. You can include one or more of the following filters: <ul style="list-style-type: none"> <li>• <code>/customer=<i>id</i></code> — Moves files for the customer with the specified id.</li> <li>• <code>/customerlocation=<i>id</i></code> — Moves files for the customer location with the specified id.</li> <li>• <code>/computer=<i>id</i></code> — Moves files for the computer with the specified id.</li> <li>• <code>/task=<i>id</i></code> — Moves files for the task with the specified id.</li> </ul> <p><i>Note:</i> You cannot include <code>/customer</code>, <code>/customerlocation</code>, <code>/computer</code> and <code>/task</code> filters in the same <code>vvmove</code> command. For example, you cannot include both a <code>/customer</code> and a <code>/computer</code> filter in a <code>vvmove</code> command.</p> <ul style="list-style-type: none"> <li>• <code>/safeset=<i>id</i></code> — Moves files in the safeset with the specified id.</li> </ul> <p><i>Note:</i> You can include only one <code>/safeset</code> filter in a <code>vvmove</code> command.</p>

	<ul style="list-style-type: none"><li>• <code>/srcraidpath=path</code> — For use with the <code>/removeraidpath</code> and <code>/consolidatetopath</code> subcommands only. Moves data with a minimum of one file in the specified path.</li></ul> <p>If you do not include a filter, the system moves all files.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each customer, customer location, computer, task and safeset. To determine the <i>id</i> for a customer, customer location, computer, task or safeset, you can run a report for the item. You can also determine the <i>id</i> of a task by viewing its properties in the Director UI.</p>
--	--

## Examples

The following command moves files from the `s:\local\index` and `s:\local\log` directories:

```
vvmove /removeraidpath="s:\local\index" /removeraidpath="s:\local\log"
```

The following command moves index and log files to their preferred storage locations:

```
vvmove /applystoragepolicy /filetype=idx /filetype=log
```

The following command moves index, pool, and log files for customer 3 and customer 5 to the preferred storage locations:

```
vvmove /applystoragepolicy/customer=3 /customer=5
```

## 14.14 vvpoolop

The `vvpoolop` command performs the following operations in primary storage pools:

- checks physical and logical data integrity
- deduplicates data
- optimizes pools
- re-enables suspect tasks

You can specify the storage pool for an operation using a selection parameter in the command.

You can also run Storage Pool Summary reports for both primary and secondary storage pools using the `vvpoolop` command.

### Syntax

```
vvpoolop subcommand selection [parameters]
```

*Note:* A *selection* must appear after the *subcommand* and before any *parameters*.

## Subcommands and Parameters

Subcommand and Parameters	Description
checkcrc [light]	<p>Checks the pool system for physical corruption.</p> <p>To only check for the presence of files, include the light parameter.</p> <p><i>Note:</i> If vvpoolop checkcrc light determines that a file is missing from the pool system, the task is not marked as suspect. Instead, log messages indicate that a file is not found and that the check failed.</p>
checksafesetpf	Checks single safeset pool files.
checkallsafesetpf	Checks all pool files for safesets.
dedup	Eliminates references to repeating data so the optimize subcommand can purge the data. For use with Version 5 pools only.
listfixed	Lists fixed files (i.e., files with no streams).
optimize [dedup] [nointerrupt] [timeout <i>minutes</i> ]	<p>Optimizes the pool system.</p> <p>To deduplicate data when optimizing the pool system, include the dedup parameter.</p> <p>To prevent incoming backups from cancelling existing optimization tasks, include the nointerrupt parameter.</p> <p>To unlock the tasks and exit after the specified number of minutes have elapsed, include the timeout <i>minutes</i> parameter, where <i>minutes</i> is the number of minutes.</p> <p><i>Note:</i> The timeout <i>minutes</i> parameter is ignored for maintenance jobs.</p>
ssiverone [ <i>id</i> /last]	<p>Checks a safeset for logical corruption: either the safeset with the specified <i>id</i>, or the last committed safeset. This verification process is more reliable but slower than the verifyidxone process.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each safeset. To determine the <i>id</i> for a safeset, you can view its properties in the Director UI.</p>
ssiverall	Checks all online safesets for logical corruption. This verification process is more reliable but slower than the verifyidxall process.



Subcommand and Parameters	Description
summary [detailed]	<p>Runs a Storage Pool Summary report for a primary storage pool system. The report provides information about the amount of disk space allocated to one task, or to each task in a vault. For more information, see <a href="#">Create a Storage Pool Summary Report</a>.</p> <p>To include information separately for each pool file in the report, include the detailed parameter.</p>
summaryall [detailed]	<p>Runs a Storage pool Summary report for both primary and secondary pool systems. The report provides information about the amount of disk space allocated to one task, or to each task in a vault. For more information, see <a href="#">Create a Storage Pool Summary Report</a>.</p> <p>To include information separately for each pool file in the report, include the detailed parameter.</p>
unsuspect	Checks the logical and physical integrity of data for a suspect task, and enables the task if no data issues are found. If data integrity issues are found, the system does not enable the task.
verifyidzone [ <i>id</i> /last]	<p>Checks a safeset for logical corruption: either the safeset with the specified <i>id</i>, or the last committed safeset. This verification process is less reliable but faster than the <i>ssiverone</i> process. Unlike <i>ssiverone</i>, this process does not read the pool files.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each safeset. To determine the <i>id</i> for a safeset, you can view its properties in the Director UI.</p>
verifyidxall	Checks all online safesets for logical corruption. This verification process is less reliable but faster than the <i>ssiverall</i> process. Unlike <i>ssiverall</i> , this process does not need pool files.

## Selections

Subcommand	Description
task { <i>id</i> /*}	<p>Performs the operation for the task with the specified <i>id</i>. To perform the operation for all tasks, enter * as the <i>id</i>.</p> <p><i>Note:</i> The system assigns an identification number (<i>id</i>) to each task. To determine the <i>id</i> for a task, you can run a report for the task or view its properties in the Director UI.</p>
location <i>physicalLocation</i>	Performs the operation for pool systems that have at least one file in the specified <i>physicalLocation</i> . The <i>physicalLocation</i> can be a local or a UNC path.

## Examples

The following command performs a CRC in primary storage for task 2:

```
vvpoolop checkcrc task 2
```

The following command optimizes and deduplicates data for all tasks:

```
vvpoolop optimize task * dedup
```

The following command deduplicates data in primary storage for task 3:

```
vvpoolop dedup task 3
```

The following command runs a Storage Pool Summary report for all primary pool systems with files in c:\Vault184170841:

```
vvpoolop summary location c:\Vault184170841
```

The following command runs a Storage Pool Summary report for both primary and secondary pool systems for all tasks:

```
vvpoolop summaryall task * detailed
```

The following command checks for logical corruption in the last committed safeset for task 3:

```
vvpoolop ssiverone task 3
```

The following command checks for logical corruption in all safesets for pool systems with files in \\computer1:

```
vvpoolop ssiverall location \\computer1
```

The following command checks for logical and physical corruption in suspect task pool systems, and enables suspect tasks that do not have data integrity issues:

```
vvpoolop unsuspect task *
```

The following command checks safeset 00000002 for task 4 for logical corruption without reading pool files:

```
vvpoolop verifyidxone task 4 00000002
```

## 14.15 vvpurge

The vvpurge command deletes log files or activity records that are older than a specified amount of time. When deleting log files, you can also specify a minimum number of each log file type (e.g., Backup log, Restore log) to keep.

Log files are typically deleted automatically, as specified by Vault Settings. Activity records are not deleted automatically and are saved in the vault database. Activity records list completion times, file sizes and success messages for backup, restore and migrate tasks.

### Syntax for Deleting Log Files

```
vvpurge [numberOfFilesToKeep numberOfDays]
```

### Syntax for Deleting Activity Records

```
vvpurge /activities /delete /age=number{h|d|y}
```

## Parameters

Parameter	Description
<i>[numberOfFilesToKeep numberOfDays]</i>	<p>For deleting log files only. <i>numberOfFilesToKeep</i> specifies the minimum number of each type of log file to keep. <i>numberOfDays</i> specifies the number of days old that a log file must be to be deleted.</p> <p>Each day represents a 24-hour period. For example, if <i>numberOfDays</i>=1, a log must be at least 24 hours old for the log to be deleted.</p> <p>If you do not include these parameters, the system deletes log files as specified by values on the Logs tab of the Vault Settings dialog box. For more information, see <a href="#">Schedule a log file purge</a>.</p>
<i>/age=number{h d y}</i>	<p>For deleting activity records only. Specifies the <i>number</i> of hours, days or years old that an activity record must be to be deleted. Include one of the following time units:</p> <ul style="list-style-type: none"> <li>• h — hours</li> <li>• d — days</li> <li>• y — years</li> </ul>

## Examples

The following command deletes log files that are older than 180 days, but keeps a minimum of 31 of each type of log file:

```
vvpurge 31 180
```

The following command deletes log files as specified by values on the Logs tab of the Vault Settings dialog box:

```
vvpurge
```

The following command deletes activity records that are more than one year old:

```
vvpurge /activities /delete /age=1y
```

The following command deletes activity records that are more than two days old:

```
vvpurge /activities /delete /age=2d
```

The following command deletes activity records that are more than four hours old:

```
vvpurge /activities /delete /age=4h
```

## 14.16 vvrecall

Recalls an archived safeset.

### Syntax

```
vvrecall [taskID] [catalog] [storagedays]
```

## 15 Carbonite Server Backup Support

If you have a question about Carbonite Server Backup that isn't covered in this guide, our frequently-updated Knowledge Base contains comprehensive information. The Knowledge Base is your first stop when searching for any Carbonite Server Backup solutions you may need. We highly recommend searching here first for the quickest answers to your questions.

**Knowledge Base:** [support.carbonite.com/evault](https://support.carbonite.com/evault)


### What can we help you with?

Popular Searches  
[pending reboot](#), [restore](#), [clnt-e-04103](#)

### 15.1 Contacting Carbonite

If you need live assistance from a qualified support agent, Carbonite Support is here for you 24 hours a day, 7 days a week (excluding US holidays). Please feel free to get in touch with us, and we'll help out any way we can! You can find the contact information for Carbonite Support in the Knowledge Base:

[support.carbonite.com/evault](https://support.carbonite.com/evault)



**Tip:** When contacting Support with a technical issue, please have both the program's log files and the store you are having difficulty with ready.

To gather log files, click **File** menu and choose *Open log folder*. Compress the contents of the folder in a .zip file and attach it to your support request.

If the log archive and/or mail store exceeds 10MB, you may not be able to send them as an email attachment. In that case, upload instructions will be provided to you upon request.